§ THE DIALECTIC REVIEW

Volume II · Winter 2024

ISSUE 04

Justice Without Consensus

Winter 2024 · 8 pieces

Eight essays on distributive fairness, punishment, and what justice can still mean once we give up hoping everyone will agree on it.

ISSN 3049-7146

thedialecticreview.org

CONTENTS

POLITICAL PHILOSOPHY 4

What Rawls Owed the Worst Off

Aditi Verma — Ashoka International School, New Delhi

ETHICS 7

Desert Without Merit

Emeka Chukwu — Federal Science College, Owerri

POLITICAL PHILOSOPHY 10

Why Majorities Don't Make Things Right

Hannah Fitzgerald — Kilbride Grammar School, Dublin

ETHICS 13

Retribution Without Anger

Rafael Duarte — Colégio das Amoreiras, Lisbon

METAPHYSICS 16

Can Justice Be Impartial About Time?

Mireille Dupont — Collège International, Geneva

EPISTEMOLOGY 19

The Epistemics of a Hung Jury

Daniyar Bekov — Almaty Science Lyceum, Almaty

ETHICS 22

Proportionality and the Problem of Incommensurable Goods

Grace Kim — Hanwoori International School, Seoul

POLITICAL PHILOSOPHY 25

Is Restorative Justice Still Justice?

Tobias Reinholt — Gymnasium an der Elbe, Hamburg

POLITICAL PHILOSOPHY

What Rawls Owed the Worst Off

The difference principle is usually defended as a compromise between equality and incentive. This essay argues the compromise is less stable than it looks, once we ask what justifies stopping where it stops.

Aditi Verma Ashoka International School, New Delhi

John Rawls's difference principle permits social and economic inequalities only when they work to the greatest benefit of the least advantaged — a formula meant to let incentive-driven inequality exist (a surgeon can earn more than an orderly) while insisting that this inequality earn its keep by actually improving the position of the worst-off group, rather than merely enriching those already ahead. The usual objection from the left is that the principle concedes too much to incentive-based inequality; the usual objection from the right is that it constrains free exchange too much. This essay raises a different problem: the difference principle's own justification, taken seriously, does not obviously stop where Rawls stops it, and the place he stops looks less like a principled boundary than a comfortable one.

Rawls's argument for the difference principle proceeds from the original position: rational parties, behind a veil of ignorance about their own place in society, would choose principles that protect them against the worst possible outcome, since they do not know whether they will end up rich or poor, talented or not. This reasoning is supposed to yield exactly the difference principle — maximize the position of the worst-off — as the risk-averse choice under radical uncertainty. But notice that the veil of ignorance, as Rawls constructs it, already assumes the parties know they will end up within a single,

bounded society whose basic structure the principles are meant to govern. Once you ask why the veil should be drawn around a national society specifically, rather than around, say, a global population, or a population extended across time to future generations, the same risk-averse logic that generates the difference principle within a nation seems to generate an argument for extending it further: a rational chooser behind a truly comprehensive veil of ignorance, not knowing whether they would be born in a wealthy or a poor country, should be just as reluctant to accept national boundaries as a limit on redistribution as they are reluctant to accept an aristocrat's birth-privilege as a limit within a nation.

Rawls himself, in his later work on the law of peoples, resisted this extension, arguing that the difference principle's logic applies specifically to the basic structure of a single, self-contained cooperative scheme with shared political institutions, and that the relationship between separate peoples does not have this cooperative structure in the same way.[1] This is a substantive and contestable empirical claim about how integrated the global economy actually is, not a further deduction from the original position's own logic, and critics like Thomas Pogge have argued persuasively that in a world of extensive trade, resource extraction, and interlocking financial systems, the cooperative structure Rawls requires for the difference principle to apply is present globally to a much greater degree than the two-tier theory admits.[2] If Pogge is right about the empirical premise, Rawls's own reasoning, applied consistently, would seem to generate a much more demanding global difference principle than Rawls himself was willing to accept — which suggests the stopping point at the national border was doing less principled work, and more work as an assumed starting condition of the theory, than the original position's own logic can independently justify.

This is not, by itself, a knock-down argument for global redistribution on Rawlsian terms — one could instead conclude that the original position's reasoning was never as clean as it first appears, and that some other, non-veil-of-ignorance justification is needed for confining principles of justice to national societies specifically. But that would be a significant concession, because it means the difference principle's celebrated derivation from pure rational self-interest under uncertainty was never actually delivering the conclusion on its own; it was smuggling in, as a background assumption about the scope of the original position, exactly the boundary that most needs justifying if the theory is meant to explain why inequality between nations raises different questions than inequality within one.

The lesson worth drawing is not that Rawls was wrong about domestic justice, but that his theory's stopping point at the nation-state is more exposed to challenge than its elegant internal derivation suggests, precisely because the derivation's force — what would a rational, risk-averse chooser accept behind a veil of ignorance — does not contain, within itself, any obvious reason to draw the veil around one population rather than a larger one. Anyone who wants to keep the difference principle's power while confining its scope to national societies needs an independent argument for that scope, stated and defended on its own terms, rather than treated as though the original position had already settled it.

NOTES

1. John Rawls, *The Law of Peoples* (1999).
2. Thomas Pogge, *Realizing Rawls* (1989) and later work on global justice.

ETHICS

Desert Without Merit

We say people deserve rewards for their talents and effort. This essay argues that neither talent nor effort survives scrutiny as a genuine basis for desert, and asks what, if anything, is left.

Emeka Chukwu Federal Science College, Owerri

The ordinary justification for unequal reward appeals to desert: the more talented surgeon, or the harder-working student, deserves more than the less talented or less diligent one, because they have earned it through some combination of natural ability and effort. This essay traces a familiar but underappreciated regress in this justification and asks what survives once the regress is taken seriously, rather than stopped short for convenience.

Start with talent. Nobody chooses their native cognitive or physical endowments; a person born with unusual mathematical aptitude did nothing to acquire that aptitude before having it, and it is hard to see how a trait a person did not choose or work for could ground a desert claim in the way we usually mean the term — desert is supposed to track what an agent has done, not the arbitrary starting conditions their agency happened to be dealt. This point is familiar from Rawls's discussion of the "natural lottery" and is widely, if reluctantly, conceded even by many defenders of merit-based reward: talent alone cannot ground desert, because talent is unearned.

The natural retreat is to effort: even if talent is unearned, surely the discipline to cultivate and apply that talent — to practice, to study, to persist through difficulty — is something a person does, and can be credited for. This is where most everyday defenses of desert land, and

it has real intuitive force. But the same regress that undermined talent applies, on inspection, to effort as well: a person's capacity for sustained effort — their baseline conscientiousness, their tolerance for delayed gratification, their resilience under difficulty — is itself substantially a product of genetic endowment and early environment neither chosen nor earned, as twin studies on conscientiousness and self-control have repeatedly shown.[1] The person who finds it easy to persist is not thereby more praiseworthy than the person who finds it hard, if the ease of persisting is itself an unchosen trait exactly analogous to the ease of native mathematical talent. Effort was supposed to be the earned remainder left over once we set aside unearned talent; the regress shows effort is, at bottom, no less a product of unchosen starting conditions than talent was.

Galen Strawson has pressed this regress to its most radical conclusion: since any trait that grounds effort must itself be either unchosen or the product of a prior choice made by a self whose character (and hence capacity to choose well) was itself unchosen, desert in the full, ultimate sense required to justify differential reward is impossible for any human trait whatsoever, all the way down.[2] This is a genuinely uncomfortable conclusion, and this essay does not fully endorse the nihilism Strawson draws from it, but the regress itself is hard to block at any principled, non-arbitrary point once you have already agreed that talent alone cannot ground desert. If we are going to keep talking about desert at all, we need either to accept ultimate desert is impossible and find a different, non-desert justification for unequal reward (incentive effects, for instance, which do not require anyone to deserve anything, only that the arrangement produces better outcomes for everyone including the worst-off), or to define desert in a deliberately more modest, comparative sense: not "deserving more than you would have without your unchosen traits," which the regress defeats, but something like "having actually done the specific, effortful thing

being rewarded, whatever produced the capacity to do it" — a genuinely thinner notion, but possibly the only one the regress leaves standing.

This thinner notion changes what desert-talk can honestly do for us. It cannot underwrite the strong claim that a talented, hardworking person has earned their advantage over someone less fortunate in the full sense that makes the disadvantaged person's worse position simply their own responsibility to bear without complaint — that claim requires the very ultimate desert the regress dissolves. What it can underwrite is a much more modest, procedural point: that rewarding actual effortful performance, rather than mere unearned potential, still serves a purpose (identifying who actually did the work, incentivizing people to do it), even once we admit that nobody, in the deepest sense, deserved to be the kind of person capable of doing it in the first place. Whether that thinner justification is enough to support the level of inequality most societies currently tolerate is a separate question this essay leaves open — but it is clearly a smaller, more defensible claim than the one desert-talk usually smuggles in.

NOTES

1. See, e.g., research summarized in Turkheimer's work on the heritability of conscientiousness and related personality traits.
2. Galen Strawson, "The Impossibility of Moral Responsibility," *Philosophical Studies* 75 (1994).

POLITICAL PHILOSOPHY

Why Majorities Don’t Make Things Right

Democratic legitimacy is often treated as though popularity itself were a form of justification. This essay separates the two and asks what majority support can and cannot license.

Hannah Fitzgerald Kilbride Grammar School, Dublin

It is common to hear a policy defended on the grounds that most people support it, as though popular support were itself a reason the policy is just, rather than merely a reason it is likely to be implemented. This essay argues these are genuinely different claims, routinely conflated in political speech, and that keeping them apart matters for how we should respond when a clear majority supports something that looks, on independent grounds, unjust.

The conflation is easy to miss because majority support does real, legitimate work in political justification — just not the work of making a policy substantively just. Majority support is strong evidence of political legitimacy in the procedural sense: that a decision was reached through a process the community has agreed to treat as authoritative, and that implementing it respects citizens’ standing as political equals whose votes count the same as everyone else’s. This procedural legitimacy is genuinely valuable, and a policy passed by fair majority vote has a real claim to being implemented that a policy imposed by a minority does not. But procedural legitimacy is a claim about how a decision was reached, not about whether its content is fair to those it affects, and the entire history of

majorities voting to entrench slavery, deny suffrage, or persecute minorities is a standing demonstration that these two things can come apart completely.

A defender of majority rule might respond that over time, and in aggregate, majoritarian decision-making tends to track justice reasonably well, because a large enough and sufficiently diverse group of people, reasoning independently, is more likely to arrive at a correct answer than a small elite — an application of the Condorcet Jury Theorem, which shows that if individual voters are each more likely than not to be correct on a binary question, and vote independently, the probability that the majority is correct approaches certainty as the number of voters grows. This is a real mathematical result, but its conditions are demanding and frequently unmet in actual politics: it requires voters to be answering the same well-defined question, to have some independent competence above chance on that question, and to vote independently rather than being swept by shared misinformation or social conformity — conditions that hold reasonably well for some questions (aggregating noisy but genuinely independent estimates of a quantity) and hold very poorly for others (questions where a false narrative has been widely and uniformly absorbed by the electorate, which produces correlated rather than independent error, exactly the condition under which the theorem's conclusion fails).

This gives us a more precise way to say what is wrong when a clear majority supports an unjust policy: it is not that democracy has malfunctioned in some rare, aberrant way that a better procedure would have prevented, but that majority support was never, even under ideal conditions, a guarantee of substantive correctness — only a statistically reasonable bet under specific epistemic conditions that frequently do not hold, particularly on exactly the questions (about the standing and treatment of minorities, about who counts as a full

member of the community) where a majority has the least first-hand stake in getting the answer right and the most structural incentive to get it wrong in its own favor.

None of this is an argument against democratic procedure, which retains real value independent of its substantive track record — respecting citizens' equal standing to participate in decisions that bind them is not merely instrumentally valuable as a truth-tracking device, and remains valuable even in the cases where the majority gets the substantive question wrong. But it is an argument against a specific, common rhetorical move: citing majority support as though it settled the question of justice, rather than settled a different, real, but more limited question about whose voice the decision-making process was required to respect. When a majority's decision is substantively unjust, the correct response is not to deny that the procedure was democratically legitimate in the procedural sense — it may well have been — but to insist that procedural legitimacy was never the same claim as substantive justice, and that citing one does not discharge the burden of defending the other.

ETHICS

Retribution Without Anger

Retributive punishment is often dismissed as institutionalized revenge. This essay argues the best version of retributivism has nothing to do with anger at all, and defends it on those terms.

Rafael Duarte Colégio das Amoreiras, Lisbon

Retributivism — the view that wrongdoers deserve punishment proportionate to their wrongdoing, independent of any further social benefit the punishment produces — is often criticized as a philosophical dressing-up of the desire for revenge, institutionalizing an emotional impulse (anger at the wrongdoer, a wish to see them suffer) rather than justifying anything a cooler, more reflective morality should endorse. This essay defends a version of retributivism that does not rely on anger at all, and argues that stripping the anger away actually strengthens, rather than weakens, the retributivist case.

The anger-based picture of retribution has real problems that its critics correctly identify: anger is notoriously unreliable as a guide to proportionate response, tends to escalate rather than calibrate, and is often directed at the most visible or available target rather than the most culpable one. If retributivism's entire justificatory foundation were "wrongdoing makes victims and communities angry, and punishment satisfies that anger," it would indeed be vulnerable to the standard consequentialist reply: satisfying anger is one psychological good among others, to be weighed against the costs of punishment (suffering inflicted, resources spent, the possibility of rehabilitation forgone) rather than treated as a freestanding moral demand.

But the more defensible retributivist tradition, developed by philosophers like Jean Hampton, does not ground punishment in the satisfaction of anger at all — it grounds punishment in the correction of a false claim implicit in the wrongdoing itself.[1] On this view, a serious wrong is not merely harmful; it is also an act of communication, implicitly asserting a false claim about the relative standing of wrongdoer and victim — that the wrongdoer's interests or will matter enough to override the victim's basic claims against being treated this way. An assault does not just cause pain; it enacts, through the act itself, the message that the victim's bodily security did not constrain the assailant's conduct the way it should have. Punishment, on this view, is justified not by anger's demand for satisfaction, but by the need to publicly and authoritatively correct the false claim the wrongdoing asserted — to vindicate, through a proportionate and consequential response, that the victim's standing was not in fact overridden, whatever the wrongdoer's act implied.

This reframing changes what proportionality is tracking. On the anger model, proportionality is calibrated to how angry the wrongdoing should make us — a psychologically unstable standard, since anger scales with vividness and salience more than with actual culpability. On the vindication model, proportionality is calibrated to how emphatically the false claim needs to be corrected in order to be genuinely undone — a more stable standard, at least in principle, because it asks a question about the wrongdoing's implicit assertion (how serious a claim about the victim's diminished standing did this act make) rather than about the psychological states of observers. This also explains a pattern the anger model struggles with: we generally think a coldly calculated, dispassionate wrong (a fraud planned over months) deserves at least as much punishment as an identical wrong committed in a moment of hot anger, even though the hot-anger case is, if anything, more likely to provoke observer anger in return. The vindication model predicts this correctly, since

the false claim about the victim's standing is, if anything, more starkly asserted by a coldly calculated wrong than by an impulsive one; the anger model, tracking observer emotion rather than the act's implicit content, has no principled reason to expect this pattern at all.

This version of retributivism is not vulnerable to the standard consequentialist reply in the same way, because it is not claiming that punishment produces a good feeling (satisfied anger) that must be weighed against punishment's costs; it is claiming that punishment does a specific piece of expressive, corrective work — vindicating the victim's standing against a false claim actually made against it — that no amount of forward-looking deterrence or rehabilitation, aimed only at future benefit, is designed to do at all. Whether this work is important enough to justify the specific costs of punishment in any given case remains a real question or, but it is a different and more defensible question than whether institutionalized anger is worth its costs, and critics of retributivism who only ever engage the anger-based version are arguing against a weaker theory than the one its best defenders actually hold.

NOTES

1. Jean Hampton, "The Moral Education Theory of Punishment," *Philosophy & Public Affairs* 13 (1984); and her work with Jeffrie Murphy in *Forgiveness and Mercy* (1988).

METAPHYSICS

Can Justice Be Impartial About Time?

Most theories of justice treat when a person exists as morally irrelevant to what they are owed. This essay asks whether that impartiality is actually coherent once future population size is itself affected by present choices.

Mireille Dupont Collège International, Geneva

A natural-sounding principle of justice holds that the mere fact of when a person exists in time should not affect what they are owed — a person born in 2200 has the same basic moral claims as a person born in 1990, and discounting future welfare simply because it is future, with no further justification, looks like an arbitrary bias, no more defensible than discounting someone's welfare because of where they happen to live. This essay asks whether this appealing principle of temporal impartiality can actually be made coherent once we admit a fact that most treatments of intergenerational justice underplay: present policy choices do not just affect how well-off future people will be, they affect how many future people there will be, and who, specifically, those people turn out to be.

The trouble is a scaled-up version of the non-identity problem: temporal impartiality, stated as "future people's welfare should count the same as present people's," implicitly assumes there is a fixed population of future people whose welfare levels are simply waiting to be affected by our choices, the way a fixed population of present people's welfare can be affected. But population size itself is a

variable under our control, at least in aggregate, through policies affecting economic conditions, healthcare, and family formation — and this means temporal impartiality has to answer a question that impartiality among presently-existing people never has to face: is it better, all else equal, for there to be more future people with good lives, or fewer future people with even better lives, or does population size not matter at all as long as some threshold of welfare is cleared?

Each answer leads to a conclusion most people find hard to accept. If we say more total welfare is always better, regardless of how it is distributed across a population whose size we can also increase, we arrive at what Derek Parfit called the Repugnant Conclusion: for any population enjoying a very high quality of life, there is some much larger population, each member living a life barely worth living, whose existence contains more total welfare and would therefore, on pure aggregation, be a better outcome — a conclusion that has struck nearly everyone who encounters it as obviously false, while proving remarkably resistant to any principled way of avoiding it without giving up something else that also seems obviously true.[1] If instead we say only average welfare matters, we get equally strange results in the other direction: a change that adds a new person with a life much better than average but slightly below the current average would count as making things worse, purely because it lowers the average, even though every actual person involved, including the new person, is thereby better off. If we retreat to a threshold view — population size does not matter as long as everyone clears some minimum bar — we need an independently defensible way of setting the bar, and any bar we set seems to reintroduce exactly the arbitrariness temporal impartiality was trying to avoid in the first place.

This is not a solved problem in population ethics, and this essay does not propose to solve it — the point is narrower and, I think, underappreciated in how intergenerational justice gets discussed outside specialist literature. The appealing slogan "future people matter just as much as present people" is not, by itself, a usable principle of justice, because it does not by itself tell us what to do about policies that change who, and how many, future people there will be, which is nearly every policy with intergenerational consequences worth discussing, from climate policy to healthcare spending to family policy. Temporal impartiality is a necessary component of any adequate theory of intergenerational justice, but it is nowhere close to a sufficient one, and treating it as though it settles practical questions about how much the present generation owes the future — rather than as the easy, uncontroversial part of a genuinely unsolved problem — understates how much philosophical work remains before we can say, with any confidence, what impartiality across time actually requires of us.

NOTES

1. Derek Parfit, *Reasons and Persons* (1984), part IV, on the Repugnant Conclusion.

EPISTEMOLOGY

The Epistemics of a Hung Jury

A jury that cannot reach a verdict is usually treated as a failure of the process. This essay asks whether persistent, well-reasoned disagreement among jurors might sometimes be the epistemically correct outcome.

Daniyar Bekov Almaty Science Lyceum, Almaty

When a jury deadlocks, unable to reach the unanimous or near-unanimous verdict most legal systems require, the standard institutional response treats this as a process failure to be corrected — through further deliberation, an Allen charge urging holdouts to reconsider, or eventually a retrial before a different jury that, it is hoped, will succeed where the first one did not. This essay asks whether this framing gets something importantly wrong: that in a meaningful subset of hung-jury cases, persistent disagreement among jurors who have each reasoned carefully from the same evidence is not a correctable malfunction but the epistemically appropriate outcome, and that treating it as a failure creates institutional pressure toward false consensus.

The relevant background is work on what epistemologists call peer disagreement: cases where two reasoners, judged to be roughly equally competent and equally exposed to the same evidence, arrive at different conclusions. A prominent view, associated with the "conciliationist" position defended by philosophers like Richard Feldman and, in different form, Adam Elga, holds that discovering you disagree with an epistemic peer is itself evidence you should weigh, typically requiring you to move your credence toward theirs,

since your peer's disagreement is itself some evidence you might be the one who has erred.[1] If conciliationism is correct, sustained disagreement among jury peers, all exposed to identical trial evidence, should indeed be somewhat rare and somewhat unstable — which fits the institutional assumption that a hung jury represents something going wrong that further deliberation should fix.

But conciliationism's force depends heavily on the peers actually being equally well-positioned with respect to the evidence, and jury deliberation routinely violates this condition in a specific, underappreciated way: jurors do not only receive the same trial evidence, they also bring different prior interpretive frameworks — different background assumptions about witness credibility, different weightings of circumstantial versus direct evidence, different thresholds for what "reasonable doubt" actually requires in practice — and these frameworks are not epistemic defects to be conciliated away, but often reflect genuinely different, individually reasonable resolutions of an underdetermined question the legal standard itself leaves open. "Beyond a reasonable doubt" does not specify a precise numerical threshold, and two jurors can each be reasoning carefully and in good faith from identical evidence while landing on different sides of a genuinely vague standard, not because one of them has made an error the other can point to, but because the standard itself does not determine a single correct answer to be converged on.

This matters because it suggests a different diagnosis for at least some hung juries: not failed convergence toward a fact of the matter both jurors are equally trying to track, but stable, legitimate disagreement about how to resolve genuine vagueness in the legal standard itself — closer to two reasonable people disagreeing about whether a barely-sunset sky counts as "red" than to two mathematicians disagreeing about a sum. Conciliationism's recommendation to move toward your peer's view makes sense when disagreement is evidence of a shared,

trackable fact one of you has likely mis-tracked; it makes much less sense when the disagreement is actually about how to resolve an underdetermined standard, where there may be no further fact for either juror's dissent to be evidence of having missed.

If this is right, the institutional pressure to eliminate hung juries — through Allen charges specifically designed to make holdout jurors reconsider their position in light of the majority's view — is applying a conciliationist logic in exactly the cases where it least applies: pressuring jurors to treat their vague-standard disagreement as though it were evidence of individual error, when it may instead be a legitimate, non-erroneous resolution of a standard that does not, on its own terms, dictate a single answer. This does not mean all hung juries reflect legitimate disagreement rather than correctable error — many surely do reflect exactly the kind of case conciliationism is right about, where one juror has simply misunderstood a piece of evidence the others tracked correctly. But it suggests the system's blanket institutional preference for eliminating deadlock, whatever its cause, may be quietly manufacturing false consensus in precisely the cases where sustained, respectful disagreement was always going to be the epistemically honest result.

NOTES

1. Richard Feldman, "Reasonable Religious Disagreements," in *Philosophers Without Gods* (2007); Adam Elga, "Reflection and Disagreement," *Noûs* 41 (2007).

ETHICS

Proportionality and the Problem of Incommensurable Goods

Justice frequently demands proportionate response, but proportionate to what, when the goods being compared share no common unit? This essay argues the incommensurability problem is more tractable than it looks.

Grace Kim Hanwoori International School, Seoul

Many judgments of justice require weighing goods against each other that appear to share no common measure: how many years of imprisonment are proportionate to a specific act of fraud; how much financial compensation is proportionate to the loss of a limb; how a nation should weigh economic growth against environmental preservation in setting policy. The philosophical worry, associated with value pluralists like Isaiah Berlin and developed further by Joseph Raz, is that these goods may be genuinely incommensurable — not merely hard to compare, but lacking any common scale along which a determinate answer to "which is greater" could even in principle be given, the way a length and a temperature cannot be meaningfully compared as "greater" or "lesser" without an arbitrary conversion imposed from outside.

If this worry is correct in its strongest form, justice's demand for proportionality looks incoherent: proportionality requires a shared scale on which to set the punishment or compensation at the correct ratio to the wrong or the loss, and if no such scale exists even in

principle, then any specific number of years or dollars we assign is not really tracking a proportionality relation at all — it is an arbitrary convention dressed up as a considered judgment.

I want to argue that this conclusion moves too fast, because it assumes commensurability requires a single, prior, independently specifiable common unit (like a shared currency existing before any particular judgment is made), when a weaker and more defensible form of commensurability is available: comparative judgments can be perfectly determinate, and rationally arrived at, even where no such prior common unit exists, as long as the comparison is made holistically, by weighing the goods against a shared, practically-oriented question rather than converting each into some third, more basic unit first. Raz himself, despite emphasizing incommensurability, allows for exactly this kind of comparison, arguing that two options can be "on a par" — not equal, not one greater than the other by a determinate margin, but roughly comparable in a way that still permits a reasoned, non-arbitrary choice between them, given the practical context the choice is embedded in.[1]

Apply this to the fraud-sentencing case. We do not need a prior conversion rate between "years of imprisonment" and "dollars defrauded" to make a comparative judgment; we need a shared practical question both goods can be weighed against — something like, what response would adequately express the community's judgment of the wrong's seriousness, deter comparable future conduct, and remain proportionate to sentences already established for comparably serious wrongs in the same legal system. This is not converting fraud and imprisonment into a common unit external to both; it is triangulating a determinate-enough answer by reference to a shared normative task both goods are being weighed in service of, and the fact that no context-independent conversion rate exists

between years and dollars does not prevent this triangulation from producing a real, defensible, non-arbitrary judgment — any more than the absence of a universal formula for combining "how talented" and "how well the piece is executed" prevents a panel of judges from reaching real, defensible, non-arbitrary rankings in a music competition.

This does not eliminate the difficulty entirely, and there will remain genuinely hard cases where the shared practical question itself is contested or underspecified, leaving real indeterminacy rather than mere difficulty. But the strong incommensurability worry, taken to its natural conclusion, would make it impossible to explain why we can recognize a grossly disproportionate sentence — life imprisonment for petty theft, a token fine for grievous injury — as clearly wrong, even without a precise conversion rate telling us the exactly correct sentence. If comparison required a prior common unit, we should have no basis for these clear judgments either; the fact that we do have a basis for them, robustly and across widely differing legal cultures, suggests comparative judgment about apparently incommensurable goods is not the impossible task the strongest version of the incommensurability thesis implies. It is a harder, more holistic task than converting currencies, but not, on inspection, a fundamentally different kind of task from the ordinary, defensible comparative judgments we make and trust in many other domains.

NOTES

1. Joseph Raz, *The Morality of Freedom* (1986), ch. 13, on incommensurability and the possibility of options being "on a par."

POLITICAL PHILOSOPHY

Is Restorative Justice Still Justice?

Restorative justice replaces punishment with reconciliation between victim and offender. This essay asks whether it can do this while still deserving the name 'justice', rather than a different, adjacent good entirely.

Tobias Reinholt Gymnasium an der Elbe, Hamburg

Restorative justice programs bring victims and offenders together, often with a trained facilitator, to discuss the harm caused, allow the offender to take direct responsibility, and agree on a process of repair — an apology, compensation, community service directed at the specific harm — in place of, or alongside, conventional punishment. Advocates present this as justice done better: more responsive to victims' actual needs (which are often not primarily about the offender's suffering), more likely to reduce reoffending, and more humane than a purely punitive system. Critics worry that restorative processes, whatever their other virtues, are not really justice at all, because they can produce wildly different outcomes for identical offenses depending on the personalities and preferences of the specific parties involved, abandoning the consistency and proportionality that justice, as opposed to therapy or mediation, seems to require. This essay argues that the critics have identified a real tension but drawn the wrong conclusion from it.

The consistency worry is genuine: two offenders who commit identical thefts might, through restorative processes, end up with very different outcomes — one victim, forgiving and interested

mainly in an apology and a small repayment, produces a light resolution; another victim, less forgiving or with a different sense of what repair requires, produces a much more demanding one for an objectively identical offense. A retributivist system, whatever its other flaws, at least aims at treating like cases alike, since the wrong itself, not the specific victim's disposition, is supposed to set the proportionate response. If restorative outcomes vary this much based on facts about the victim rather than facts about the offense, it is fair to ask whether the resulting process is really justice, in the sense of a system committed to like cases receiving like treatment, or something else — a valuable form of victim-centered repair, perhaps, but adjacent to justice rather than an instance of it.

I think the sharpest response available to restorative justice's defenders is not to deny the inconsistency, but to question whether consistency, understood this strictly (identical offense, identical response, regardless of relational context), was ever the right standard for the specific kind of wrong restorative justice is best suited to address. Conventional criminal justice treats an offense primarily as a wrong against the state or the public order, which is exactly why consistency across cases, regardless of the specific victim's wishes, makes sense: the wrong being addressed is impersonal and general, so the response should be too. But many offenses restorative justice is applied to — interpersonal harms, offenses within existing relationships or communities — are also, and perhaps primarily, wrongs against a particular person, whose relationship to the offender, whose specific needs, and whose own sense of what repair requires are not incidental noise to be averaged away, but part of the actual content of what was wronged and what repairing it requires. On this view, the fact that restorative outcomes track victim-specific facts is not a failure of consistency in a system that was supposed to ignore those facts; it is evidence the system is

tracking a different, genuinely victim-relative dimension of the wrong that offense-based, victim-blind consistency was never designed to capture and, arguably, systematically ignores.

This does not fully dissolve the tension, because plenty of offenses genuinely do have both a personal and an impersonal dimension — a theft harms a specific victim and also violates a general social norm against theft that the community has an independent interest in seeing consistently enforced, and a purely victim-relative process risks neglecting this second dimension entirely, especially where a victim's personal forgiveness might produce an outcome too lenient to serve the general deterrent and expressive functions punishment is also meant to serve. The honest position, I think, is that restorative and conventional justice are answering different questions that most offenses actually raise simultaneously — what is owed to this particular victim, and what response the general social order requires — and neither system, used exclusively, adequately answers both. A restorative process that operates entirely outside any framework tracking the general, victim-independent seriousness of the offense risks exactly the inconsistency critics worry about; a purely retributive system that ignores the particular victim's relationship to the offender and their own account of what repair requires risks the opposite failure, treating a deeply personal wrong as though it were owed only to an abstraction. The right question is not whether restorative justice is really justice, full stop, but how a system might hold both dimensions — the impersonal claim of consistent, proportionate general treatment, and the personal claim of a specific victim's actual relationship to the harm — without collapsing either into the other.

MASTHEAD

Justice Without Consensus — Issue 04

EDITOR-IN-CHIEF

Dominic Alvares

St. Xavier's Collegiate School, Kolkata

DEPUTY EDITOR, POLITICAL PHILOSOPHY

Hannah Fitzgerald

Kilbride Grammar School, Dublin

FACULTY ADVISOR, POLITICAL PHILOSOPHY

Prof. Elena Marsh

Princeton University

FACULTY ADVISOR, ETHICS

Prof. Naomi Katz

University of Toronto

REVIEWING EDITOR, EPISTEMOLOGY

Dr. Amit Sharma

University of Oxford