

A SYSTEM FOR FAIRNESS IN SENTENCING

KENNETH HALPERN

February 7, 2012

CONTENTS

1. Abstract	1
2. Introduction	1
2.1. Motivation	2
2.2. Forms of Unfairness	3
2.3. Why Fairness Matters	4
3. The Ingredients of Our Method	5
3.1. Observer	5
3.2. Sentence	5
3.3. Case	6
3.4. Ordering	7
3.5. History and Paradigm	8
4. Formal Framework	9
4.1. The Sentencing Function	10
5. Our Procedure	11
5.1. Insertion	11
5.2. Procedure for Sorting Sentences	12
5.3. Coverage	12
5.4. Degeneracy of S	13
5.5. The Purpose of Anchor Points	13
5.6. Summary of the Procedure	14
6. Finite vs Infinite Representations	15
7. Aggregation	16
7.1. Decision Point Aggregation	17
7.2. Ordering Aggregation	18
8. Metric	18
8.1. A Class of Swap Metrics	19
8.2. Comparison of Case Orderings	20
8.3. Comparison of Sentence Assignments	20
8.4. Comparison of Sentencing Histories	21
9. Dependence on Sequence of Presentation	21
10. Algorithmic Observers	22
11. Main Conclusions	23
Appendix A. Review of Mathematical Orderings	24
References	25

1. ABSTRACT

We ask how judicial sentences can be dispensed in a manner that is consistent with a given individual's sense of fairness. After discussing the importance of this question, we propose a simple mathematical solution. We then examine various metrics for comparing sentencing assignments by different individuals as well as practical issues with the implementation of our methodology¹.

2. INTRODUCTION

Dostoevsky famously noted that “The degree of civilization in a society can be judged by entering its prisons.” Perhaps an equally important measure is its consistency in the administration of justice. Do materially identical cases result in the same verdict and punishment? In practice, the same case could result in a broad array of outcomes depending on many factors. The greater the prevalence

¹We thank Professor Eugene Volokh for his helpful comments on this piece.

of such inconsistencies, the greater the apparent unfairness of the judicial system. This does not necessarily mean that the system *is* unfair² or undesirable³.

We do not examine all components of the judicial process, focusing instead on the assignment of sentences. Our purpose is to offer a specific means of minimizing the perceived unfairness of sentencing, subject to a number of assumptions. Though our discussion is motivated by and centered on the present criminal justice system of the United States, its application is general.

The method we shall describe may appear similar to certain real systems, such as the Federal Sentencing Guidelines. The latter uses a base penalty range for each crime, with fixed adjustments for aggravating or mitigating factors. To preemptively clarify the distinction, we note that our approach requires only the ability of a judge to compare pairs of cases. A discretization of cases into sentencing classes is unnecessary, though helpful. Where present, it depends on the judge rather than statutory guidelines and will be far more refined than the coarse reliance on crime category employed by the Federal Sentencing Guidelines. If deterministic sentencing arose, it would do so as a natural consequence of deduction; once the judge has established a sufficient body of precedents, future decisions would be precisely constrained. The nature of these differences will be apparent as we present the method in detail.

In sections 2.1 we provide some background and motivation, discuss the sources and consequences of unfairness, and argue for the importance of its redress. Sections 3 - 5 describe our method for assigning sentences in a manner consistent with a given individual's sense of fairness. We consider various aspects of this system including possible convergence to deterministic behavior (section 6) and questions of robustness (section 9). We then discuss the aggregation of individual judgments (section 7), possible metrics for comparing such judgments (section 8), and the construction of an implied judge whose sense of fairness most closely resembles those of the real world (section 10). Finally, in section 11 we provide brief suggestions for practical use of our method. For the reader interested only in a summary, sections 5.6 and 11 should serve. A brief review of mathematical orderings is provided in appendix A.

2.1. Motivation. It is not uncommon for us to hear of criminal sentences that offend our individual senses of fairness. This is true even if there is no doubt in our minds of the defendant's guilt⁴. Two equivalent crimes result in radically different sentences, or a mild transgression draws a greater penalty than a severe offense, or a sentence seems vastly disproportionate to the crime. We shall generically refer to such instances as "unfairness", with the understanding that there is an implicit observer by whose standards they are deemed unfair⁵. The prevalence and degree of unfairness vary between observers. It is reasonable to suppose that, given any individual, there exists a set of real cases – large or small – that he would deem unfair.

It may be objected that most of us only are privy to highly publicized cases, that our knowledge of such cases is informed by the popular media, and that consequently our view of judicial fairness as lay observers is subject to a significant selection bias toward sensational stories or those which arouse a strong sense of indignation. This would be a valid concern were such knowledge the basis for our analysis rather than simply a motivation. While the extent of such unfairness may reflect on the utility of our approach it does not determine its validity or necessity. The ubiquity of such instances may not be constant with time. There is no telling when they could proliferate, and the

²For example, we tend to presuppose that there is an absolute truth in each case, that it is determinable, and that it will be arrived at unerringly by any jury. On reflection, there are many reasons why this is false. However they do not change our underlying perception. Similarly, a random disposition would appear capricious and offend our sense of justice even if it statistically treated everyone as equal.

³There are many theories of justice, and not all value the individual to the same degree.

⁴Of course, this does not mean that he or she actually *is* guilty. As mentioned, our focus is on the sentencing phase. We do not address questions of guilt or innocence here.

⁵Note that our use of the term "unfairness" is confined to sentencing.

present judicial system is sufficiently complex that a small legal or administrative change could lead to a rapid growth in such cases. Moreover there are strong reasons to believe that, by most of our standards, instances of unfairness are rife⁶.

We won't delve too deeply into the myriad sources of unfairness. A few obvious culprits suffice to illustrate their nature and probable frequency.

- **Specific Variations:** The most basic source is simple variation between circumstances and individuals. It also is the most difficult cause to adjust for, particularly if we require determinism. As will be discussed shortly, one requirement for fairness is that irrelevant factors play no role in the outcome. In reality there are many such factors which do, including countless aspects of the people involved in all facets of the case.
- **Variation in Jurisdiction:** Another important source of unfairness, at least in the author's view, is the large number of legal and enforcement frameworks. These include multiple levels and countless geographies. What is considered a major crime in one area may be deemed lawful or of minor import in another⁷.
- **Prosecutorial Power:** The recent transition of power from judges and juries to prosecutors furnishes another source of unfairness. Legislative expansion of the number of criminal offenses and a hardening of penalties across the board give prosecutors great leverage to coerce plea bargains. A commonly cited Bureau of Justice statistic is that over 90% of cases in the United States [1] are resolved through plea bargains⁸. If a plea bargain does not fix a precise sentence, it certainly affects the applicable range of sentences. Plea bargains depend on the precise personalities involved, otherwise inconsequential details of the case, and the specific laws that can be used as leverage. The net effect of all three of these phenomena is that many cases result in unfair punishments.

Neither of the extremes of full discretion or mandatory sentencing alleviate the problem of unfairness. The former allows corruption and wide individual variation while the latter fails to distinguish between material differences in cases and generates sentences without benefit of common sense. Both sources of unfairness are problematic, and the balance between them is a timeless question of social structure.

2.2. Forms of Unfairness. In order for a judicial framework to appear fair, it should possess certain properties. Some of these are universal and some vary with the observer. Most could equally well apply to the entire judicial process, though we only focus on the sentencing component. As mentioned, we do not consider whether these requirements are necessities for various general judicial approaches. Rather, they reflect the author's – and, we believe, most people's – intuitive sense of fairness.

- (1) **Consistency under repetition:** The same case must result in the same outcome with little variation when tried repeatedly in the same venue, possibly by different individuals. This is a very basic form of fairness; it's absence leads to a sense that the system is arbitrary.
- (2) **Independence from immaterial factors:** Two cases which differ only in details that we would deem irrelevant to sentencing should be considered materially similar and result in similar outcomes⁹.

⁶It also could be argued that the frequency of unfair outcomes should be higher in the general population than for persons worthy of media attention. Consequently, our bias would underestimate rather than overestimate the prevalence of such cases. We present no evidence either way.

⁷The possession of drugs and guns furnish two current examples.

⁸It is not our purpose here to comment on this fundamental change in the nature of our judicial system, but simply to mention it as a major source of perceived unfairness.

⁹What we consider "immaterial" depends on the observer. Note that factors which could affect the verdict may be deemed irrelevant to sentencing and vice versa. Perhaps whether a man has a family may affect his believability or likelihood to

- (3) Independence from jurisdiction: The same case should result in the same outcome independent of where it is tried or by what authority. This requires uniformity of law and its application¹⁰. It really is a form of independence from immaterial factors where jurisdiction, in the view of the author, should be an immaterial factor.
- (4) Relative consistency: Given two cases, the greater offense should get the greater penalty. Materially similar cases should result in similar sentences.
- (5) Proportionality: No punishment should be significantly disproportionate to the offense¹¹.

The first three requirements are general, though not everyone would agree on them. The latter two depend on the specific observer. The assessment of whether a punishment is “disproportionate” or whether one offense is worse than another varies between individuals. The nature of the perceived unfairness has great bearing on its effect and possible remedy. Relative consistency and proportionality reflect on the body of law and require legislative solutions. They are not failings of the judicial system per se. In a democracy their remedy involves lobbying for legal reform by the individual whom they offend. Also, what one person considers excessive another may deem lenient¹². However a lack of consistency under repetition makes the judicial system seem arbitrary and incompetent, while lack of independence from immaterial factors implies either similar incompetence or – worse – various forms of prejudice. A lack of independence from venue is something that we are accustomed to, but can likewise lead to a sense of arbitrariness and therefore distrust of the legal system¹³.

Our proposed system directly addresses all of the above sources of unfairness except the issue of proportionality, though this may be indirectly remedied as well if the spectrum of penalties is sufficiently broad and the set of prior cases sufficiently large.

2.3. Why Fairness Matters. There are a number of reasons we should be concerned about cases that violate our sense of fairness. As individuals, we may feel that we share moral culpability for such outcomes even if we do not directly influence them. From a broader standpoint, our faith in our government depends on a sense of consistency and fairness¹⁴. Among the factors that form this perception is our view of the judicial process. After all, “due process” really means “due and fair process” to most of us. When we lose faith in our judicial system, an important cornerstone of respect for our government is removed. Depending on the political and economic environment, this may feed perceptions of inequality based on wealth or race or other characteristics.

Even if the more insidious forms of inconsistency were completely eliminated, there is no sentencing system which everyone would deem fair. Nor need such fairness even be the goal of a penal framework. For example it is perfectly reasonable to adopt a “broad brush” legal policy that harms

commit certain crimes, but should have no bearing on his sentence *when* guilty (all else being equal). And perhaps a man who hits another man unprovoked may deserve a different sentence than one who hits a man who is harassing a woman.

¹⁰This does not mean that offenses need be independent of geography. For example, it is perfectly reasonable for certain actions to be prohibited in cities or in parks. That is part of the definition of the offense itself. However, a sentence of ten years in one State and life in another for the same offense would violate this tenet of fairness.

¹¹As mentioned, this is a requirement for a sense of fairness by common current standards. There are judicial philosophies for which it is not a desirable constraint.

¹²We need not go far to find individuals whose real sense of justice radically differs from our own. A simple example (for most of us) would be Draco, from whose name the word ‘draconian’ derives. He established a written code of law for Athens that prescribed the death penalty for the majority of offenses. When asked why, he replied to the effect that the lesser ones deserved death and he had no stronger punishment for the greater ones (Plutarch, Life of Pericles). That others did not share this sentiment was established when Pericles softened those laws, to popular acclaim.

¹³It is particularly problematic when the various applicable statutory frameworks are so massive that no individual can reasonably be expected to know the laws that apply where he stands at a given moment. That is indeed a frightening state of affairs. There is an excellent discussion of the general phenomenon of excess criminalization in [2].

¹⁴As Sun Tzu noted, one of the ways to tell the relative strengths of adversaries is to determine “In which army is there the greater constancy both in reward and punishment?”. The same could be said of nations.

some number of people disproportionately but has a deterrent effect overall¹⁵. It is not our purpose to address the philosophical questions behind the definitions of “fairness” and “justice” or discuss penal theory in general. We simply examine the possibility of satisfying a single individual’s sense of fairness, were this the goal.

3. THE INGREDIENTS OF OUR METHOD

Our procedure allows an observer to transform his sense of fairness into a means of assigning actual sentences. For an automated observer, the distinction is irrelevant and we simply impose a set of consistency constraints. However for a human observer, our method provides a non-trivial means of consistently administering justice.

There are several components to our system, and we now introduce these before providing a formal construction of the framework in section 4.

3.1. Observer. We define an “observer” to be an individual whose sense of fairness we seek to accommodate. This is accomplished through a series of comparisons on his part. The observer need not be a person; it could be any agent which yields a set of suitable responses consistent with our assumptions. For example, one could employ a machine learning or other algorithm that attempts to construct an “implied” sense of fairness based on existing outcomes. Any non-human observer will either be termed “algorithmic” if purely computerized or “aggregate” if composed of multiple humans. A human or aggregate observer may be assisted by a computer in various ways – such as through a choice of comparisons to make or information regarding consistency constraints – but retains ultimate discretion. Unless the context demands otherwise, we assume a single, fixed observer.

It is important to note that our observer is arbitrary. We make no assumptions regarding his particular sense of fairness, whether it reflects any prevailing social norm, or whether it is remotely compatible with our own sensibilities. Consequently, the range of information he uses¹⁶ may be quite different from what we would consider reasonable or appropriate. We simply examine the means of reflecting his particular standard in sentencing without reference to its nature. To construct observers that reflect broader social values, we must ensure that qualities unacceptable for consideration in sentencing are excluded.

3.2. Sentence. A “sentence” s_i bears the ordinary connotation: a specific penalty that may be applied to an offender. The set of sentences S_{law} consists of all penalties that are available to an observer for consideration. This is independent of the observer and ordinarily would be dictated by the judicial system in question. For example, in the United States at this time we would have a range of prison sentences, fines, death, registration as a sex offender, probation, driving restrictions, residence restrictions, community service, and a variety of less common punishments. However, mutilation, whipping, enslavement, and exile are not considered acceptable penalties in our society, though they were in many other times and places. We do not focus on the merits or moral palatability of different penalties, nor do we consider their precise administration. Any controllable variation in harshness – such as the type of prison – is considered part of the sentence, and any uncontrollable component – such as possible guard brutality – may be included in a determination of the relative severity of sentences by the observer.

¹⁵This approach – among others – was embodied in the “Bloody Code” of England in the early 1800’s. By 1815, there were over 200 crimes for which capital punishment was mandatory – including minor theft [3].

¹⁶That is, the information included in the definition of a “case” as discussed below.

We should mention that our use of S in what follows implicitly assumes that a given sentence is of the same severity to whomever it may be applied. In reality this is not true. A middle class offender would find a month in prison far more devastating than would a repeat offender. Similarly, a poor man would find a given fine more onerous than a rich man. There are various ways to incorporate this information, but the easiest is through the use of mitigating factors as part of the case information discussed below¹⁷.

A given observer may choose to only consider a subset $S \subseteq S_{law}$ of penalties that are consistent with his sensibilities.

3.3. Case. A “case” refers to a particular defendant, set of circumstances, and trial in which a penalty phase is present. By the latter, we mean the presence of a punishment whether following from a plea, an administrative action, a judgment, or any other mechanism¹⁸. As such, its form is unconstrained. Any meaningful comparison of cases requires a representation of some sort. This embodies all information relevant to a sentencing decision. For an algorithmic observer the representation must be precisely codified, while for a human observer it could be a heuristic description based on an unambiguous use of common language. There is a trade-off that is required, however. If we include too much information, then the representation grows complex and it would prove difficult to fruitfully compare cases. The more readily comparable cases are, the greater our chances of meaningfully imposing an ordering corresponding to our intuitive sense of fairness. On the other hand, too simple a description may omit information that the observer would deem material, hobbling our ability to distinguish between cases. This could effect a near-automatic mandatory sentencing system. For a given observer, the ideal balance would be a minimal description that includes all material information. As an overly simplistic example, a case could be described by “Good character. Sucker punched a stranger in a bar fight, breaking his nose. Unprovoked. No priors.” as opposed to “A 39 year old guy named John — walked into a bar and had a whiskey and” followed by 40 pages of narrative on one hand or “class 2 assault” on the other.

As a very simplistic mathematical example, a representation could be the set $[0, 1] \otimes [0, 1] \otimes [0, 1]$ – three numbers measuring how much other people were hurt, the defendant’s recidivism, and the social severity of the type of crime – and a map $C \rightarrow R$ that extracts this information from each case. Of course, any real codification would be far more complex.

While the representation is observer-dependent, it may be useful to define a universal representation that contains all information that any observer we may reasonably consider could care about. For example, any real attempt to codify and record cases for use by an unknown set of future observers would require some form of universal representation. Observer-dependent representations could then be obtained as projections (or subsets) of this.

We shall use the term “cases with outcomes” to refer to cases along with their assigned sentences. If the latter have been imposed by an actual court or administrative entity then we shall refer to them as “real cases with outcomes,” otherwise they are the result of another observer’s decisions. It is the purpose of our method to review real cases with outcomes and determine appropriate adjustments to those outcomes. Except in theoretical exercises, all cases that are processed by an observer will be real in the ordinary sense of the word. However for most purposes we do not require knowledge of the assigned sentence. It is only when we do that cases with outcomes are considered. We also distinguish between cases that are used for analysis or learning and those which an observer processes in earnest. The latter result in sentencing adjustments that either are recommended or ordered depending on the role assigned the observer, and we refer to them as “active” cases. Note

¹⁷It would be more intuitive to modify the sentence ordering, but we could only do so only by introducing multiple copies of each sentence, associating them with different types of offender, and then assigning them varying positions in the ordering. This quickly grows far less intuitive than our use of cases as proxies, and it introduces a host of problems.

¹⁸It is irrelevant to us whether human discretion is involved or the sentence is mandatory or agreed upon.

that active cases may be old or new, though it is unlikely that cases whose sentence has long since been served would merit review.

3.4. Ordering. Our method calls for an observer to sort two sets of elements. The observer must perform a one-time ordering of the set of sentences by severity, and he must continually maintain an ordering on the set of processed cases by seriousness of offense. These orderings will be discussed in detail shortly; for now, we simply note that they are of the same type.

In either situation we require that for any two elements (cases or sentences) a and b , one of three relations hold: $a < b$, $a > b$ or $a \sim b$. We use the notation $a \sim b$ for equivalence under the ordering and reserve $a = b$ for set-theoretic equality. The important thing to note is that multiple sentences may be deemed equally severe and multiple cases may be deemed equally serious. A typical ordering (such as those of the integers or real numbers) would require that $a = b$ if $a \sim b$. Such an ordering is called a 'linear' or 'total' order. The properties of this and other orderings we shall discuss are reviewed in Appendix A.

There are three basic ways to order the set in question. The first two use relaxed variants of familiar relations and are entirely equivalent: we could define our ordering using a strict $<$ operation or a weak \leq operation. However we choose to define our ordering, we demand certain basic properties. Some of these are redundant from a mathematical standpoint but we list them all for clarity.

- For any element a , we require $a \leq a$ and $\neg(a < a)$.
- For any elements a and b we demand that $a \leq b$ or $b \leq a$ or both. We likewise demand that one of $a < b$, $b < a$, or $a \sim b$ holds.
- For any elements a , b , and c , we require transitivity. If $a \leq b$ and $b \leq c$ then $a \leq c$. Likewise, if $a < b$ and $b < c$ then $a < c$.

We've mixed two ordering notations and an equivalence (\sim) relation in this list of requirements. This isn't a problem because the two means of defining the ordering are entirely equivalent. An ordering using the $<$ operator and subject to our requirements is called a 'strict weak order', while an ordering using the \leq operator is called a 'total preorder'. Given either operator, we define the other using $a \leq b \Leftrightarrow \neg(b < a)$. The two orderings differ in the way they treat elements that are equivalent. The strict weak order defines equivalent elements to be noncomparable: $a \sim b$ iff $\neg(a < b)$ and $\neg(b < a)$ and a and b are sandwiched by the same elements in the order. On the other hand, the total preorder treats elements as equivalent in a more familiar manner: $a \sim b$ iff both $a \leq b$ and $b \leq a$.

We mentioned that there are *three* ways to order the set. The third is based on the total preorder \leq . Given such a relation on a set, we can define an order on the set of equivalence classes¹⁹. Denoting by $[a]$ the set of all elements b such that $a \sim b$, the induced order is defined by $x < y$ iff $\neg(b \leq a)$ for any (or all) $a \in x$ and $b \in y$. This order is a total order and therefore is more intuitive to manipulate. To summarize, the following are equivalent:

- The strict weak order $<$ on the set.
- The total preorder \leq on the set.
- The combination of \sim on the set to define a partition into equivalence classes and a total order $<$ on the latter.

Before proceeding, we should explain why we chose not to numerically index our set instead of imposing an explicit ordering. We could assign each element of the set a numeric value. This would allow for distinct equivalent elements and rely on a familiar proxy for our ordering. There are several reasons why we do not adopt this approach. First, we wish to impose the minimum

¹⁹We could equally well use the equivalence classes of the strict weak order because they are the same.

set of requirements relevant to our system. Any numerical set that we would consider for use has substantial additional structure. There is a significant danger of inadvertently appropriating this structure for other purposes. For example, we do not have a natural metric or algebra on the sets of sentences or cases. However, if we employed a numerical index then the temptation would be to define a distance by subtracting values. Moreover, our method involves a very basic assumption: that a human can directly compare any two cases. If we employ numbers, then the natural operation is to assign a value rather than compare cases. While this is a heuristic distinction, it is an important one.

3.5. History and Paradigm. Before proceeding, we should clarify an important distinction in terminology. At any given point, the result of our method is the disposition of a particular set of cases. We term this a “Sentencing History” and define it to consist of a set of cases, an ordering of them by seriousness, and an assignment of sentences to some or all of them²⁰.

The Sentencing History contains no reference to the means by which the ordering was imposed or the sentencing decisions arrived at. There is no concern about consistency because we have not alluded to any ordering on the set of sentences. In this sense, the Sentencing History purely is an outcome. Note that from a judicial standpoint, the outcome entirely consists of the sentences assigned to cases. Our sole addition in the definition of the Sentencing History is the requirement of an appropriate ordering on the set of cases.

Our “method”, alluded to above, is the general idea of imposing an ordering on the set of sentences and requiring that the Sentencing History always remain consistent with that ordering. As such, it includes no specific procedure or constraint on the observer or means by which he renders decisions. As long as the Sentencing History is suitably consistent, we are happy.

As discussed above, there are various manners in which an observer may be defined or constructed to use our method. An algorithmic observer has a pre-defined set of instructions for doing so. It may employ machine learning, regression, or statistical algorithms, among others.

A human observer requires a means to translate his sense of fairness into actual decisions regarding cases. We define such a mechanism below and term it our “Procedure”. It is not an algorithm for rendering decisions per se, but rather one for enabling a human observer to do so based on his judgment.

By a “Sentencing Paradigm”, we refer to a mechanism for making future choices. Its decisions must be consistent with, and hence constrained by, the existing Sentencing History and the ordering on sentences²¹. Therefore a Sentencing Paradigm must include these elements as well as the observer itself. For an algorithmic observer, this is sufficient. However for a human one, we must also include our Procedure as well as any additional information it relies on.

The decisions of a human observer are not predictable. In principle, he must always remain present to issue future ones. However it is conceivable that after a sufficiently large Sentencing History has been established, all choice may be constrained away and the Procedure itself could function as an algorithmic observer²². We shall see that this is the situation if the representation set is sufficiently small, but not if we merely have covered the set of sentences²³. In the latter case, the frequency with

²⁰In most situations, all of them would be assigned sentences; but it proves useful to allow greater generality.

²¹Though it is conceivable that a minimal state representation of these could be used, this has no real implication for our method.

²²We shall see that this can only be true to within equivalence classes of sentences; there always is a choice of sentence within those unless a predefined algorithm for selecting one has been adopted.

²³That is, each sentence has one or more cases to which it has been assigned.

which the human observer must get involved in the actual assignment of sentences grows very small but remains finite. He must continue to order the cases, however.

To summarize our terminology:

- **Sentencing History:** A set of cases, ordered by seriousness, and with sentences assigned to some or all of them.
- **Our Method:** The idea of maintaining a Sentencing History consistent with a pre-imposed ordering on the set of sentences.
- **Sentencing Paradigm:** A mechanism for making future decisions (ordering of cases and assignment of sentences) in keeping with our method (i.e. consistent with any existing Sentencing History and ordering on sentences).
- **Our Procedure:** A mechanism we propose for converting a human observer's sense of fairness into specific decisions in a manner that yields a valid Sentencing Paradigm.

The basic idea behind our method is simple. The observer maintains an ordering by seriousness on the set of cases he has processed so far. He also has assigned specific sentences to some or all of those cases. This constitutes our Sentencing History. When a new case arises, the observer determines its seriousness relative to the existing cases and hence its position in the ordering. This then constrains the set of sentences which may be assigned in a manner consistent with his current Sentencing History. Depending on the context, he may or may not then select a specific sentence from this set. Either way, the Sentencing History is updated to reflect his new choices.

4. FORMAL FRAMEWORK

Now that the various ingredients of our system have been discussed, we are in a position to formally define it. We distinguish between two types of components. 'Stationary' components are defined initially and rarely change:

- (1) A set $S \subseteq S_{law}$ of sentences allowed by society and which the observer would consider applying.
- (2) A total preorder \leq on S , corresponding to the observer's view of the relative severity of sentences and the resolution to which he feels he can discern it. As discussed above, the ordering on S has the following properties:
 - Reflexive: $a \leq a$ for all $a \in S$.
 - Transitive: $a \leq b$ and $b \leq c$ implies $a \leq c$ for all $a, b, c \in S$.
 - Total: $a \leq b$ or $b \leq a$ or both for all $a, b \in S$.

On the other hand, nonstationary components evolve as new cases are introduced and processed:

- (1) A set of cases C that already have been processed by the observer.
- (2) A total preorder \leq on C which corresponds to the observer's assessment of the relative seriousness of offenses. As discussed in appendix A, this ordering has the following properties:
 - Reflexive: $a \leq a$ for all $a \in C$.
 - Transitive: $a \leq b$ and $b \leq c$ implies $a \leq c$ for all $a, b, c \in C$.
 - Total: $a \leq b$ or $b \leq a$ or both for all $a, b \in C$.
- (3) A set of anchor points $A \subseteq C$ which may or may not equal C .
- (4) A sentencing function $P_0 : A \rightarrow S$ for the anchor points. We require that if $c_1 \leq c_2$ then $P_0(c_1) \leq P_0(c_2)$.

The nonstationary components constitute the Sentencing History. In addition to a mechanism for rendering decisions, any Sentencing Paradigm must include both the stationary and nonstationary components in order to know the constraints to which it must adhere.

It is possible for a stationary component to change; for example, legislation may alter S_{law} . However such basic modifications may require a reconstruction of the entire Sentencing History²⁴, or at least large parts of it. Although any practical implementation must provide for this²⁵, we shall assume that the stationary components remain constant once established.

Rather than directly using S , wherever possible we prefer the induced total order on equivalence classes as discussed earlier. We define S' to be the set of equivalence classes under \sim in S . Each element of S' is a maximal set of sentences of equal severity. We will use S' to simplify our discussion, and revert to S only where the precise choice of sentence is relevant. In the context of S' , the expression $a \leq b$ means either $a < b$ or $a = b$. The analogue of P_0 when using S' is a function $P'_0 : A \rightarrow S'$ which takes us to the equivalence class of the sentence assigned by P_0 . That is, $P'_0(c) \equiv [P_0(c)]$.

Our introduction of anchor points may seem unnecessary. If the observer processes cases and assigns sentences from scratch then every case is an anchor point. We will see that anchor points serve an important purpose in the initial stages of the construction of a Sentencing History (or when reconstructing one after a change to a stationary component). Before doing so, however, we must discuss the sentencing function.

4.1. The Sentencing Function. The constraint on the sentencing function P_0 that $c_1 \leq c_2 \Rightarrow P_0(c_1) \leq P_0(c_2)$ is the key consistency requirement of our method. It embodies two relations. First, $c_1 < c_2$ does not simply imply that $P_0(c_1) < P_0(c_2)$. They also may be equal. Second, if $c_1 \sim c_2$ then $P_0(c_1) \sim P_0(c_2)$. This immediately follows from the symmetry of $c_1 \sim c_2$ and the consequent requirement that both $P_0(c_1) \leq P_0(c_2)$ and $P_0(c_2) \leq P_0(c_1)$. Together, these two relations comprise our sense of consistency. No crime may incur a greater penalty than does a greater offense, and two offenses of equivalent seriousness must be assigned equivalent sentences. In terms of P'_0 , the requirements are that if $c_1 < c_2$ then $P'_0(c_1) \leq P'_0(c_2)$ and if $c_1 \sim c_2$ then $P'_0(c_1) = P'_0(c_2)$. Because there is a total order on S' , the equivalence \sim is replaced by equality = in this latter relation.

We may construct an extension of the sentencing function P'_0 to all non-anchor cases in C by defining a function $P' : C \rightarrow I(S')$ where $I(S')$ is the set of all intervals on S' (open or closed at either end). That is, P' assigns each existing case a range of sentences. We require that:

- (1) $P'(c) \neq \emptyset$ for any $c \in C$. We cannot constrain away our ability to assign a sentence to every case.
- (2) $P'(c) = \{P'_0(c)\}$ for any $c \in A$. The extension P' must agree with the original sentencing function at the anchor points.
- (3) If $c_1 \leq c_2$ then $x \leq y$ for all $x \in P'(c_1)$ and $y \in P'(c_2)$. We cannot allow a case to receive a harsher punishment than a more serious case (they can be equal, though).
- (4) $P'(c)$ is the maximum subset of S' which obeys the above constraints. We impose no constraints other than those dictated by our ordering requirements and the prior choices (i.e. anchor points). $P'(c)$ clearly is an interval.

²⁴In the case of an algorithmic observer it may also entail relearning from scratch, affecting the Sentencing Paradigm as a whole.

²⁵Doing so is an extremely complex exercise because it may involve altering existing sentences and have profound effects on many active cases. Such changes run the risk of engendering distrust in the justice system, hence they must be conducted rarely and with great care.

Certain properties immediately follow from these requirements.

Lemma: If $c_1 < c_2$ then $P'(c_1)$ and $P'(c_2)$ may overlap in at most one point. *Pf:* If $c_1 < c_2$ then no sentence allowed by $P'(c_1)$ may be more severe than any sentence allowed by $P'(c_2)$ (though they may be of equal severity). Suppose that $x, y \in P'(c_1) \cap P'(c_2)$. Because S' is totally ordered, $x < y$ or $y < x$ for $x \neq y$. If $x < y$ then we could assign $a \in y$ to c_1 and $b \in x$ to c_2 , violating our requirement that $P(c_1) \leq P(c_2)$. The same holds if $y < x$.

Lemma: If $c_1 \sim c_2$ then $P'(c_1) = P'(c_2)$ *Pf:* c_1 and c_2 are subject to the same constraints from all other cases and impose no constraint on one another. Since $P'(c_1)$ and $P'(c_2)$ are maximized subject to these constraints, they must be the same.

Of course, if $P_0(c_1)$ is known then we are guaranteed that $P'(c_2) = [P_0(c_1)]$.

To use S instead of S' , we note that the analogue of P' is a function $P : C \rightarrow I(S)$ which takes us to the union of equivalence classes $P(c) = \bigcup_{x \in P'(c)} (x)$. The properties listed for P' all hold for P as well, except that if $c_1 < c_2$ then $P(c_1)$ and $P(c_2)$ may overlap in a set of sentences of equivalent severity rather than one point.

The extension P' is useful even when anchor points are not at issue. When a new case c arises and has been inserted at the appropriate place in the ordering, there is a non-anchor point in the new $C = C_{existing} \cup \{c\}$. Then P' reflects the constraints imposed by all of $P(C_{existing})$ on the possible values of $P(c)$. That is, we can treat any new case as a non-anchor point and use the extension function for it. More accurately, the opposite is true. We ordinarily would only define P' in the context of new cases. However we trivially may extend this to the more general situation where all cases have been ordered but only some have been assigned sentences. As we will see, this painless generalization can be of use.

5. OUR PROCEDURE

Our development of the formal framework applies equally well to any Sentencing Paradigm. We now consider the specific needs of a human observer. He is assumed to possess both reason and an invariant sense of fairness. Our Procedure provides a mechanism for him to translate these into specific choices. Taken together with both the stationary and nonstationary components of the formal framework described above, a given human observer along with the Procedure forms a Sentencing Paradigm.

5.1. Insertion. The insertion procedure for a new case is conceptually simple. Existing cases are presented to the observer and he determines whether the new case is less serious, more serious, or equally serious. When he has fully constrained the location, the procedure is complete. Computer algorithms are perfect for this. As with sorting, they reduce the procedure to a sequence of simple comparisons. The human observer serves the role of the comparison function. There are two important considerations. First, we must determine a sequence of cases to present for comparison with the new one. Any given computer algorithm has an approach to this. For example, we could use a balanced binary tree to represent the ordering and then choose the middle of each constrained region until we are very close. Then we would present a much denser set of cases until we've locked in a location. This should be at most $O(\log N)$ where $N = |C|$ is the number of cases already processed. The exact insertion algorithm is immaterial for our procedure.

The second issue is more of a problem. The human observer may be inconsistent. We assumed consistency, but there is every reason to believe that close comparisons may yield conflicting results.

This is true even if the observer is rational and constant²⁶. The algorithm must accommodate this. That is, if two comparisons conflict then they must both be presented to the human for reconsideration. He may then choose which to rectify. In more complex or egregious cases²⁷ he may choose to be presented with a sequence of cases (either the same sequence or a different one) from scratch. Accommodation must be made for this possibility.

Note that our use of a sequence of binary or ternary choices may seem pedantic. It could be easier for a person to simply scan the sequence of ordered cases and locate the correct position, and this may be the approach used in practice. There is nothing wrong with this²⁸. The reason we use comparisons is not purely formal, however. We assume that it always is possible for a person to make a choice between two cases. This is important for cases that are close together or that may not easily seem to fit into the existing ordering. While a quick and simple placement may suffice much of the time, the binary comparison method can always be relied on. It also places both human and non-human observers on an equal footing; as far as the ordering goes, they differ only in the choice of the comparison operator.

5.2. Procedure for Sorting Sentences. Though we glossed over it by presupposing the existence of an ordering on S , there must be a procedure for the observer to impose that ordering. This can be accomplished in an analogous manner to the ordering of a large prior set of cases. That is, we could do it all at once or through a series of binary comparisons. The latter is the preferred manner both because it is algorithmic in nature and because it guarantees a result under the simple but plausible assumption that a human observer could decisively compare any two sentences and remedy any inconsistency that arises²⁹.

The observer is performing a one-time sort of the set of sentences. This is a basic problem of computer science, and a plethora of algorithms exist to address it. We will not delve into them here.

5.3. Coverage. The set S is known at the outset, though C rarely is³⁰. We define S as “covered” if for every sentence there exists at least one case in C to which it has been assigned. That is, P is surjective. The inverse image of any sentence is non-empty.

Once S is covered, future sentencing decisions are constrained to at most two equivalence classes of sentences. Suppose that a new case c arises. If $c \sim c'$ for some c' to which a sentence has been assigned then we require $P'_0(c) = P'_0(c')$ per our earlier result. This leaves the case where c belongs to an equivalence class with no anchor points (or is in its own equivalence class). Let $c_1 < c$ be the “nearest” anchor point below c (that is the one the fewest ‘<’ steps away). If $P_0(c_1) = s_1$ then $[s_1] \leq P'_0(c)$. Likewise, if c_2 is the nearest anchor point above c and $P_0(c_2) = s_2$ then $P'_0(c) \leq [s_2]$. However every value of s has coverage, so there are no values of s between s_1 and s_2 . If $s_1 = s_2$ then we are constrained to one value $P'_0(c) = [s_1]$. If not, then $[s_1] < [s_2]$ and there are two possibilities: $P'(c) = \{[s_1], [s_2]\}$. This is not surprising. It always is possible to create a new equivalence class between any two existing ones.

While coverage does not fully determine the appropriate sentence equivalence class, it constrains it considerably. The observer need only choose between two adjacent equivalence classes when the case falls between the inverse images of those classes. Such instances should grow increasingly

²⁶The human observer is not a machine and may make mistakes or lack precision. However if presented with an inconsistency, his reason comes into play and he will address it.

²⁷For example, if the observer loses his way and starts changing decisions in a manner that progressively worsens his constraints.

²⁸This is, in fact, a set of comparisons done very quickly.

²⁹I.e., if a new comparison would prove inconsistent or unacceptably constrain the observer into a corner, a reworking of existing decisions may be necessary.

³⁰Except for ex-post-facto disposition of a set of cases, C grows as new cases arise and are inserted.

infrequent. For example, if there are N cases and M equivalence classes of sentences then the probability that a randomly located new case is between two classes is $\frac{M-1}{N+1}$. While there is no reason to expect such a uniform distribution for the relative seriousness of new cases, it nonetheless illustrates the benefit of coverage.

5.4. Degeneracy of S . There are two sources of degeneracy in S , and they are closely related. First, there is the actual ordering that the observer would impose on S in the absence of any other consideration. This leads to a set of equivalence classes of sentences that would truly be equal in severity according to the observer. However we do not need or wish for arbitrary resolution in distinguishing severity levels. For example, in theory fines could take a continuum of severity values that are strictly ordered. However such exactness is pointless. Among other things, it would make coverage impossible.

To avoid introducing spurious precision in sentencing decisions, it may be expedient to aggregate sentences into severity bins. Each bin would represent all sentences with a certain range of severities. The severity bins need not be uniform; they represent a partition of the set of sentences into coarse-grained severity classes.

For example suppose there are 200 sentences, and that if the observer were forced to impose an order on them the result would be 58 equivalence classes. The observer may believe that he only can meaningfully distinguish between 8 levels of severity. If he used all 58, then he would inject a false precision and unnecessarily constrain future choices. He either could sift the 58 severity classes into 8 groups through a new equivalence relation or embody that same coarse-graining in his original ordering on S . The distinction is irrelevant if the same person is performing the ordering and coarse-graining. However the distinction becomes meaningful if we attempt to use a universal ordering of sentences. That is, the set S_{all} could be ordered in a way that everyone agrees on. This isn't far fetched. Even if different observers would only consider the use of certain subsets of S_{all} , the relative severity of penalties itself is uncontroversial for the most part. Both prison sentences and fines are explicitly ordered by a numerical quantity. Indefinite sentences and other penalties or variants would be less trivially related, but may not engender differences of opinion³¹. The process involves two stages: the ordering on S_{all} is initially imposed by a universal standard, and then each observer constructs bins by modifying the order – but only through aggregation. That is, they may convert $a < b$ to $a \sim b$ but *not* to $b < a$. In this case, the dichotomy of role is important even if it doesn't have mathematical consequences.

A similar consideration applies to the granularity with which the observer distinguishes case severities. This would be embodied in his ordering of them rather than the explicit construction of bins. However in this instance there is no universal ordering that could be considered, and the dichotomy is entirely irrelevant. A coarse graining of the cases themselves (as opposed to their relative severities) could also be effected through an appropriate discretization of any parameters present in their representation.

5.5. The Purpose of Anchor Points. As mentioned earlier, our use of anchor points may seem superfluous. If the observer processes cases and assigns sentences from scratch then every case is an anchor point. They are important whenever the set of cases that has been ordered differs from the set to which sentences have been assigned. This could happen in learning the proper ordering to establish coverage. It is perfectly reasonable to build coverage based on previous known cases before processing active cases. For one thing, the sentencing decision is much simpler as there remain at most two equivalence classes to consider – and in a growing majority of cases there is no choice at all.

³¹That *someone* would disagree is a foregone conclusion – but whether any relevant observer would is less certain.

There is another important reason to pre-learn the seating of one's case ordering on the set of sentences. When only a few cases have been processed, sentencing decisions require great discretion. They constrain the possible future judgments. It would not be uncommon to wish for a readjustment of prior choices to better accommodate new cases. Once coverage is well established this ceases to be a concern. However the cost of adjusting past decisions for active cases can be very high. It may involve resentencing, which hurts the view of the judicial system and may involve legal or political challenges as well. A much easier solution is to develop a Sentencing History by processing past cases until coverage is well established. Active cases then run little risk of instigating a reassessment of prior decisions.

A similar situation would arise if a new member of a revolving panel needs to catch up by studying the cases disposed of by his colleagues. To do so, he may order a large set of cases and then apply sentences according to his own standard as an exercise. The basic problem is that every assigned sentence imposes a constraint on future ones. If he orders the entire prior case history and then seats it on S through sentencing assignments he can maintain consistency, but this is a very difficult problem to solve³². On the other hand, if he processes the cases as he would new ones then he has the problem of constraining himself unnecessarily.

5.6. Summary of the Procedure. A possible implementation of the process from the vantage point of the observer is as follows:

- (1) Determine S_{all} from the ambient legal framework.
- (2) Determine the subset S of sentences that are palatable.
- (3) Impose an ordering by severity on S using any sorting procedure and the observer's judgment, to within a reasonable precision.
- (4) Further aggregate the elements of S into groups of sentences that can reasonably be considered equivalent for his purposes.
- (5) Settle on a fixed algorithm A for presenting cases for comparison. The algorithm must guarantee an appropriate insertion after a reasonable number of comparisons.
- (6) Settle on a fixed algorithm B for selecting a sentence from within an equivalence class. Such a choice has no effect on constraints. Some possibilities include a preferred sentence in each class, a random choice, a choice that attempts to evenly distribute the sentences across cases, or a choice that attempts to use the most common previously-applied sentence in that class.
- (7) If desired (or necessary, as in the case of an algorithmic observer), determine a means of converting actual case records into a representation that contains all info that the observer deems relevant. This could be a set of mathematical parameters or a heuristic summary.
- (8) Select a set C of prior cases of size $N \gg M$ where M is the number of sentence severity equivalence classes (that is, $M = |S'|$).
- (9) Process the set C in some preselected sequence (perhaps random or alphabetical). For each case, perform a series of comparisons using the selected presentation algorithm A . The result is an insertion point for the new case in the existing ordering.
- (10) Process the set C in some preselected sequence that may or may not reflect the newly imposed order³³. For each case:
 - (a) Determine the set of allowed sentences subject to the choices already made.
 - (b) If that set is empty or if the presentation of this case illustrates a shortcoming in the prior set of choices, modify that set of choices as little as possible and always maintain consistency with the constraints. Modifications include shifting existing anchor point assignments or possibly removing certain anchor points altogether. When done, verify that all constraints are still satisfied. If this proves difficult because of the interdependent constraints, it always is possible to clear the slate and restart this step. In that

³²Obviously it can only be done ex-post-facto, and wouldn't work for assigning sentences on an ongoing basis.

³³In fact, it probably is best that it be some other order – perhaps random or alphabetical.

- case the order is preserved, but the assignment of sentences begins anew – now with the benefit of hindsight. Note that the separation of the entire ordering process into the previous step already gives the observer some sense of what he is working with.
- (c) Select a sentence from the allowed ones. If multiple equivalence classes are involved, this requires a real choice that will affect future constraints. Within an equivalence class, however, the choice is made by the selection algorithm B .
- (11) Continue to add cases to C as described above (thus increasing N) until there is significant coverage. That is, there are many cases that have been assigned each sentence.
 - (12) Begin to process new cases for real sentencing. When a new case comes in, do the following:
 - (a) Perform a series of comparisons using the selected presentation algorithm A , thus inserting the new case into C at the appropriate point.
 - (b) Determine the set of allowed sentences subject to the choices already made. Because coverage exists, this is either one or two equivalence classes of sentences.
 - (c) If only one equivalence class is allowed – as in the vast majority of instances – no human choice is needed. If two equivalence classes are allowed, choose one of them.
 - (d) Assign a sentence to the case using selection algorithm B .

6. FINITE VS INFINITE REPRESENTATIONS

For a human observer there is no constraint on the representation of cases. Each could consist of an arbitrary set of court documents written in common language. However the comparison of cases would then be time consuming and difficult. It is almost certain that a case representation would prove necessary. As discussed earlier, such a representation consists of a map from every case to a point in a set of codified information. Most likely, the latter would be some subset of a direct product of the sets of possible values for the relevant parameters. The map could involve a program that translates court records into points in that set (that is, n -tuples of values). Or this could be performed by humans. The precise definition of this map is one of the most difficult challenges of our method. It also is unnecessary for the purpose of our discussion, and we leave it an abstraction for now. Instead, we focus on the representation set.

Once we have a representation, the set of all possible cases is well-defined. Every case will map to some point in that representation. Cases that have the same representation are indistinguishable from the observer's standpoint. It is not strictly necessary that a representation set be finite, though it is hard to envision a situation where this is not preferable. Moreover any actual representation really is finite. Numerical codifications have a finite range and finite precision and occupy a finite number of bits. Of course, the resulting finite set may be enormous.

Unlike a human observer, an algorithmic one requires a precisely codified case representation. This is defined by the algorithm itself. Whether for a human or algorithmic observer, the choice of representation implicitly incorporates the precision to which parameter values can meaningfully be distinguished. This resolution limit is reflected in the discretization of the representation set.

A concrete example may be illustrative. Suppose we have 3000 cases, each with a set of descriptive court documents of varying detail, quality, and comprehensiveness. In reality an observer may consider 20 facets of each case important, but for simplicity let's assume there only are 3. Each of these takes values from some set, either discrete (numeric or non-numeric) or continuous. A human clerk may read through the documents associated with each case and assign an appropriate triplet of values. Denoting these x , y , and z , let us assume that they are represented by a 32-bit real value, a 16-bit positive integer, and an element of a 5-member set of categories. The direct product set would have around 1.4×10^{15} elements. Even if the representation set R were only a subset of this, it likely would be huge. However, such precision in the parameters x and y is greater than possibly could be of use. If the observer pretends to be able to discern the 4 billion values of x or 65536 values of y –

or is algorithmic and naively uses a formula to do so – then the 3000 cases almost certainly would have distinct representations. More likely an observer would wish to coarse-grain these into bins, let's say 22 for x and 8 for y . There would remain 880 possible representations. The 3000 cases then would prove a generous sample for getting a flavor of the representation set; they may even cover it.

If a new case has the same representation as an existing one, it is identical in the eyes of the observer. Therefore it must be equivalent under the case ordering and must be assigned an equivalent sentence (if a sentence has been assigned). This has an important consequence for a reasonably sized finite representation set R . Suppose the cases we have processed (and assigned sentences to) cover R . Every representation is associated with at least one case, therefore the disposition of all future cases is completely determined. An infinite or very large R can not be covered, but a small R may be.

If coverage is achieved, then the Sentencing Paradigm becomes deterministic regardless of whether the observer is human or algorithmic. A human observer would be removed from the process altogether. This may or may not be desirable. If R is large enough to meaningfully distinguish all properties that the observer cares about, then sufficient precedent relieves him of the need to judge future cases and the entire exercise becomes clerical. This is a mandatory sentencing (or sentence remedy) system, but one which distinguishes cases to the resolution desired by the observer. As such, it is not an abrogation of responsibility but a natural fruit of the observer's extensive prior labor.

An overly simplistic R may lead to an undesirable mandatory sentencing system which does not serve to remedy unfairness in the eyes of the observer. However such a representation is problematic even in the absence of determinism; the observer cannot distinguish cases beyond the detail provided by R , and the continuing use of his judgment is of no benefit.

In summary:

- If R is too large, then comparison of cases is difficult and perpetual involvement of the observer is certain.
- If R is too small, then the observer's judgment is subverted by a lack of meaningful information.
- If R is of reasonable size, then comparison of cases is manageable, the observer's judgment will be meaningfully incorporated, and after a reasonable number of cases are processed it is likely that his involvement will no longer be necessary.

7. AGGREGATION

For obvious reasons, a real implementation of our system is unlikely. Among other problems, the choice of observer would prove contentious. A concentration of power in the hands of a particular individual would be unpalatable and dangerous. Perhaps the best approach would be to empanel a group of individuals to serve as a review committee. Unless determinism or near-determinism has been achieved, distributing cases amongst the committee members would be tantamount to having several independent observers, and could lead to the very disparities which motivated our method. The alternative is to aggregate the individual opinions into a single effective observer. There is a danger with aggregation, however. Even if the individual members adhere to the ordering constraints of our method, an aggregate may not. Even a simple aggregation technique such as 'majority rule' could lead to violations of transitivity. Aggregation can be performed in two ways: by decision and by Sentencing History. They are identical for unrestricted sequences of choices, but differ when subject to the constraints of our method.

7.1. Decision Point Aggregation. We could aggregate individual opinions every time a choice must be made. This is much like the operation of the Supreme Court. Every comparison and every sentencing assignment would be performed by a committee. For example, a panel of judges could use a majority voting rule to make each choice. We assumed earlier that when presented with an inconsistent decision, an individual observer would amend either it or his prior choices to remedy the problem. This is a fundamental requirement, and easily could be violated by an aggregate decision maker. While individual choices would be consistent, there is no guarantee that the aggregate would be.

Consider the case of majority rule. Let us suppose there are n participants³⁴. Is the relation on C defined by their aggregate comparison choices a total preorder as needed? Denote the relation \leq_M and define it as $a \leq_M b$ iff a majority of the participants agree that $a \leq b$.

Lemma: The Majority Rule relation \leq_M is reflexive. *Pf:* For every participant and for all $a \in C$ we have $a \leq a$, so $a \leq_M a$.

Lemma: The Majority Rule relation \leq_M is total. *Pf:* Given any $a, b \in C$, for each participant either $a \leq b$ or $b \leq a$ or both. Therefore at least one of the two holds. If we separately count each instance of $a \leq b$ or $b \leq a$ then we obtain at least n . This means that either $a \leq b$ or $b \leq a$ (or both) has at least $\frac{n+1}{2}$ votes.

The key requirement common to all the relations listed in Appendix A is transitivity. In fact, from this and totality, reflexivity would follow. Unfortunately, the majority-rule relation \leq_M may be non-transitive.

In order to analyze transitivity, we must consider 3 elements. There are 27 possible choices of ternary comparisons among these, and only 13 are transitive³⁵. The following is an example where three observers use a majority-rule vote and transitivity is not preserved.

TABLE 1. Example of Failure of Transitivity in Aggregation

Choice	o_1	o_2	o_3	Majority
$s_1 \leq s_2$	T	F	T	T
$s_2 \leq s_1$	F	T	F	F
$s_2 \leq s_3$	T	T	F	T
$s_3 \leq s_2$	F	F	T	F
$s_1 \leq s_3$	T	T	T	T
$s_3 \leq s_1$	F	T	T	T
Order	$s_1 < s_2 < s_3$	$s_2 < s_1 \sim s_3$	$s_3 \sim s_1 < s_2$	$s_1 < s_2 < s_3 \sim s_1$
Transitivity?	OK	OK	OK	FAILS

For n observers, there are 13^n sets of ordering choices that preserve individual transitivity. For large n we expect at least $13/27$ of the majority-rule aggregates to be transitive³⁶. In the case of $n = 3$, there are $13^3 = 2197$ sets of individual observer choices that preserve transitivity, and only 1753 of these yield majority-rule relations that do so³⁷.

Sentence assignment from among the constrained options available could be voted on in many ways. For a well-established programme with a large Sentencing History and coverage of the set of sentences, the choice is at most a binary one between adjacent equivalence classes – and this, rarely.

³⁴We require that n be odd to ensure a decision each time.

³⁵These are $(a \sim b \sim c)$, $(a < b \sim c)$, $(b < a \sim c)$, $(c < a \sim b)$, $(a \sim b < c)$, $(a \sim c < b)$, $(b \sim c < a)$, $(a < b < c)$, $(a < c < b)$, $(b < a < c)$, $(b < c < a)$, $(c < a < b)$, and $(c < b < a)$.

³⁶The fraction actually is higher than this for various reasons.

³⁷This number is best arrived at through numerical calculation.

In those latter instances, a simple majority-rule vote would serve. As always, choice within an equivalence class is arbitrary and we do not concern ourselves with it.

The fact that a majority-rule aggregate *can* violate transitivity need not deter us from its use. If individual panel members subject their choices to the constraints imposed by the aggregate Sentencing History, transitivity is never in jeopardy. However, those very constraints may force individuals to perform comparisons in a manner inconsistent with their individual senses of fairness. This is inevitable, as the aggregate ordering certainly will diverge from those which individual members would have imposed. Thus each member must either endorse undesirable comparisons or introduce inconsistencies, rehashing the committee's previous deliberations. In practice, aggregation would prove most important during an initial learning period. The individual senses of fairness would serve to form early constraints, and would grow progressively less important as those constraints tightened. Subsequent administration would be near-deterministic, and could in fact be distributed amongst individual members. Only when a human observer needs to make a choice, would the panel convene and vote.

7.2. Ordering Aggregation. Instead of aggregating choices, individuals could maintain their own separate Sentencing Histories and aggregate these as needed. If we decompose the individual orderings into lists of comparisons as in Table 1, then the result is the same as that of decision point aggregation. Alternately, we could assign numerical positions to the various cases in the orderings and then average the position locations. This is very likely to violate transitivity for obvious reasons.

There is a basic problem with ordering aggregation. Individual Sentencing Histories impose differing constraints on future choices. In order to meaningfully aggregate them, they must also adhere to the constraints of the aggregate Sentencing History. For example, consider two observers *A* and *B* whose Sentencing Histories are to be aggregated. When processing a new case, the sentencing choices for observer *A* could be quite different from those available to observer *B* or to the aggregate. Moreover, the intersection may be empty. There may be no choice that *A* can make which is consistent with both his own Sentencing History and that of the aggregate. There may be ways around this, but in general aggregation of independently maintained Sentencing Histories is not a viable approach.

8. METRIC

Different observers produce different Sentencing Histories pursuant to their respective senses of fairness. It is natural to ask whether we can compare the corresponding Sentencing Paradigms in a meaningful way. That is, can we measure a “distance” between them? At minimum, we require a mathematical metric³⁸. However to make a comparison, we first must consider what exactly it is we wish to compare.

Unless the decision making process has become deterministic as discussed in section 6, we cannot compare the Sentencing Paradigms themselves³⁹. The real product of our method is a set of sentence assignments to cases – that is, cases with outcomes – and we could attempt to compare these. However there is no natural means of doing so. We are not guaranteed that the same representation is shared by observers or even that the same sentence ordering is. Though we often assume the latter, we must recognize it as an additional requirement. There is no distance metric on the set of cases or the set of sentences.

³⁸Given elements x , y , and z of a set G , a distance function $d : G \times G \rightarrow R$ is positive definite ($d(x, y) \geq 0$, with $d(x, y) = 0$ iff $x = y$), symmetric ($d(x, y) = d(y, x)$) and obeys the triangle inequality $d(x, y) + d(y, z) \geq d(x, z)$.

³⁹And if it is deterministic, then doing so is no different than comparing Sentencing Histories.

The only additional structures present in our Sentencing Paradigm are the orderings of cases and sentences. If we are to have any hope of comparing Sentencing Paradigms, we must rely on these orderings to do so. Of course, we could introduce additional structure – such as metrics on the sets of cases or sentences – but then we would not be comparing Sentencing Paradigms. Moreover we would have to assume that any such structure is common to both observers. Even if we restrict ourselves to comparing Sentencing Histories, we are not free of the need to make certain assumptions of commonality. At the very least we must restrict ourselves to comparing a set of cases shared by both Sentencing Histories. For our purpose, we shall assume that the Sentencing Histories have identical case sets C ; however, if this is not true, then we must restrict C to their intersection before any comparison is made.

There are two components to a Sentencing History: the ordering of cases and the assignment of sentences. We may choose to compare one or both of these, as suitable for the task at hand. For a comparison of case orderings, no additional assumptions are needed. However, for a comparison of sentencing assignments we must assume that the observers agree on the ordering of sentences to within coarse-graining. This means that we allow one ordering of sentences to be a refinement of the other. In that case every equivalence class of observer A is a subset of an equivalence class of observer B, the associated strict total orders are compatible, and we use the coarser partition (that of B) for our common ordering. We do not require that the sets of sentences employed by the observers be identical, only that the observers agree on an ordering of S_{law} as described. For convenience we will refer to S instead of S_{law} where unambiguous.

Note that if additional structure *is* present and common to the observers under consideration, then it may prove of practical use for comparison. The result would not be a comparison of Sentencing Paradigms or Sentencing Histories per se. But that need not matter from the standpoint of utility. For example, if an external metric has been imposed on the set of sentences then we may use it to compute a distance between Sentencing Histories. This is perfectly valid if we recognize that the metric constitutes additional information and that we are no longer strictly working within our method.

8.1. A Class of Swap Metrics. Consider two orderings of the same type on a given set G ⁴⁰. If we require that they be strict total orders, then we have two different lists of the elements of G . Let us denote these two sequences x and y . Through a series of swaps of elements in sequence x we may eventually arrive at sequence y . There are various ways to accomplish this, and the procedure is symmetric – if we apply the reverse sequence of swaps to y then we get x . Let us define a quantity d that counts the minimum number of such swaps needed to convert sequence x into sequence y . This minimum sequence of swaps may not be unique, but the minimum count will be. It is not hard to see that d is a true metric. It trivially is positive definite and symmetric (just reverse the series of swaps as mentioned). Denoting by z a third sequence drawn from the same set, the triangle inequality follows from the observation that we always can concatenate the sequences of swaps from x to y and y to z to get from x to z . This places an upper bound on the minimal length sequence of swaps to get from x to z . Note that we said nothing of the type of swap used. Any class of swaps that can move us between arbitrary sequences will work. The two natural choices are adjacent swaps and arbitrary swaps. We shall denote the metric obtained using adjacent swaps d_1 and that using arbitrary swaps d_2 .

Our metrics d_1 and d_2 only are defined for strict total orders, and may not seem applicable to either the case or sentence orderings. For orderings which allow equivalent elements, how would we define a swap? We could attempt to use the equivalence classes of sentences or cases – which do have strict total orders. However, those equivalence classes need not be the same for both observers. In fact,

⁴⁰For example, the orderings $a < b < c < d < e$ and $b < a < d < c < e$ both are strict total orders on the set $\{a, b, c, d, e\}$.

they likely won't be unless the two orderings are the same⁴¹. So how do we get around this? A simple modification of our swap metric suffices. The problem is that there is no strict order within an equivalence class. A total preorder can be thought of as a set of strict total orders, with the same strict order among equivalence classes and differing strict orders within them. For example, $(a \leq b < c \leq d)$ can be thought of as the set of four strict orders $(a < b < c < d)$, $(b < a < c < d)$, $(a < b < d < c)$, $(b < a < d < c)$. Our sequences x and y have total preorders on them, and each may be considered such a set of strict total orders. We then simply define our distance d to be the minimum number of swaps when comparing any of the strict orders corresponding to x with any of those corresponding with y . Suppose x corresponds to 4 strict total orders (as in our example above) and y corresponds to 18. Then we obtain 72 values of d , the metric between strict total orders, and select the smallest. The resulting distance function d' is a true metric. It trivially is positive definite and symmetric. The triangle inequality follows from the same argument as for d . As their use is unambiguous, we will use d_1 and d_2 to denote the metrics between total preorders obtained from the d_1 and d_2 defined above for strict total orders.

We note in passing that we may use a strict total order on a specific set to construct a metric, in this case the number of steps along the ordering between elements. Though this metric has no obvious analogue for total preorders, counting the strict steps between elements can be useful anyway⁴².

8.2. Comparison of Case Orderings. The comparison of case orderings may directly be made using our metrics d_1 and d_2 for total preorders. Whether arbitrary or adjacent swaps are appropriate depends on the context. It is likely that most of the time adjacent swaps will prove more suitable. Arbitrary swaps can make quite different systems look similar. For example, if there are two strict total orderings of N cases and one is the reverse of the other then $d_2 = \frac{N}{2}$ while $d_1 = \frac{N \cdot (N-1)}{2}$. Clearly d_1 would better represent the difference for most purposes. On the other hand, suppose that two sequences only differ in that the end values are exchanged. Then $d_1 = 1$ and $d_2 = 2N - 3$. In that case we may or may not consider the sequences very far apart, and the choice is less clear cut.

8.3. Comparison of Sentence Assignments. In order to compare sentence assignments, we require a common sentence ordering as described earlier. In the presence of one, we may use the aforementioned counting method. We cannot define a metric if we attempt to distinguish between equivalent sentences, so we consider equivalence classes instead. That is, we use the strict total order on S' . Given a set of cases C and two sentence assignment maps $P'_x : C \rightarrow S'$ and $P'_y : C \rightarrow S'$, we may count the total distance between elements as $d(x, y) = \sum_{c \in C} n(P'_x(c), P'_y(c))$ where $n([s_1], [s_2])$ is the number⁴³ of strict steps between $[s_1]$ and $[s_2]$. Clearly, d is positive definite and symmetric. Consider a single case c and three sentences $\{s_1, s_2, s_3\}$. If $P'_x(c) = s_1$, $P'_y(c) = s_2$ and $P'_z(c) = s_3$ then $d(x, y) + d(y, z) \geq d(x, z)$ regardless of how the sentences are ordered. If this is true for each individual case, then it is true for the sum. Therefore, the triangle inequality holds.

As with any use of ordering metrics to compare sentence assignments, d has a critical shortcoming. It takes no account of the actual severity of sentences. Suppose we have three sentences: s_1 is 1 month in prison, s_2 is 2 months in prison, and s_3 is life in prison. Clearly $s_1 < s_2 < s_3$ to any rational observer. However, such an observer would also agree that s_3 is much farther from s_2 than s_2 is from s_1 . This is not accounted for by an ordering metric. Of course, it *cannot* be accounted for in our system, and requires additional structure as discussed.

⁴¹Specifically, they won't be the same unless the two observers' orderings happen to define the same equivalence classes and differ only in the strict total order imposed on these.

⁴²For example, if $a \leq b < c < d$ we would count 1 step between a and c and 2 steps between b and d .

⁴³There are other possible metrics. We probably could use n^α for $\alpha > 1$, but would need to prove that it gives rise to a metric.

8.4. Comparison of Sentencing Histories. If we wish to compare entire Sentencing Histories, then we must be careful. Suppose we have two Sentencing Histories, x and y , with corresponding (specific) sentence assignment maps P_x and P_y . As always, we assume that they share the same case set C , and denote the corresponding ordered sequences C_x and C_y . Because we are working with Sentencing Histories produced by Sentencing Paradigms, we know that P_x and C_x are consistent with the order on S , as are P_y and C_y . If we were to transform C_x into C_y through a series of swaps without changing P_x , then we would have C_y with P_x and it is highly unlikely that this would be consistent with the ordering on S . We could then measure the number of swaps of assigned sentences needed to convert P_x into P_y or return it to an ordering consistent with S . There are various ways to do this, and we won't delve into them here. We simply note that this is not a particularly enlightening approach. The same problems encountered with the comparison of sentence assignments are worsened by the need to incorporate the distance between case orderings as well.

9. DEPENDENCE ON SEQUENCE OF PRESENTATION

One important question is whether the Sentencing History that results from an application of our method depends on the sequence in which cases are processed by the observer. As before, let the set of cases encountered be C . If we consider only the ordering imposed on C , then we have good reason to expect a unique result. It is reasonable to expect that the observer will arrive at the same conclusion given any comparison at any point. Moreover the observer's attitude won't change as he processes cases. A comparison of two cases is based on his existing sense of fairness and ability to distinguish between them⁴⁴.

However, when it comes to sentencing there is a problem of hysteresis. This arises for two reasons, one human and one mathematical. Unlike comparisons, which are relative, assignments are absolute. Unless an observer has an existing absolute sense of the appropriate sentence for each offense, it is likely that his view will change as cases are processed. This does not reflect a change in his sense of fairness or an inconsistency. Rather, it is the natural consequence of a growing familiarity with the range of offenses, their frequencies, and the sentences that may be applied. Even if he knows all of these at the outset, it is only in the personal application of this knowledge that his sense of proportion evolves⁴⁵. This attitude will stabilize after a sufficient case history has been processed.

The mathematical hysteresis problem is intransigent, however. Initial sentence assignments are relatively unconstrained, but they constrain future choices. If readjustment is allowed, then those choices may be modified or the entire ordering reseated. However changing prior choices requires much more impetus than choosing from amongst a set of available sentences. Only when the observer feels a strong need will he do so. Therefore it is likely that he will adhere to constraints – even if not optimal – as long as he can. If the deviation from his sense of fairness grows too great he may reconsider past choices. As such, the initial choices carry far greater weight than later ones. Therefore a hysteresis issue exists. Moreover, humans are not entirely rational and the sentencing choices may not be identical even if presented with the same sequence of cases⁴⁶.

We could mitigate this to some extent by requiring a learning period prior to any application of the system. A large set of existing cases would be processed to achieve coverage of the set of sentences (or possibly even the representation set itself). During this period, revision and reseating has no cost other than effort. This allows a human observer who otherwise would always arrive at the

⁴⁴This is subject to human error, which could lead to apparent inconsistencies. However on being presented with these, the observer would make the appropriate correction. His internal criteria for choosing never change, even if they are imperfectly reflected in action.

⁴⁵This is much like a teaching assistant grading papers. At first, he has no idea how many points to deduct for each error. He develops a view with time depending on the prevalence of certain errors or the presence of worse ones.

⁴⁶Note that all our references to sentences really are to equivalence classes. We do not care about different choices amongst equivalent sentences.

same conclusion to avoid being blocked by constraints he injudiciously imposed early on. That is, it circumvents the purely mathematical hysteresis issue.

Taking this a step farther, we may remove most of the human hysteresis issue as well by separating the learning process into two parts. First, the set of cases is ordered. Only then are sentences assigned. The latter is accomplished in one pass, and with as much revision as desired. We could even require that it be done twice, to guarantee that the observer knows what he will encounter the second time around. This removes the effect of the sequence of presentation, subject only to the inherent vagaries of human thought.

In summary, while we expect the ordering of cases by seriousness to be independent of the sequence in which those cases have been presented, the sentencing assignments need not be. Though we can take steps to reduce the potential dependence, we may not eliminate it altogether. At best we can hope that, if reseating and revision are encouraged or mandated, the same observer will reach a similar – if not identical – Sentencing History regardless of the sequence in which cases are presented.

10. ALGORITHMIC OBSERVERS

Throughout our discussion we have touched upon the use of algorithmic observers. Unless we presuppose a complete solution of the problem – an ordering and sentence assignment map on the entire set of possible cases – then the underlying algorithm must be adaptive and learn from the cases it processes. After sufficient experience, its adaptation may cease or slow down considerably⁴⁷, but this is a consequence of learning rather than explicit design. As with a human observer, each sentencing choice constrains future options. We may allow for modification or reseating of existing sentences when processing prior cases. However it is not permissible to do so with active cases. We focus on generalities concerning the uses and behavior of an algorithmic observer, and do not delve into the various approaches to its design.

While it is unlikely that an algorithmic observer would be considered for a real decision-making role, there are a number of uses to which it is well suited. These pivot on its ability to evolve a specific Sentencing Paradigm from experience processing a set of existing cases. Such learning requires certain information. If this information must be supplied by a human observer on an ongoing basis, then the algorithm is at most an advisory tool. Only if it can make decisions on an ongoing basis without intervention can it be considered autonomous.

There is no natural means for an algorithm to order the set of sentences, and a human observer must do so instead. However, this is provided once at the outset, and the ordering of sentences likely is universal as discussed earlier. Aside from a resolution-limit, it shouldn't depend on the observer. Therefore, an algorithmic observer may require such information, yet remain independent.

A similar situation holds for the ordering of cases. An algorithm has no natural means of imposing such an ordering *ab initio*. This is a much more serious concern than the sentence ordering. Unless the representation set is covered, ordering choices continually must be made. If the algorithm is serving in an auxiliary capacity or is in a learning mode, then a human observer could make these. But to process active cases autonomously, the algorithm must make these decisions on its own. However unlike the sentence ordering – which is imposed without any prior knowledge at the outset – the case ordering may be learned. This requires a source of ordered examples. Such examples either may be deduced as described below or provided by a human observer. For example, a simple regression algorithm could attempt to assign a single number to each element of a multi-parameter representation set, and then numerically order the cases.

⁴⁷As discussed earlier, complete coverage of the representation set may lead to a deterministic outcome for each new case.

Sentencing information may also be learned from. By processing a prior set of cases with outcomes, an algorithm may deduce the appropriate decisions to make when presented with future choices. It also could make random choices or attempt to evenly distribute sentences, for which it would only require cases without outcomes. However such an approach is not likely to be of much use. Only if additional structure were present – such as metrics on C and S or numerical indices representing absolute scales of seriousness and severity – would such an approach conceivably prove fruitful.

It may also be meaningful to use an algorithm to learn from existing cases (with or without outcomes) until it establishes a deterministic Sentencing Paradigm (that is, the representation set is covered). Such an algorithm could “fill in” the blanks of what a human observer would do, given a small sample of his actual decisions⁴⁸. It also could be used to deduce an “implied” Sentencing Paradigm based on the real sentences assigned cases by judicial institutions, an application we will consider shortly. Note that at any point, a truly autonomous algorithmic observer can deterministically map the entire representation set to sentences. However that map may change as new cases are processed and the algorithm evolves. Only when the Sentencing History of the algorithmic observer itself has covered the representation set are its choices constrained away. From then on, no learning may have any impact on the outcome of a processed case. The map has ossified.

An implied Sentencing Paradigm is one which reflects the implicit sense of fairness embodied in real world case dispositions. It would be learned solely from these and would require an external ordering on the set of sentences. As mentioned earlier, the latter should be fairly universal aside from a possible coarse-graining. The comparison of two real cases would be accomplished by comparing their sentences, the reverse of our usual procedure. This yields a set of case comparisons that most likely will be horribly inconsistent⁴⁹. To obtain a Sentencing Paradigm, we must transform this into a suitable ordering on the set of cases. Even if this were accomplished, the real sentence assignments undoubtedly would violate the rules of our method. The algorithm must reseat the ordered set of cases on the ordered set of sentences in a way that does not. Because there are multiple orderings we could arrive at by changing different comparisons, the two steps are interdependent and we must combine them. By penalizing ourselves for various adjustments to orderings and sentence assignments, we could attempt to construct the Sentencing Paradigm that accomplishes this while most closely adhering to the real case dispositions. This requires a utility metric that measures our distance from the latter⁵⁰. The means of doing this could be quite complex and may involve relaxation or randomization algorithms to find an optimal solution.

The construction of an implied Sentencing Paradigm could serve several purposes. First, it could be used to measure the inconsistency amongst real world decisions. It also furnishes a Sentencing History that can be compared with a given observer’s in order to measure the difference. This could not be accomplished with raw real-world decisions, only with a valid Sentencing History.

11. MAIN CONCLUSIONS

A practical implementation of our method to reduce the unfairness of real sentences would face many challenges. The greatest obstacle would be the deference paid the sovereignty of various jurisdictions. In the absence of major and improbable structural changes, this would preclude enforcement authority by a central review board. As is evident from our discussion, a single such review board would be necessary. While various judicial units could employ our method, and even individual judges could make use of it to ensure consistency, only with complete scope could an observer maintain overall consistency. With proper data gathering and retention requirements – which could be mandated within the current constitutional framework – a central review board may still

⁴⁸Though such a prediction cannot be relied on, or human observers would be superfluous.

⁴⁹The corresponding relation on $C \times C$ will look nothing like an ordering.

⁵⁰The metrics discussed in section 8 aren’t applicable because they compare Sentencing Histories and assume an ordering on the set of cases.

serve a valuable advisory capacity. It could measure the extent of unfairness, detect jurisdictions which appear to foster the most egregious or frequent instances, and track what the regional and subordinate authorities within our nation are doing to its people. It also could issue recommendations for sentence adjustments where significantly unfair outcomes are encountered and for overall reform where systematic inconsistency appears to be present. The mere presence of such a review board, even if merely advisory, would expose the behavior of our overall judicial framework to the public eye and create a better impression of the government's concern for fairness. The following are some conclusions from our analysis, and may be of use in an implementation of our method.

- If the set of sentences is covered, there remain only progressively less frequent binary choices.
- If the representation set is covered, future decisions are deterministic.
- It is best if the representation set is mid-sized and statistically distributes cases uniformly among the majority of its elements.
- When first processing cases, the choice of sentence has more freedom and greater significance to future decisions than later assignments do. There also is a greater chance that future cases will require revision of those choices. It is best to perform this startup phase on a set of existing cases, and only proceed to active cases once the set of sentences has been covered.
- It would be reasonable to use a revolving committee of judges, with long or permanent tenures, as an aggregate observer. Majority rule would be used for decisions, constrained by prior choices to avoid inconsistencies. Once near-determinism is achieved, cases could be distributed among justices and need only be brought before the committee if they are on or near a boundary between sentence equivalence classes.
- Because determinism or near-determinism eventually is achieved by almost any Sentencing Paradigm, provision must be made for a periodic reseating or revision process to accommodate changing sensibilities or new panel members. Any such alterations should be slow and carefully considered, and their application to already processed cases restricted only to the greatest disparities between the old and new programmes.

APPENDIX A. REVIEW OF MATHEMATICAL ORDERINGS

We provide a brief overview of orders. The interested reader should consult on of the many excellent textbooks on Set Theory or Algebra for a proper treatment of the subject. We assume basic familiarity with set theory.

Given a set S , a binary relation R is a subset of $S \times S$. We write aRb for the pairs $(a, b) \in R$. For example the relation $<$ defined on the integers consists of all pairs (a, b) for which $a < b$. For ordering relations, two elements a and b are said to be comparable if aRb or bRa or both.

First, we list the primary properties that may be of interest. A given relation may satisfy some or none of these.

- Reflexive: aRa .
- Irreflexive: $\!(aRa)$ (i.e. we never have aRa).
- Symmetric: aRb implies bRa (but neither need be true).
- Asymmetric: aRb implies $\!(bRa)$ (but neither need be true).
- Antisymmetric: aRb and bRa implies $a = b$.
- Transitive: aRb and bRc implies aRc .
- Transitive Incomparability: If a and b are incomparable and b and c are incomparable, then so are a and c .

- Total: For any pair a, b , we either have aRb or bRa or both. I.e. any two elements are comparable. Note that Total implies Reflexive.
- Trichotomous: For any pair a, b , either aRb or bRa or $a = b$.

The following are some common relations that can be defined on S , their typical notation for R , and their set of properties. Properties in parenthesis follow from the other properties listed.

- Equivalence (\sim): Reflexive, Symmetric, Transitive.
- (Weak) Preorder (\leq): Reflexive, Transitive.
- Total Preorder (or Weak Order) (\leq): (Reflexive), Transitive, Total.
- (Weak) Partial Order (\leq): Reflexive, Antisymmetric, Transitive.
- (Strict) Partial Order ($<$): Irreflexive, Asymmetric, Transitive.
- Total Order (or Linear Order) (\leq): (Reflexive), Antisymmetric, Transitive, Total.
- Strict Weak Order ($<$): Irreflexive, (Asymmetric), Transitive, Transitive Incomparability.
- Strict Total Order ($<$): Irreflexive, (Asymmetric), Transitive, Trichotomous.

Some important relationships follow:

- An Equivalence \sim partitions S into equivalence classes, each of whose elements obey $a \sim b$.
- The usual relation $<$ on real numbers or integers is a Strict Total Order.
- A Preorder induces an Equivalence defined as $x \sim y$ iff $x \leq y$ and $y \leq x$. There is an induced Total Preorder on the set of equivalence classes.
- There is a one-to-one correspondence between Strict Weak Orders (with relation $<$) and Total Preorders (with relation \leq) on a set S . It is defined by $(x \leq y)$ iff $!(y < x)$. The same correspondence holds between Strict Partial Orders and Weak Partial Orders. Incomparable elements of the Strict Weak Ordering are treated as equivalent elements (using the induced Equivalence) in the Total Preorder.
- The following three are interchangeable (these are the orderings we make use of):
 - Strict Weak Order $<$ on S .
 - Total Preorder \leq on S related by $(x \leq y)$ iff $!(y < x)$.
 - Total Order \leq on the equivalence classes of S using the induced equivalence relation under the total preorder's \leq .
- When we speak of $<$ and \leq on the reals or integers, they are the related Total Order and Strict Total Order. The induced equivalence classes each contain a single element.

REFERENCES

- [1] Lindsey Devers, "Plea and Charge Bargaining - Research Summary," U.S. Department of Justice, 2011
- [2] Douglas Husak, "Overcriminalization: The Limits of the Criminal Law," Oxford University Press, 2008
- [3] Frank McLynn, "Crime and Punishment in Eighteenth-Century England," New York: Routledge, 1989