CHAPTER 20

Purpose and Punishment

Presupposes: the dissolution of creation ex nihilo (Ch. 9); the response-constitutive account of value and the commitment result (Ch. 14); the primary/added value taxonomy and field-enrichment (Ch. 15); the great moral conflation (Ch. 16); the unified causal model and perspectival asymmetry (Ch. 18); the evaluative constitution of purpose (Ch. 19).

Establishes: that purpose requires neither God nor immortality; the failure of divine retribution's desert basis, by two independent routes; restorative over retributive justice; autonomy-preserving universalism.[1]

In brief. Can anyone be justly punished for what they did not ultimately author? Purpose and punishment are two faces of one coin: the modality that constitutes meaning without cosmic endorsement is the one whose involuntary delivery makes ultimate accountability impossible. Against Craig: significance requires neither permanence, endorsement, nor duration — a sunset's beauty is not diminished by its ending, and immortality invites evaluative saturation — Makropulos's tedium, given its mechanism. Against divine retribution: the wrong of eternal damnation is that it is punishment — a response to a desert no agent possesses and a value no register supplies — not that it is eternal. A judge who understood the architecture would not punish; he would educate. The terminus — every soul reconciled — is proposed for the companion God alone, as an aim rather than a guarantee.

---

1 Introduction

Can anyone be justly punished for what they did not ultimately author? This chapter applies the results of the preceding chapters to two theological claims that have shaped the lives of billions: that purpose requires God, and that God is justified in punishing those who choose wrongly. The results it applies are the evaluative constitution of purpose, with its bipartite structure of object and sense, and bounded authorship as the structural condition of any possible agent.

Purpose and punishment belong together, and the connection is structural rather than merely thematic.

The same evaluative modality that constitutes purpose (Chapter 19) is the modality whose involuntary delivery constitutes the bounded authorship that undermines absolute moral responsibility (Chapter 18). They are two faces of one coin: the first shows what the evaluative modality makes possible, meaning without cosmic endorsement, while the second shows what it makes impossible, ultimate accountability for choices the agent did not ultimately author.

The chapter argues four things. Purpose does not require God's existence, cosmic endorsement, or immortality (§§2–3). The theological claim of religious dissatisfaction rests on a misunderstanding of infinity (§4). Divine retribution lacks the desert basis it requires, given the bounded authorship of any possible agent; the strongest theological defence, Swinburne's self-separation, fails on both its premises; Plantinga's free will defence presupposes a degree of freedom the structural architecture does not provide; and religious obedience admits of four structurally distinct forms, only one of which constitutes moral goodness in its own right (§5). Finally, the structural analysis supports restorative over retributive justice and, in the divine case, education over punishment altogether (§6).

2 Purpose Without God

2.1 The structural argument

Purpose is constituted by the conjunction of an object, the target toward which the agent is oriented, and a sense, the evaluative modality's emotional engagement with that target, as Chapter 19 established. None of the structural conditions mentions a deity. Purpose obtains wherever these conditions are jointly satisfied — wherever a conscious system with a functioning evaluative modality identifies an object and constitutes an oriented response toward it — regardless of whether the universe was designed, whether the differences are cosmically endorsed, or whether a creator intended the system to have purpose.

The first section of this book established that the grounding terminus is unintelligent and that creation ex nihilo is structurally incoherent (Chapters 4, 6, 9, 10). God cannot be the origin of existence from non-being, and whether or not existence depends on a God (Chapter 24, §4.1), purpose arises from the structural architecture of conscious evaluation, not from the intentions of a cosmic designer.

2.2 Craig and the absurdity of life without God

Craig's (2008) argument is the most developed contemporary version of the theistic claim and deserves sustained engagement. Craig identifies three conditions he claims are jointly necessary for life to be meaningful: lasting consequences, cosmic endorsement, and personal survival. A godless universe satisfies none. The argument is serious because it gives voice to a real evaluative orientation — a longing for permanence, endorsement, and survival — that many people share, whether or not they share Craig's theological commitments.

The first strand is the absence of lasting consequences. Craig argues that in a godless universe all human achievements will eventually be erased, by personal death, by extinction, by the heat death of the cosmos, so nothing we do has lasting significance.

The response identifies the assumption: that significance requires permanence. The assumption is itself an evaluative claim the modality does not support, since most of what matters to most people — a child's first steps, a conversation with a friend, a problem solved — matters in the present, without reference to cosmic duration, and a sunset's beauty is not diminished by its ending (Chapter 19, §7.2). Craig conflates two claims: that significance is real, and that it requires permanent objects. The first is true; the second fails.

Scheffler (2013) develops a striking thought experiment: imagining that humanity will go extinct thirty days after one's death. Scheffler argues that this prospect would drain present activities of significance far more than the prospect of one's own death, suggesting that what matters is not personal immortality but the continuation of the human project. The present account explains Scheffler's finding: the evaluative

modality constitutes significance through its engagement with non-evaluative differences, and many of those differences presuppose the continuation of evaluative beings who can register them. The "afterlife" that matters for purpose is not personal survival but the persistence of the evaluative community: the ongoing existence of beings whose evaluative modalities can register the differences in virtue of which one's engagement is significant.

The second strand is the absence of cosmic endorsement. Craig argues that without a God who values what we value, our valuations are subjective preferences projected onto an indifferent universe. The response draws on the response-constitutive account (Chapter 14), since evaluative differences are response-dependent realities on the secondary-quality model, no more unreal than colours. Craig's demand for cosmic endorsement is a demand for a further layer of derived value — not merely that the modality constitutes significance, but that a cosmic being confirms the constitution — so the demand adds a condition purpose does not require. A mother comforting a frightened child constitutes real significance through her empathic engagement, whether or not God witnesses the act.

The third strand: the inevitability of personal annihilation. Section 3 addresses this strand.

2.3 Swinburne and the value of a calling

Swinburne (2004) offers a more nuanced version: God's existence provides human life with a point it would otherwise lack, the opportunity to develop one's character in response to a divine calling. The response operates at two levels.

First, the "calling" is a form of purposive orientation, an evaluative target the modality is oriented toward. But the calling's purpose is constituted by the modality's oriented response to it rather than by its cosmic origin. A person who feels called to alleviate suffering is oriented toward a real evaluative target, the difference between suffering and its absence, and the orientation constitutes purpose

whether the calling comes from God or from the person's own developmental history.

Second, Swinburne's argument depends on what the value taxonomy (Chapter 15) calls the confusion between total value and primary value. Swinburne is right that the total value of a calling may be greater if God exists, since the derived value of divine authorship is a real addition. But the total value's being greater does not mean the primary value is absent without it.

Metz (2019) provides the most comprehensive recent analysis of the relationship between God and meaning, distinguishing several ways in which God could contribute to meaningfulness: as creator, as moral legislator, as object of a relationship, as guarantor of cosmic justice. Metz argues that the relational contribution is the most defensible. The structural account agrees that a relationship with a being of extraordinary evaluative richness would contribute real meaning, which is what the companion God of Chapter 24 provides. What it denies is that this relational contribution requires God to be the creator or the moral legislator.

3 Immortality and Evaluative Saturation

3.1 The duration assumption

A common religious argument ties purpose to the afterlife: if we all die and there is no heaven, what is the point of existence? The argument assumes that purpose requires infinite duration, and the present account dissolves the assumption. Purpose requires the evaluative modality's oriented engagement with real differences rather than the infinite continuation of that engagement. The modality operates in the present, constituting the oriented difference between better and worse, desirable and undesirable, now, and the constitution is purpose regardless of how long it continues. Duration determines how long purpose obtains.

A forced choice makes the point vivid: annihilation on one side, an eternity of complete solitude on the other. The comparison shows where worth is located. Whether a life spans a day, a thousand years, or

forever, its worth is carried by the engaged responses that fill it, and duration without them adds nothing.

3.2 Williams and evaluative saturation

Williams (1973) develops the counter-intuitive consequence of the duration assumption in his discussion of Elina Makropulos, who is granted immortality at the age of forty-two. Three hundred years later she has become intolerably bored, emotionally flat, and indifferent to experience. Williams argues that her condition is a feature of immortality rather than a contingent failure of imagination.

The value taxonomy (Chapter 15) provides the mechanism. Primary value diminishes through habituation, since the evaluative modality's oriented response to a stimulus weakens as the stimulus is repeated, because the contrast between the stimulus and the subject's prior experiential field diminishes. Added value exhausts its associations, since over infinite duration every possible derived and consequential association would eventually be formed, explored, and exhausted. Total value, the composite of primary and added, therefore approaches zero as duration approaches infinity. This is evaluative saturation: the erosion of the contrast on which evaluative engagement depends.

One premise of this argument is an empirical wager. That the supply of real novelty is bounded is a claim about the universe, since whether it affords finitely or infinitely many distinguishable states is a question of physics, entangled with the temporal disjunction of Chapter 5, and the chapter does not settle it. The saturation argument accordingly holds on the finite branch and is offered conditionally there, while §7.2 states what survives on the other branch.

Fischer (1994) provides the most sustained philosophical response, arguing that Williams's argument depends on the assumption of a fixed character, since an immortal person with sufficient flexibility could continually develop new interests and thereby avoid exhaustion. Fischer's point is stronger than a fixed-character reply allows, since memory is finite too: an immortal who forgets re-encounters, and a bounded field that clears makes room for the same difference to register

again, so bounded novelty drives total value to zero only if the memory that would recognise the repetition is unbounded, which a bounded substrate forbids. The saturation argument therefore needs either an unbounded memory, which it cannot have, or the weaker claim that survives without one: within the span a memory holds, the field fills and its novelty is spent, and across spans the immortal repeats — total value cycling, with the tedium Williams describes being the registration of the cycle as a cycle. The chapter takes the weaker claim and marks it. In a finite universe with a finite number of distinguishable states, the supply of real novelty is bounded, so infinite time will exhaust it within every span a memory can hold, which is what the weaker claim needs and all it claims.

Cottingham (2003) argues that engagement with an infinitely rich God could provide inexhaustible evaluative novelty. The answer: this is possible in principle. A being of infinite diversity could sustain eternal evaluative engagement. But this defence trades on infinite diversity, the same infinity that is incompatible with the divine simplicity classical theism attributes to God. The defence works only for a God of infinite diversity that classical theism rejects. The irony is sharp for the classical tradition, and only for it, since the open theist and the personalist do not hold the simplicity doctrine it turns on: the theological claim that immortality with God provides inexhaustible purpose is coherent only on a conception of God that the dominant traditions deny.

The physics behind the wager. The physics behind the finite-novelty premise should be named, because the premise is an empirical wager and the wager has identifiable stakes. The bound is of the Bekenstein kind: a physical system confined within finite energy and radius carries at most a finite information content (Bekenstein, 1981), so a post-mortem career realised in any bounded physical substrate has finitely many distinguishable states to cycle through. On the cosmological side, analyses of long-run information processing in an ever-expanding universe approaching a de Sitter horizon find that only finitely many thoughts can be completed (Krauss & Starkman, 2000), against Dyson's (1979) earlier argument that an open universe might permit unending, ever-slowing experience; Dyson (2001) replied that

the objection binds only digital processing, a reply not conceded, so the wager is live. On eternal-inflation or cyclical scenarios the supply of novelty may be unbounded at the global scale, and the premise then falls back on its conditional form: what saturates is a bounded experiencer's field, whatever the universe at large does. The saturation argument is therefore strongest under de Sitter futures and bounded substrates, and the theological work it does in §3 inherits that grade.

3.3 Finitude and evaluative intensity

Finitude may be a structural condition for certain forms of evaluative richness. The knowledge that time is limited makes present engagement more precious: scarcity intensifies the evaluative modality's response. A conversation that could be the last is experienced with greater evaluative intensity than one among infinitely many. This is essentially Heidegger's (1927/1962) insight about Being-toward-death: the awareness of mortality is not a defect of human existence but an existential structure that intensifies present engagement. The analysis provides the mechanism: finitude intensifies evaluative response because scarcity is itself a difference that the modality registers. Mortality is not merely compatible with purpose; it is, for finite beings, one of the conditions under which the most intense forms of purpose are constituted.

Becker (1973) argues that the denial of death is the primary engine of human culture, since mortality anxiety generates "immortality projects" designed to transcend it. The present account accommodates this, because mortality anxiety is the evaluative modality's negative response to the prospect of evaluative cessation, a real and powerful oriented difference. But the response itself is a form of purposive engagement rather than its negation, since the person who builds an immortality project is constituting purpose through her engagement with the project. What the present account adds to Becker is the observation that the anxiety is mislocated: the evaluative modality responds to the prospect of evaluative cessation as though permanence were a condition for present significance, and the analysis shows it is not.

People of faith may struggle to see it, though belief in God is neither unique nor necessary for purpose. God can serve as a dominant emotional anchor and a powerful source of meaning, while purpose can be found wherever emotional resonance arises: in family, friendship, art, science, or a walk beneath the stars. The account does not adjudicate between hierarchies, since a life oriented toward God and a life oriented toward artistic creation are both purposive, provided the modality is engaged with real non-evaluative differences.

Wielenberg (2005) argues that value and virtue are possible in a godless universe because the conditions for moral goodness do not depend on a divine source. The analysis agrees and provides what Wielenberg's argument lacks: an explanation of why the conditions are sufficient. Evaluative differences are response-constituted realities rather than brute facts, located in the relational structure between the modality and its input rather than in the input alone.

Cottingham's (2003) more moderate position is that theistic commitment provides resources for meaning — communal worship, contemplation, a narrative framework — that secular alternatives may not fully replicate. The account is partially sympathetic, on the ground Chapter 19 (§7.3) states: the resources are real, and they work because they engage the modality rather than because a deity authorises the engagement.

4 Religious Dissatisfaction

4.1 Religious dissatisfaction and the misunderstanding of infinity

"Without God there can be no true happiness." Beliefs of this form teach their holders to expect dissatisfaction, since joy is referred to the afterlife and the present life is read as incomplete by design. The expectation rests on a misunderstanding of infinity entangled with the doctrine of God's infinite nature, and the reasoning runs: "God is infinite, and I need God; therefore I must need something infinite."

But the two infinities (Chapter 8) are structurally opposite conditions. Infinite diversity is the endless extension of internal structure; infinite equality is the dissolution of all structure. If the God in question is infinitely diverse, the felt lack the belief reports is a desire for inexhaustible evaluative engagement — real, legitimate, and satisfiable wherever the modality meets a field rich enough to sustain it, with a deity as one candidate among others. If the God in question is infinitely equal, the lack admits of no divine filling at all, because zero complexity has no evaluative content to engage.

The remedy is therefore not the acquisition of faith but the cultivation of the engagement that faith, at its best, was always meant to provide: deeper empathic connection, richer aesthetic experience, more purposive orientation toward the real non-evaluative differences that populate one's life. Whatever sits at the top of a person's evaluative hierarchy — God, family, justice, creative work, or self-respect — becomes her dominant purpose (Chapter 19, §3.2). God can serve as a powerful dominant purpose, while the structural conditions for purpose are satisfied by any engaged hierarchy rather than only a theistic one.

If purpose does not require God, what about justice? The Abrahamic traditions hold that God is the ultimate moral judge: that the righteous are rewarded and the wicked punished, if not in this life then in the next. This section applies the unified causal model and the perspectival asymmetry (Chapter 18) to this doctrine.

Where the argument stands: purpose has been shown to require neither God nor cosmic endorsement nor immortality (§§2–3), and the felt lack that religious dissatisfaction reports has been traced to a confusion between the two infinities (§4). What follows turns to punishment, and asks what desert a divine judge could be answering to.

5 Divine Retribution

5.1 The structural argument

Chapter 18 established three results. The evaluative field is populated by subconscious processes consciousness does not control,

for any conscious system, physical or non-physical. Intentional behaviour shares the chain architecture of unintentional behaviour, differentiated only by the evaluative modality. And freedom is phenomenologically real from the first-person present while metaphysically determined from the third-person retrospective perspective.

Divine retribution presupposes absolute moral responsibility (Chapter 18, §7.2): complete control over the factors that determined the choice, over the evaluative hierarchy that populated the deliberative process rather than only over the process itself. That is a presupposition no secular philosopher holds, so what the chapter delivers against the live theistic defences is the luck objection in Pereboom's (2005) form together with the value route of §6.

The value taxonomy (Chapter 15) sharpens the diagnosis. The choice is determined by total value, the composite of primary and added, and the agent chose neither component, since the primary responses were constituted involuntarily by the modality and the added-value profile is a product of developmental history and involuntary association. The total value that determines her choice is therefore the product of an evaluative architecture she did not design.

Adams (1999) argues that "horrendous evils", sufferings so severe that they threaten to render the sufferer's life meaningless, present a challenge to any theodicy grounded in proportionate response to free choices. Adams's analysis converges with the bounded-authorship thesis, because if some evils are so severe that they overwhelm the agent's capacity for evaluative engagement, then the choices the agent makes under their influence are products of an evaluative hierarchy distorted by conditions she did not create.

Consider two people raised in radically different environments: one in a loving religious household where God is associated with warmth, security, and community, the other in an abusive religious household where God is associated with fear, punishment, and control. Both "freely choose" to accept or reject God, while the evaluative architectures from which they choose are radically different, and neither

person authored her architecture. Holding both agents to the same standard of eternal accountability ignores the structural conditions under which the choices were made.

5.2 Conditional love

Freedom under sanction. It is often said that God has left us free to accept or reject Him. But a freedom whose every exercise is met with judgement is freedom only in name, since where one option is answered with eternal sanction the agent is not so much free to choose as free to incur the penalty.

The deeper objection, however, is not that the stakes render the freedom illusory but that the sanction has no warrant to begin with. Retributive punishment, the imposition of suffering because it is deserved, presupposes that the agent is the ultimate author of the self from which the wrongful choice issued, and §5.1 and Chapter 18 deny that any agent is. The wrong of eternal damnation is therefore not, in the first instance, that it is eternal; it is that it is punishment. The absence of the authorship that desert requires withdraws the warrant for retribution at every magnitude, so the infinite form merely extends without limit a sanction that was unfounded at any.

The same analysis reaches the claim that divine love is unconditional. A regard that is withdrawn on the outcome of a judgement the agent lacked the ultimate authorship to meet is conditional in the strict sense, since it is extended on a condition the agent was never positioned to control. To make eternal punishment the consequence of falling short is to make acceptance contingent on a performance that an involuntarily constituted hierarchy may or may not have allowed, which is regard conditioned on an outcome the agent did not author rather than unconditional regard. If hell is on the table, the love that tolerates it is conditional, where the condition is one no agent meets, or fails to meet, by anything that traces to him.

The point does not reach every version, and does not need to. On the self-exclusion models (§6; Kvanvig, 1993; Walls, 1992) nothing is withdrawn, since the love continues to be offered and the damned

decline it, so the regard is unconditional in the sense the objection requires; escapism (Buckareff & Plug, 2005) sharpens the model by keeping the door open without end, so that the offer of reconciliation is never withdrawn. What those models cannot decline is the prior argument. A soul that turns away turns away from within a hierarchy it did not author, and the suffering consequent on the turning is not thereby deserved — which is why the wrong of damnation is that it is punishment rather than that it is imposed.

An omniscient being who understood the structural architecture of consciousness, seeing the unified chain from the third-person retrospective perspective with perfect clarity, would recognise that the desert retribution presupposes — an ultimate authorship of the self — belongs to no conscious agent. The infinite stakes of divine retribution, eternal punishment for wrong choices and eternal reward for right ones, are on this view unfounded rather than merely excessive: not a proportionate response overshot, but a response to a desert that is not there.

The strongest theological defence reconstructs eternal punishment not as externally imposed torture but as the natural consequence of free rejection of God. Swinburne (1998) argues that hell is the condition of a soul that has definitively turned away from the source of all goodness. God does not send anyone to hell; the agent sends herself. The punishment is the natural consequence of the rejection, just as loneliness is the natural consequence of pushing away everyone who loves you.

The self-separation defence. Swinburne's defence is the most sophisticated available, and the analysis must engage it at its foundation. The argument depends on two premises: that the agent's rejection of God is free in the full sense, and that the consequences of this free rejection are proportionate.

The unified causal model engages the first premise directly. The choice to reject or accept God is itself a product of the agent's evaluative hierarchy: her emotional responses to the religious evidence she has encountered, her developmental history with religious authority, her subconsciously constituted orientation toward or away from the

divine. The two-environment illustration (§5.1) makes the point vivid, since a person raised in a loving religious household and a person raised in an abusive one "freely choose" from radically different evaluative architectures. Both choices reflect real evaluative hierarchies, and both are real at the phenomenological level, while neither person authored the hierarchy from which the choice flows.

The perspectival asymmetry (Chapter 18) identifies the mechanism. The rejection is phenomenologically free, since the agent really is evaluating, and metaphysically determined, since it is produced by a chain of evaluative forces she did not initiate. Swinburne's defence operates at the first-person level, where the rejection is free, and draws consequences at the metaphysical level, where eternal separation is said to be proportionate, so the two levels deliver divergent reports.

Eternal consequences for such a choice therefore answer to a desert the choice cannot carry, as a sentence premised on the offender's authorship of his own will is undercut once that authorship is shown to be absent. Duress is the familiar, local case in which conditions outside the agent's making defeat the desert a sanction would require; bounded authorship is the general case, obtaining for every agent and every choice.

The universality objection. The generalisation invites a stock compatibilist reply; it should be met where the analogy is used. A mitigating condition that obtains for every agent and every choice, the reply runs, mitigates nothing — a currency issued to everyone is worth nothing — so bounded authorship, being universal, cannot do the exculpatory work that duress, being exceptional, does.

The reply mistakes the logic of the two notions it compares. Mitigation is comparative and admits of a curve; desert is not. Duress defeats desert in the local case not because it is rare but because it severs the premise a retributive sentence requires, the agent's authorship of the will from which the act issued, and a severing that turns out to be universal does not restore the premise; it universalises its failure. What is graded on a curve is practical responsibility, which the analysis preserves intact and which continues to distinguish the coerced from the

willing, the reflective from the impulsive (Chapter 18, §7.1). Basic desert is pass–fail, and bounded authorship fails every agent alike, which is not the inflation of an excuse but the discovery that the sentence's premise was never available to be satisfied.

The second premise, proportionality, is undermined by the same analysis. Even if the agent's choice is "hers" in the practical sense that it reflects her evaluative hierarchy, the hierarchy's constitution is bounded, and the ethereal mind argument (Chapter 18) shows this holds for any possible agent, physical or non-physical. Eternal consequences for a choice produced by bounded authorship therefore answer to a basic desert that bounded authorship withholds, regardless of how sincere or reflective the choice appeared from the inside.

The problem of evil in its traditional form, the tension between divine omnipotence, omniscience, omnibenevolence, and the existence of suffering, does not arise if there is no omniscient, omnipotent, omnibenevolent creator at the origin of existence. But Plantinga's (1974) free will defence deserves engagement on its own terms, because the analysis provides a critique that goes beyond the standard objections.

Plantinga argues that a world containing creatures who freely choose good is more valuable than a world of automata, and that God cannot create free creatures while guaranteeing they always choose well. The suffering in the world is therefore the cost of a freedom God could not eliminate without destroying the value it produces.

The analysis targets Plantinga's central premise — not that free will exists, since it agrees that choice is phenomenologically real, but that the creatures' freedom is of the kind that justifies the cost. If the agent's evaluative hierarchy was involuntarily constituted by processes she did not control, then the disobedience Plantinga invokes is the product of an evaluative architecture the agent did not design rather than of the absolute free choice the defence requires.

The placement matters, because the objection belongs to a lineage the defence has met before. Pereboom (2005) argues that rejecting libertarian free will removes the free-will theodicy, since creatures who

lack that freedom cannot carry the cost the defence asks them to bear, and the present analysis reaches his conclusion by its own route.

Where it differs from the standard incompatibilist reply, which denies free will under determinism, and from the logical problem of evil, which finds a contradiction among God's attributes and suffering, is in the ground it occupies. It accepts the reality of choice — the agent chose, and the choice reflects her hierarchy — while denying that this reality suffices for the absolute responsibility the proportionality claim requires, because the hierarchy from which the choice flows was not itself chosen. This is the hard-incompatibilist point of Chapter 18 brought to bear on the defence, with the substrate-generality of the ethereal mind argument closing the libertarian escape routes a determinism-based version leaves open. The "risk" Plantinga identifies, that God cannot guarantee free creatures will choose well, is therefore not a risk freely accepted by the creatures; it is a risk imposed on creatures whose evaluative hierarchies were constituted by processes they did not control.

Van Inwagen (2006) has developed the "expanded free will defence", on which the first human beings, in turning from God, lost the protection they had enjoyed from the hazards of the natural order, so that what human beings suffer from natural causes traces to a misuse of free will as moral evil does; Plantinga (1974) had allowed that natural evil might be the work of fallen angels. The analysis applies to both expansions as it applies to Plantinga's original, since the evaluative hierarchies of the first humans and of the angels were themselves involuntarily constituted, and each "rebellion" flowed from a hierarchy its bearer did not author.

Stump (2010) offers a theodicy that shifts the ground from freedom to love: suffering is the occasion for redemptive love rather than the cost of freedom. The engagement is more sympathetic than with Plantinga, because Stump grounds the justification in empathic concern. But two problems remain. Stump presupposes a creator God who designed the world in which suffering occurs, and whether a person responds to suffering with deepened compassion or embittered withdrawal depends on her evaluative hierarchy, which was not authored by the sufferer.

The insight that suffering can, under the right conditions, deepen empathic engagement is preserved; the theological claim that this deepening was divinely orchestrated is not.

Four forms of obedience. If morality is constitutively empathic (Chapter 16), what is the moral status of religious obedience, compliance with divine commands out of reverence for God? The three-goodness taxonomy (Chapter 14) provides the diagnostic.

Empathic-alignment obedience occurs when the agent obeys because she empathically grasps the command's content: she understands why cruelty causes suffering, and her evaluative modality responds to that suffering. The command and the empathic response converge. This is the best form of obedience, since it combines orthotic compliance with empathic engagement.

Fear-based obedience occurs when the agent obeys because she fears divine punishment. The preference for cruelty is overridden by fear rather than transformed. This is shallow persuasion (Chapter 21): the underlying preference survives, suppressed by a stronger emotional response.

Reward-based obedience occurs when the agent obeys because she seeks divine reward: paradise, blessing, salvation. The evaluative engagement is with the reward, not with the persons the command affects. The structure is identical to fear-based obedience: the preference is overridden, not transformed.

A fourth form must be added, because it is the one the tradition itself puts first. Love-of-commander obedience occurs when the agent obeys from love of God, neither to escape punishment nor to secure reward, but because the beloved asks. The tradition is right that this differs in kind from fear and reward, since it is empathic in structure — a real oriented response, engaging and transformative where fear and reward merely suppress — and nothing in the account denies its reality or its beauty.

What settles its moral standing is its object. The response is directed at the commander's state — His will, His pleasure, the relationship —

rather than at the persons the commanded conduct affects, while morality's subject matter is conduct's bearing on those it affects (Chapter 16, §2.1).

Love-of-commander obedience therefore inherits its moral credit from the alignment of the command's content: where God commands kindness, the lover of God does good at one remove. Where command and cruelty coincide, love of the commander is the devoted perpetrator's structure — the compliant perpetrator of Chapter 16 (§5.2) with love in place of procedure, the most sympathetic member of that family and still within it. The nobility of the motive is real; it is not the moral fact.

Only empathic-alignment obedience constitutes moral goodness in its own right, since love-of-commander obedience is moral where, and because, the commands it follows are empathically aligned. A theology that evaluates persons by their obedience — whether the obedience is empathically aligned, love-based, fear-based, or reward-based — is conflating orthotic compliance with moral goodness. This is the great moral conflation (Chapter 16) in its most intimate religious form.

6 Restorative Justice and the Empathic Framework

Two independent routes. Retributive justice holds that wrongdoing deserves punishment proportionate to the wrong, so the punishment is backward-looking. The analysis faces retribution with more than a limit on its degree, meeting it on two grounds that do not lean on each other.

The first concerns the value the suffering is said to have. Basic desert — the non-instrumental desert that would make suffering owed for its own sake — claims that the wrongdoer's suffering is good in itself, and the metaethics has a test for claims of that shape: ask who it is good for. Developed below, the test returns the vacancy result for the response-independent reading and the anchor's verdict for the response-dependent one, and it uses nothing the applied sequence does not already rest on (Chapter 14).

The second concerns the responsibility the suffering is said to answer to. Basic desert requires the ultimate authorship of the self that Chapter 18 denies to any agent, on the sourcehood premise Chapter 18 (§8.1) names as contested. Either ground removes retribution's basis, so a reader who refuses the sourcehood premise meets the first ground unchanged. What is withdrawn is therefore not a quantum of proportionate punishment, leaving a smaller measure intact, but the desert basis as such. Practical responsibility is preserved, and with it the forward-looking response examined below; backward-looking retribution is not.

Restorative justice holds that the appropriate response is the repair of the relationship between the wrongdoer, the victim, and the community rather than punishment. The response is forward-looking, since it aims to transform the wrongdoer's evaluative hierarchy so that the preference for the wrongful conduct is itself changed — deep persuasion (Chapter 21) applied to the specific conflict the wrongdoing created.

Duff (2001) develops a communicative theory of punishment, on which punishment is justified as a form of moral communication: the community's expression to the offender that her conduct was wrong, inviting recognition and repentance. Tadros (2011) argues that punishment is justified by the duty a wrongdoer incurs to protect her victim, and others, from further harm. Both accounts are structurally aligned with the empathic framework, since Duff's communicative punishment is deep persuasion, engaging the offender's evaluative modality through moral communication, while Tadros's victim-centred approach prioritises the empathic relationship between offender and victim over the orthotic satisfaction of a retributive standard.

A third family should be named as surviving untouched: quality-of-will accounts, on which blame responds to the meaning of the agent's conduct, the ill will or indifference it expresses, without invoking basic desert at all (Scanlon, 2008; McKenna, 2012). The value route's test must be run on this family too, or the instrument cuts unevenly: blame is unpleasant to its object, so a blamee's registration returns "bad" as a sufferer's does. What distinguishes blame that survives from suffering

that does not is that blame is address — a communication whose point is the blamee's uptake, and which succeeds as blame only if it is understood — whereas retributive suffering is imposed for its own sake and succeeds whether or not anything is understood. The first is a registration offered to a register; the second is a registration inflicted on one. The test "good for whom?" returns, for blame, the practice of address and both its parties, since address is something done with the blamee and not to her, and for retribution, the persecutor alone; the distinction is drawn against the instrument, and the family survives on it. Communicative, victim-centred, and quality-of-will accounts neither need the desert eternal retribution requires nor lose anything to its removal. The lineage runs to P. F. Strawson (1962), whose reactive attitudes ground the practice, with Darwall (2006) supplying its second-personal grammar; both survive on the same terms.

One further family sits at the fork the two routes create, and its placement completes the taxonomy. The fair-play tradition (H. Morris, 1968) grounds punishment in a distribution: the law-abiding carry the burden of self-restraint, the offender takes the benefits of the system while shedding the burden, and punishment restores the balance.

Read as a fairness rationale, the account uses no basic desert, since the debt it collects is anchored in the registrations of those who bore the restraint, and the response it licenses is answerable to the same mutually informed field as every other claim about burdens (Chapter 15, §4.4). It therefore survives beside the communicative, victim-centred, and quality-of-will accounts, at whatever weight the field assigns it.

Read as a desert rationale — the offender's suffering owed because taking the advantage made suffering his due — it falls to both routes above, since the owing is the vacant clause under a distributive description and the taking issued from a hierarchy the taker did not author. Berman (2011) draws the line within retributivism that this fork tracks, separating consequentialist versions, on which deserved suffering is a good the state has reason to bring about, from non-consequentialist versions, on which the desert claim does its work without routing through the goodness of outcomes, and Husak (1992) argues from inside the retributive literature that desert alone cannot

carry the justification of punishment — a conclusion the present analysis reaches with the further claim that nothing else restores what desert was to supply.

Braithwaite (1989) provides the most influential criminological evidence: reintegrative shaming, which engages the wrongdoer's empathic modality rather than merely inflicting pain, produces lower recidivism than retributive punishment. The structural prediction aligns with Braithwaite's empirical finding, since restorative approaches that engage the offender's evaluative modality through empathic confrontation with the victim's suffering produce more durable behavioural change than punitive approaches that merely override the offender's preferences through fear.

The structural analysis supports restorative over retributive justice on three grounds. It is aligned with deep persuasion, aiming to transform the preference rather than override it. It is compatible with bounded authorship, presupposing no absolute control over the hierarchy. And it engages the modality in its empathic deployment, bringing the wrongdoer into contact with the victim's suffering so that the hierarchy is transformed from within.

The strongest retributivism. This leaves the strongest form of retributivism, which rests on neither deterrence nor reform: the claim that the guilty simply deserve to suffer, that their suffering is good in itself and owed independently of any further end (Kant, 1797/1996; M. S. Moore, 1997). The forward-looking case for restoration does not touch this position, which grants that punishment serves no purpose and embraces it regardless.

It is met instead at two points, each sufficient on its own. First, basic desert of this non-instrumental kind requires the ultimate authorship of the self that the bounded-authorship thesis withholds (Chapter 18), so there is no desert for retribution to answer to. Second, "deserved suffering is good in itself" faces the fork on which the metaethics is built. Read response-independently, good in a way that would obtain apart from every response, the claim is the vacancy

result's target by name (Chapter 14, §3.6): a clause marking no contrast between worlds.

Read response-dependently, the claim does not vanish but changes hands. The retributivist's satisfaction at the wrongdoer's suffering is a real oriented response, and the account does not deny it; it weighs it. A response that takes another's suffering as its object with positive orientation is the persecutor's structure (Chapter 16, §6.3), and what it constitutes is what every such response constitutes: that the suffering is good to the one approving.

What it cannot constitute is that the suffering is owed, because the moral fact about the sufferer is anchored in the sufferer's own registration (Chapter 14, §6.4), and hers returns the opposite verdict. The intrinsic-desert retributivist must therefore locate the good of punishment either in a value the vacancy result has shown to be unconstituted, or in an orientation the anchor has already outranked — the persecutor's, offered as a theory of justice.

The response-dependent reading also admits of a test the account has used before. The informed-registration anchor (Chapter 14, §6.5) asks whether a registration survives accurate acquaintance with its object, and the retributivist's satisfaction is a registration whose object is the wrongdoer's suffering. Does it survive the victim's full registration of that suffering?

The restorative-justice trials have put the question directly. Across the controlled comparisons reviewed by Sherman and Strang (2007), victims who met their offenders face to face reported less desire for violent revenge, and more satisfaction with justice, than victims processed conventionally. The satisfaction at suffering thinned under acquaintance with the sufferer, which is what the account predicts of a registration whose object it had not fully registered. Where a victim's satisfaction does survive acquaintance, the account does not dismiss it; it weighs it as one registration against the sufferer's, in the mutually informed field of Chapter 15, and what such fields settle on is the chapter's evidence for restoration over retribution: two screens and no third, rather than a scale that reads "deserved".

Nor can he retreat from constitution to commitment, to the claim that desert-talk at least means the response-independent owing, since the commitment result (Chapter 14, §A.3) forecloses that retreat for "owed" as for "wrong": a vacant clause conditions nothing, so the discourse of desert is answerable only to its contrastful remainder, which the anchor has already adjudicated.

The case against retribution accordingly stands on two routes, and the relation between them should be stated.

	The value route	**The authorship route**
Rests on	The response-constitutive account of value: the vacancy result (§3.6), the anchor (§6.4), the commitment result (§A.3) — nothing the applied sequence does not already use (Chapter 14)	Bounded authorship and its sourcehood premise (Chapter 18)
What it removes	The value the suffering is said to have: a suffering good for no register is, in the model of Chapter 14 (§A.5), a field with no body	The responsibility condition: no agent is the ultimate source of the self its desert requires
If the other route fails	Were ultimate authorship available, the responsibility condition would return and the value claim would stand where it stood	Were the response-independence clause contentful after all, this route would still deny the source
Whom it meets	The compatibilist about responsibility, who refuses the sourcehood premise	The retributivist who grants bounded authorship — met by both routes

Each route removes the desert basis by itself. What survives the surrender of one premise is the result by the other route, and only the surrender of both would give retribution its basis back.

The nearest development in the literature reaches the verdict by the second route alone. Caruso (2021) argues from scepticism about basic desert to the rejection of retributive punishment, with a public-health model supplying the forward-looking remainder. The convergence is recorded together with the difference: the value route carries the result here independently, so the verdict survives a reader who keeps ultimate sourcehood, while the surviving practices are those the chapter names — communicative, victim-centred, quality-of-will, restorative — rather than a quarantine model alone.

The theological placement. Traditions that propose finite, remedial punishment — purgatory in Catholic theology, Gehenna as purification in some Jewish traditions, temporary punishment in some Islamic traditions — stand closer to restorative justice than traditions that propose eternal retribution, as do the self-exclusion models on which the damned are not punished but withdraw from God of their own accord (Kvanvig, 1993; Walls, 1992; cf. §5.2).

But the bounded-authorship thesis does not challenge merely the infinite escalation of retributive consequences, as though finite retribution were sound and only its unbounded form excessive. It challenges retributive desert as such. The target of the argument is the eternal retributive hell of the strong tradition; annihilationist and remedial models, which abandon eternal retribution, already fall outside it; and the model toward which the argument's own logic points is neither, but universalism — the eventual reconciliation of every soul.

The universalism the argument points to has a theological lineage of its own, and the convergence should be recorded together with the divergence of route. Within the Christian tradition the position runs from Origen's apokatastasis to its contemporary defenders. Talbott (2014) argues that a God who loves each creature cannot will the final loss of any, so eternal damnation contradicts the tradition's own first premise. D. B. Hart (2019) argues that an eternal hell is morally

unintelligible on any candid accounting of finite guilt, and that a creation purchased at the price of one soul's endless torment could not be called good.

The present argument reaches the shared destination from outside the tradition's premises, needing neither the perfection of divine love nor the moral grammar of creation, only bounded authorship and the educative means developed below — a convergence of differently premised routes of the kind this book treats as evidence of structure. And the theological universalists retain the creator God whose coherence Chapter 9 dissolves, so their conclusion survives here with its ground exchanged.

Why the divine case differs. For the divine case in particular, even the forward-looking justification that survives for human institutions falls away. Human justice retains restorative and protective measures because human agents cannot reach into a wrongdoer and remake the hierarchy from which the wrong issued; they must work through sanction, constraint, and the slow labour of moral communication. An omnipotent God is under no such limitation. What restoration aims at — the transformation of the evaluative hierarchy so that the preference for the wrong is itself changed — such a God could accomplish directly, and without suffering.

The decisive qualification is that this is education, with no part for coercion. To overwrite a soul from without, replacing its responses wholesale, would destroy the very autonomy that makes it a person; that is brainwashing. It is what a God who valued persons would not do.

Autonomy-preserving transformation is instead deep persuasion at its limit: a working through the soul's own evaluative responses, so the change is one the soul undergoes and comes to own rather than one imposed upon it. Bounded authorship is what makes this coherent against the free-will defence, since there is no libertarian self-authorship for such transformation to violate, and divine education is then one more formative input to a hierarchy that was never self-originated, differing from ordinary moral formation only in its perfect

responsiveness. A God so placed has no remaining reason to impose suffering at all, so punishment, finite or infinite, is gratuitous.

The constructive consequence is conditional. If souls survive death and are not punished, what remains is for them to be brought to peace, and the only autonomy-respecting route to that peace is education: the gradual turning of every soul toward reconciliation. That reconciliation is not evaluative quiescence but engagement with a being of inexhaustible diversity (Chapter 24), the one form of unending life §3 found immune to evaluative saturation. This is the universalism the argument points to — not salvation imposed, but the eventual reconciliation of all. A divine judge who understood the architecture would educate rather than punish.

Containing the consequence. One consequence of the argument must be contained before anything is banked. If an omnipotent God could bring every soul to reconciliation by autonomy-preserving education, the question why the education waits — why any pre-mortem suffering exists at all — arises at once, since the same grant that makes punishment gratuitous makes the world's suffering unaccounted for. Nothing in this section is proposed on that grant.

The chapter's dialectical position should therefore be stated with care. Omnipotence is granted arguendo, to the theist who asserts it, and the omnipotent educator is a reductio addressed to him: if there is an omnipotent God, punishment is without warrant, because a better means is by hypothesis available. The reductio proposes nothing; it removes the retributive theist's footing on his own assumptions, and the problem of evil the grant reopens stays with the assumption's owner, since the chapter has taken nothing from it.

On the framework's own theology the grant is taken up by one survivor and declined by the other. The sustainer of Chapter 24 (§4.1) holds the causal order in being and has the means the conditional assumes, so its force falls on him undiminished. The companion God has no such means, since he cannot remake hierarchies or the natural order at will, and he answers for no suffering he lacks the power to prevent; his educative relation is persuasion across time rather than

transformation on demand (Chapter 24, §6.5). The universalism this section proposes is accordingly the companion's alone. It is conditional on the survival of persons beyond death, and, since the companion cannot remake a hierarchy on demand, it is a hope with a mechanism rather than a prediction: restoration as what such a God keeps working toward across time, never an outcome he guarantees (Chapter 24, §§5.2, 6.5).

7 Objections and Replies

7.1 The cosmic significance objection

Objection. Local purpose is not cosmic significance. The theist's claim is not that people without God cannot care about things but that their caring has no ultimate grounding: that in a godless universe, even the deepest engagement is a brief flicker in an indifferent cosmos.

Reply. The desire for cosmic significance is real and should not be dismissed. But significance is constituted by the modality's engagement with real differences rather than by a further layer of cosmic endorsement, so demanding endorsement from the cosmos adds a condition the account shows unnecessary. The desire itself is the modality's oriented character reaching beyond its structural conditions, where the reaching is itself evaluative engagement, and therefore a form of purpose.

7.2 The Williams counterargument

Objection. An immortal being who develops radically new interests need not experience evaluative saturation.

Reply. The question is whether real novelty is inexhaustible, which is §3.2's question and carries §3.2's answer: on the finite branch the supply is bounded and infinite time exhausts it. Fischer's response is strongest in a universe of infinite diversity and weakest in a finite universe approaching heat death.

The reply, carrying the hedge of §3.2, is accordingly conditional. On the finite branch saturation wins and the flexible immortal runs out

of world. On the infinite branch Fischer wins the duration question, and even there what survives for the chapter is the conditional Craig must deny: purpose requires engagement rather than permanence, so mortality threatens nothing that infinite duration would rescue.

7.3 The proportionality defence

Objection. Even if freedom is bounded, proportionate punishment is still possible. Finite punishment for finitely free choices may be perfectly just.

Reply. The objection conflates retributive desert with forward-looking sanction. What bounded authorship withdraws is the desert basis of retribution as such (§6); finite retribution is no better grounded than infinite, only smaller. What is coherent is the forward-looking, restorative response that does not rest on desert, and in the divine case not even that (§6). Practical responsibility is preserved throughout; retributive desert, at any magnitude, is not.

7.4 The mercy objection

Objection. Theological traditions incorporate divine mercy alongside divine justice, accommodating bounded authorship through doctrines of grace.

Reply. Traditions that qualify retribution with mercy have already moved in the direction the analysis recommends, since they have recognised, implicitly, that absolute responsibility is too strong a standard. The analysis provides a principled foundation for the theology of mercy: if absolute responsibility is structurally unavailable, then mercy is not a divine concession to human weakness but the only just response to the conditions under which human beings actually make choices.

7.5 The moral motivation objection

Objection. Without divine punishment, what motivates moral behaviour? Fear of God may not be the noblest motivation, but it is effective.

Reply. The objection confirms the chapter's analysis. Fear of divine punishment is shallow persuasion: it overrides the preference for wrongdoing without transforming it. Deep persuasion (the transformation of the preference through empathic engagement with the suffering wrongdoing causes) is more stable and more morally adequate. A society that relies on divine fear for moral behaviour is a society in managed conflict, not one in structural peace — the condition in which the conflict has been resolved rather than managed (Chapter 21).

7.6 The practical value of religion objection

Objection. Religious traditions provide practical resources for meaning that secular alternatives struggle to match. Dismantling the theological justification may leave people without the infrastructure that sustains purpose.

Reply. The chapter agrees that religious traditions provide valuable practical resources and does not argue for their dismantlement. The practical resources work because they engage the evaluative modality, so they can in principle be preserved even if the theological metaphysics is revised. What the chapter denies is that the resources derive their value from a divine source rather than from their capacity to engage the evaluative modality with real differences.

7.7 The religious experience objection

Objection. A feeling of divine forgiveness is evidence of divine mercy.

Reply. A feeling is evidence only of a feeling: a vivid experience of divine forgiveness is evidence that the agent experienced being forgiven (a real and potentially transformative evaluative event) not that God exists, forgives, or grounds the forgiveness. The transformation is real; the theological interpretation is a further claim the feeling itself does not establish.

8 Conclusion

The chapter yields three results, all constructed.

Purpose without God or immortality. Craig's three strands each fail because they add conditions the analysis shows are not required: permanence, endorsement, and infinite duration. Swinburne's "calling" is a purposive orientation whose primary value is independent of its cosmic origin. Williams's argument converges with the account, since evaluative saturation is the mechanism that explains the tedium of immortality — and the irony stands: sustaining purpose through immortality would require a God of infinite diversity, the very infinity incompatible with classical theism's simple God.

The failure of divine retribution's desert basis. Divine retribution presupposes absolute moral responsibility, a form of freedom Chapter 18 shows unavailable to any possible agent, and a value in the sufferer's suffering that no register supplies (Chapter 14); either lack removes its basis.

The Swinburne self-separation defence fails on both premises, since the rejection is phenomenologically free while metaphysically determined, while the eternal consequences answer to a desert no agent possesses — the defence failing not because the penalty is too large but because none is owed. Plantinga's free will defence presupposes a degree of freedom the architecture does not provide, which is the hard-incompatibilist point his existing responses do not address. Religious obedience admits of four structurally distinct forms, and only empathic-alignment obedience constitutes moral goodness in its own right, love-of-commander obedience borrowing its standing from the alignment of what is commanded.

Restorative justice and universal reconciliation. The structural analysis supports restorative over retributive justice, and in the divine case dispenses with punishment altogether. What is abandoned is not a measure of retribution but its desert basis, since even an omnipotent agent who could transform a hierarchy by autonomy-preserving education would have no warrant to punish at all — the wrong of

eternal damnation being that it is punishment rather than merely that it is unending.

The chapter's constructive terminus is the educative reconciliation of every soul, universalism of the autonomy-preserving kind, proposed for the companion God alone: conditional on survival, and offered as an aim that God pursues by persuasion across time rather than an outcome he can guarantee (§6). This stands on two independent routes, the value route through Chapter 14 and the authorship route through Chapter 18, each of which removes the desert basis by itself, so only the surrender of both would give retribution its footing back.

If purpose does not require God, and divine retribution has lost its desert basis, what happens when these insights meet the conditions for democratic peace? The next chapter develops the three modes of conflict resolution — enforcement, shallow persuasion, and deep persuasion — and argues that only deep persuasion produces structural peace.

Grades and Dependencies

1. Conditions. The first three results are *constructed*; universalism is *proposed* for the companion God alone: conditional on post-mortem survival, and an aim pursued by persuasion across time rather than a guaranteed outcome (Chapter 24, §§5.2, 6.5); the omnipotent educator of §6 is a reductio addressed to the retributive theist and proposes nothing. The retribution case stands on two independent routes — Chapter 14 (vacancy, anchor, commitment) and Chapter 18 (bounded authorship) — each sufficient; the value route carries the result for a reader who refuses the sourcehood premise, and the victim-satisfaction finding (§6) is an empirical stake. The saturation argument is an empirical wager (§3.2). The Premise Register carries the chain.

2. What survives the desert argument (§6). Removed: basic desert, hence retribution at any magnitude. Untouched: practical responsibility; communicative (Duff), victim-centred (Tadros), quality-of-will (Scanlon, McKenna); fair-play on its fairness reading (H. Morris), its desert reading falling with the routes. Outside the target: annihilationist and remedial models. Met at the prior step: self-exclusion models — the suffering still

undeserved. Converging by another route: Talbott and Hart, ground exchanged.

References

Adams, M. M. (1999). *Horrendous evils and the goodness of God*. Cornell University Press.

Becker, E. (1973). *The denial of death*. Free Press.

Bekenstein, J. D. (1981). Universal upper bound on the entropy-to-energy ratio for bounded systems. *Physical Review D*, 23(2), 287–298.

Berman, M. N. (2011). Two kinds of retributivism. In R. A. Duff & S. P. Green (Eds.), *Philosophical foundations of criminal law* (pp. 433–457). Oxford University Press.

Braithwaite, J. (1989). *Crime, shame and reintegration*. Cambridge University Press.

Buckareff, A. A., & Plug, A. (2005). Escaping hell: Divine motivation and the problem of hell. *Religious Studies*, 41(1), 39–54.

Caruso, G. D. (2021). *Rejecting retributivism: Free will, punishment, and criminal justice*. Cambridge University Press.

Cottingham, J. (2003). *On the meaning of life*. Routledge.

Craig, W. L. (2008). The absurdity of life without God. In *Reasonable faith* (3rd ed., pp. 65–90). Crossway.

Darwall, S. (2006). *The second-person standpoint: Morality, respect, and accountability*. Harvard University Press.

Duff, R. A. (2001). *Punishment, communication, and community*. Oxford University Press.

Dyson, F. J. (1979). Time without end: Physics and biology in an open universe. *Reviews of Modern Physics*, 51(3), 447–460.

Dyson, F. J. (2001, March 13). Is life analog or digital? *Edge*. https://www.edge.org/conversation/freeman_dyson-is-life-analog-or-digital

Fischer, J. M. (1994). Why immortality is not so bad. *International Journal of Philosophical Studies*, 2(2), 257–270.

Hart, D. B. (2019). *That all shall be saved: Heaven, hell, and universal salvation*. Yale University Press.

Heidegger, M. (1927/1962). *Being and time* (J. Macquarrie & E. Robinson, Trans.). Harper & Row.

Husak, D. (1992). Why punish the deserving? *Noûs*, 26(4), 447–464.

Kant, I. (1797/1996). *The metaphysics of morals* (M. Gregor, Trans.). Cambridge University Press.

Krauss, L. M., & Starkman, G. D. (2000). Life, the universe, and nothing: Life and death in an ever-expanding universe. *The Astrophysical Journal*, 531(1), 22–30.

Kvanvig, J. L. (1993). *The problem of hell*. Oxford University Press.
McKenna, M. (2012). *Conversation and responsibility*. Oxford University Press.
Metz, T. (2019). *God, soul and the meaning of life*. Cambridge University Press.
Moore, M. S. (1997). *Placing blame: A theory of the criminal law*. Clarendon Press.
Morris, H. (1968). Persons and punishment. *The Monist*, 52(4), 475–501.
Pereboom, D. (2005). Free will, evil, and divine providence. In A. Dole & A. Chignell (Eds.), *God and the ethics of belief: New essays in philosophy of religion* (pp. 77–98). Cambridge University Press.
Plantinga, A. (1974). *God, freedom, and evil*. Eerdmans.
Scanlon, T. M. (2008). *Moral dimensions: Permissibility, meaning, blame*. Belknap Press of Harvard University Press.
Scheffler, S. (2013). *Death and the afterlife*. Oxford University Press.
Sherman, L. W., & Strang, H. (2007). *Restorative justice: The evidence*. The Smith Institute.
Strawson, P. F. (1962). Freedom and resentment. *Proceedings of the British Academy*, *48*, 187–211.
Stump, E. (2010). *Wandering in darkness: Narrative and the problem of suffering*. Oxford University Press.
Swinburne, R. (1998). *Providence and the problem of evil*. Clarendon Press.
Swinburne, R. (2004). *The existence of God* (2nd ed.). Clarendon Press.
Tadros, V. (2011). *The ends of harm: The moral foundations of criminal law*. Oxford University Press.
Talbott, T. (2014). *The inescapable love of God* (2nd ed.). Cascade Books.
van Inwagen, P. (2006). *The problem of evil*. Oxford University Press.
Walls, J. L. (1992). *Hell: The logic of damnation*. University of Notre Dame Press.
Wielenberg, E. J. (2005). *Value and virtue in a godless universe*. Cambridge University Press.
Williams, B. (1973). The Makropulos case: Reflections on the tedium of immortality. In *Problems of the self* (pp. 82–100). Cambridge University Press.