CHAPTER 21

Conflict and Deep Persuasion

Presupposes: the primary/added value taxonomy (Ch. 15); the great moral conflation (Ch. 16).

Establishes: the three modes of conflict resolution; their hierarchical relationship; deep persuasion as the only source of structural peace.

In brief. Most of what passes for peace is not peace at all — it is managed conflict, held in place by force, incentive, or fatigue, and stable only while its overrider holds. A conflicting preference can be addressed at exactly three points: the capacity to act on it (enforcement), the calculation around it (shallow persuasion), or the preference itself (deep persuasion) — and there is no fourth. Only the third leaves nothing suppressed, nothing waiting on a change of circumstance to re-express itself; the strict definition of peace is forced, not stipulated. Peace and freedom turn out to be one limit approached from two directions. Deep persuasion is field-enrichment between agents, and its empirical signature is durability: a shift that survives the removal of what sustained it. The clock of peace advances one transformed preference at a time.

---

1 Introduction

Most of what passes for peace is not peace at all. It is managed conflict: the suppression of disagreement through force, incentive, or fatigue. A society in which dissenters are silenced is subdued. A society in which dissenters comply because the cost of defiance exceeds the cost of submission is stable. Neither is peaceful. Structural peace requires something neither force nor incentive can produce — the transformation of the preferences that generated the conflict in the first place.

The chapter commands nothing. Before the structural question can be asked, a prior one must be met: does the chapter command that peace be sought? It does not. The analysis is descriptive rather than prescriptive, and its central claims are conditional.

If aligned social peace is the goal, then deep persuasion is the only route to it, and the conditions developed in the next chapter are what that route requires.

Whether to adopt the goal is left open, and deliberately so, because the goal is itself a target of the kind this chapter analyses: to come to want social peace is to undergo deep persuasion rather than to accept an axiom. The empathic account of Chapter 16 does not compel the aim; it explains why the persuasion so reliably succeeds, since an agent whose empathic modality engages the suffering that conflict produces is, on exposure and reflection, moved toward its relief. The chapter therefore earns its organising value the way the next will argue all values must be earned: by persuasion rather than by authority. What follows maps the mechanism for those the persuasion has reached, and commands nothing of the rest.

What, then, is the peace so reached? Strictly, it is the limit at which emotional alignment is maximal: the state in which each person can act on her preferences because no one's preferences stand against them, so doing what one wants and doing what others would have one do coincide. The freedom this names is unconditioned by others, since nothing is being suppressed; it is not licence, because the wants that would collide with others have been dissolved in the very alignment that produces it — the empathic constitution of morality (Chapter 16) having already shaped each agent's wants to include the others'. Difference is not erased but freed to act without collision, so the suffering conflict produces falls away with the opposition that generates it.

Real communities only approach this limit, so the degree of alignment a society reaches is the degree of freedom its members enjoy. The limit is a limit, and the argument needs less than it. What §2.1 states, and what the results of this chapter and the next require, is alignment at some level: agreement on the framework under which the lower-level dispute is conducted. The complete alignment the limit names is the direction in which the clock of peace points (§4) rather than a condition a society must meet for the analysis to apply.

Peace and freedom are one limit approached from two directions, within the domain the identity concerns, which is constraint by other agents. Nature, scarcity, and the limits of one's own body constrain action too; they are not conflict, and no alignment dissolves them. The identity is between peace and social freedom, not between peace and omnipotence. Within that domain the only constraint on action is another's opposed preference, so alignment removes the opposition and the completion of peace and of social freedom coincide. This reframes a tradeoff much political theory takes for granted, on which social harmony is purchased at the cost of individual liberty: here the harmony and the liberty arrive together, each the measure of the other.

The republican objection. One tradition contests the identity rather than the tradeoff; it should be named where the claim is made. On the republican account (Pettit, 1997; engaged at Chapter 22, §2), freedom is non-domination rather than the absence of opposed preference: the contented slave of a kindly master holds no opposed preference and is unfree, since the master retains the power to interfere at will.

The reply is not stipulation but Chapter 15 (§6.6). Her contentment is an adaptive preference sustained by foreclosed alternatives, and the informed-registration anchor (Chapter 14, §6.5) returns the opposed preference the domination suppressed, so the case is not alignment but its counterfeit. Where alignment is real, surviving her accurate and complete registration of what is done to her, the retained power is a capacity she does not want exercised. The republican will not accept that description of his complaint, since for him domination is unfreedom now; what the reply achieves is location — the accounts divide on whether a standing power to interfere, unexercised and unwanted, is itself a constraint on action or a prediction about one, and the identity of peace with freedom holds on the second reading only.

This freedom-through-alignment also places peace within the book's metaphysics, and it forestalls a misreading the framework's own vocabulary invites. Given the partition of Chapter 3 it is tempting to picture peace as an equalisation of the evaluative field, the dissolution of the differences in which conflict consists. The picture is inverted, and Chapter 6 says why: equality is the unintelligent terminus, the state of

zero distinction that can host no mind and so no value, so a peace that equalised the field would abolish the parties rather than reconcile them.

Conflict is not difference but opposed difference: two evaluative orientations bearing on one behaviour from contrary directions, the one preferring its occurrence and the other its omission. Resolution removes the opposition rather than the difference. When two agents come to prefer the same of the same behaviour their preferences align without fusing, since there remain two evaluative processes, one in each person, reaching a shared direction rather than collapsing into one.

What peace preserves is therefore not a difference of direction, because on the contested behaviour the directions coincide, but the irreducible difference of bearer: two minds that remain two even in perfect agreement. Peace never approaches the equality pole, because the plurality of persons is never crossed; it approaches only the alignment of distinct processes. This is the content of a difference without opposition — the social form of the book's commitment to difference, not its dissolution.

The question this chapter asks is not what political institutions best promote peace, which is a question for political science. The question is structural: what must be true of a process of conflict resolution for the resolution to be stable? The answer is that conflict resolution admits of exactly three modes, that only one achieves lasting stability, and that the three modes stand in a hierarchical relationship grounded in the evaluative priorities of the agents involved.

The argument defines conflict and peace structurally (§2), then develops a three-mode taxonomy of conflict resolution — enforcement, shallow persuasion, and deep persuasion — and argues for its exhaustiveness (§3). It establishes the hierarchical ordering among the three modes and the progressive trajectory of the clock of peace (§4), examines lawmaking and law enforcement as distinct operations, including the vulnerability of law to emotional fluctuation (§5), and engages Habermas's ideal speech situation and the deliberative democracy tradition of Fishkin and Dryzek (§6).

2 The Structure of Conflict and Peace

Conflict is, strictly, the co-presence of two preferences bearing on a single behaviour from contrary directions: one agent prefers that the behaviour occur, another that it not occur. The preferences are moral preferences — evaluative orientations constituted by the empathic modality, generating dispositions that move agents toward or away from specific conduct.

Two features of this structure govern what follows. The negative preference is a preference in full rather than the absence of one, since to prefer that an action be omitted is to occupy a determinate position, whether the omission is secured by acting otherwise or by doing nothing at all. Declining to act, where another prefers the action, is therefore the holding of one pole rather than a neutral standing-aside.

And because a behaviour either occurs or does not, two agents engaged with the same behaviour hold preferences that are exhaustively concordant or contrary — aligned on a single pole or set against each other — with no intermediate relation available. Preferences that merely fail to collide are not an intermediate case, since they bear on different behaviours, and where the behaviours differ no shared object exists over which a conflict could arise. The binary is not a simplification imposed for tractability but the contrastive individuation of Chapter 2 at the level of preference: a preference for an action is the preference it is only against the preference for its omission, as a content is the content it is only against what it excludes.

Scarcity is not conflict. One clarification forestalls a natural objection to the definition's scope. Not every situation in which two agents cannot both have what they want is a conflict in the present sense. Two agents who both prefer that a single indivisible good come to them — one kidney, one appointment, one seat — are, at the level of behavioural preference, aligned: each prefers that the needy be helped, and neither prefers that the other be denied. Their predicament is scarcity rather than opposed preference, since the frustration is imposed by the world's not affording both satisfactions rather than by either holding the pole the other rejects.

The distinction matters because the two call for different treatment, and the chapter's central claim is confined to the second. Scarcity is addressed by allocation, by procedures for dividing or sequencing a good the parties do not evaluatively oppose, and a just allocation can leave both parties in full evaluative alignment, wanting for each other the very thing the shortage denies. Opposed preference is addressed by the three modes below, because there the parties want incompatible things of the same behaviour, and no division of a good dissolves the opposition.

The two can of course combine, since scarcity often generates opposed preference, as when the parties come to prefer their own claim over the other's and the shortage hardens into rivalry. When it does, the opposed-preference component is what the analysis treats, and deep persuasion — here, the parties' coming to weigh each other's need as their own — is what resolves it. The residual allocation problem remains, while ceasing to be a conflict in the strict sense: it is a distribution to be arranged among agents who no longer stand against one another. The peace–freedom identity of §1 is accordingly a claim about opposed preference rather than about scarcity, since alignment removes the opposition that constrains action without conjuring a second kidney. What it secures is that where goods must still be divided, they are divided among agents whose preferences no longer collide.

Peace is the presence of shared behavioural preferences, or the absence of conflict. The definition is strict, with the strictness derived rather than stipulated. A community is peaceful when no one dislikes anyone else's behaviour, when everyone already behaves as everyone else wishes they would. This is the positive condition set out in §1: with alignment maximal, each can act on her preferences without standing against another's, so peace is at once the absence of conflict and the fullest exercise of freedom.

Galtung's (1969) distinction between negative peace, the absence of direct violence, and positive peace, the absence of structural violence and the presence of social justice, maps onto the present account directly. Negative peace corresponds to enforcement, since the capacity for violent action has been removed; positive peace corresponds to deep

persuasion, since the preferences that generated the conflict have been transformed. The strict definition used here corresponds to Galtung's positive peace and identifies the structural conditions that make it possible.

Why strictness is forced. That the strict definition is the right one, and not merely one stipulation among several, can be shown rather than asserted. Consider the looser definition a critic is likely to prefer: peace as stable coexistence, in which compliance is maintained whether or not the underlying preferences have been reconciled. Galtung's negative peace and the Rawlsian modus vivendi are of this kind.

Any such state is held in place by something — by the capacity to enforce, by an incentive structure, or by fear or fatigue — and each of these is contingent. Enforcement capacity erodes, incentives shift, fear fades. When the sustaining factor weakens, the conflicting preference, which was never transformed and so never went anywhere, re-expresses itself as conflict. A coexistence sustained this way is therefore not a distinct terminal condition but a metastable one: conflict in suspension, stable only while its contingent overrider holds. By the analysis of §3 it is managed conflict, and it belongs on the conflict side of the ledger rather than the peace side.

The only condition not held in place by a contingent overrider is the one in which there is no suppressed preference left to re-express: the condition in which the conflicting preference has itself been transformed into a shared one. This is the strict definition; it is now seen to be forced rather than chosen. It names the unique fixed point of conflict resolution — the only state from which a change of circumstance produces no reversion to conflict, because nothing is being held down.

The apparent counterexample confirms this rather than threatening it. A constitutional or federal arrangement that both parties accept as fair, though neither has changed the underlying preference, is peaceful just to the extent that the parties share a preference for the arrangement itself: a shared preference at the level of the framework, reached by deep persuasion about its fairness, even where first-order disagreement

persists beneath it. To the extent that acceptance is instead merely strategic, better than the alternative given present power, it is shallow persuasion, and it reverts when the balance of power shifts.

2.1 Levels of alignment

The strict definition invites an objection with institutional weight behind it, and meeting it here fixes the definition's scope. Rawls's (1993) fact of reasonable pluralism holds that some evaluative disagreements survive ideal deliberation among reasonable agents, not through ignorance, bad faith, or failed communication, but through the burdens of judgement: differences of experience, weighting, and interpretive standpoint that no procedure eliminates. If some first-order disagreement is permanent in principle, then a peace defined as the transformation of every conflicting preference is not a distant ideal but an incoherent one, and the strict definition names a state no possible society occupies.

Alignment is levelled. The reply is that alignment is levelled, and the strict definition, correctly read, requires alignment at some level rather than at all of them. Preferences bear on behaviours, while arrangements for handling disagreement about behaviours are themselves objects of preference, and so upward. Call first-order behavioural preferences level zero, preferences over the frameworks that govern level-zero disagreement level one, and preferences over the choice among frameworks level two.

The federal case displays the mechanism. Two communities that disagree about a behaviour, and cannot be brought to agree, may nonetheless come — by deep persuasion about fairness rather than by strategic calculation — to a shared preference for the arrangement under which each decides for itself. The level-zero conflict persists; the level-one preference is shared; and the analysis of deep persuasion applies at the higher level as at the lower, since nothing is suppressed there, no overriding response holds anything down, and no change of circumstance finds a conflicting level-one preference waiting to re-express itself. A conflict is structurally resolved when there is some

level at which the parties' preferences concerning it are shared and every level below is governed by that shared preference.

The ceiling principle. The governance clause deserves its own statement, because it is what blocks the manufacture of "peace" by ascent. Agreement at some lofty level counts for nothing unless it regulates the levels where the conflicting conduct lives: the shared preference must be the one under which the lower-level dispute is conducted, constrained, and, when it flares, adjudicated. A shared abhorrence of mutual annihilation that leaves the parties' fighting just as before governs nothing and resolves nothing. A conflict is managed when every level exhibits either opposition or merely strategic acceptance.

The strict definition is thereby preserved rather than weakened, since transformed shared preference remains the unique fixed point and the complete alignment §1 named as the limit is the limit of this fixed point rather than a threshold. What is fixed is its object: the transformation peace requires is located wherever in the hierarchy it can be had, and the burdens of judgement, which concern the recalcitrance of level-zero differences, give no reason to expect failure at every level.

Rawls's own overlapping consensus, on this reading, is a special case of the account rather than a rival to it: deep persuasion succeeding at level one, about the terms of association, among parties whose level-zero hierarchies remain distinct. Rawls's insistence that the consensus be moral rather than a modus vivendi is the deep/shallow distinction of §3, drawn at the framework level.

Two consequences follow. First, the limit at which alignment is maximal is now doubly indexed: a society approaches peace both as more conflicts are resolved and as the level at which each is resolved descends — framework alignment where preference alignment is unavailable, preference alignment where it is. The clock of peace (§4) advances along both dimensions, so a society whose deepest disagreements are held within shared frameworks is at peace with respect to them in the strict sense, however lively the level-zero contest remains.

Second, the diagnosis of managed conflict gains sharpness. What condemns enforcement and shallow persuasion is not that first-order disagreement persists beneath them, since it may persist beneath peace itself, but that no level exists at which the parties' preferences are shared rather than suppressed. Agonistic democracy (§7.2) is accommodated on the same terms: the ongoing contest it prizes is level-zero difference exercised inside level-one alignment — not the elimination of politics but its housing.

3 Three Modes of Conflict Resolution

When two agents hold opposing behavioural preferences, the conflict can be resolved in exactly three ways.

The trichotomy is a partition rather than a taxonomy read off the cases, and its exhaustiveness can be derived. For a conflict to end, the colliding conduct must cease, and either the preference behind it persists or it is transformed; there is no third condition for a preference to be in.

If the preference persists, compliance is externally maintained; there are exactly two handles by which conduct can be secured against a standing preference: the agent's capacity, what she is able to do, or her calculation, what doing it costs or yields — incentive, threat, and deception included, since deception operates by corrupting the calculus. Enforcement is the capacity route; shallow persuasion is the calculation route. If the preference is transformed, the resolution is internal and self-maintaining: deep persuasion.

The exhaustiveness of the modes therefore rests on one premise a challenger must reject — that conduct contrary to a standing preference is sustained only through capacity or calculation — and the cost of rejecting it is a form of compliance that is neither compelled, nor chosen on balance, nor wanted, which is not a form of compliance at all. That only the third mode yields structural peace then follows from the strict definition of §2: a resolution is structural when no suppressed preference waits on a change of circumstance, and on the first two routes the preference, by hypothesis, persists.

The partition has neighbours in the conflict-resolution literature, recorded here with the difference marked. Kelman (1958) distinguishes compliance, identification, internalisation as processes of attitude change; Deutsch (1973) separates destructive from constructive processes of resolution; Burton (1990) grounds settlement in needs rather than negotiating positions. Each classifies observed change, whereas the three modes fall out of the persistence of a preference under the two handles a party can hold, so the exhaustiveness claimed above is structural rather than empirical.

3.1 Enforcement

Enforcement alters an agent's behaviour without changing her underlying preference. The attacker is disarmed or imprisoned. He may still want to harm you, but he is no longer able to act on that desire. Enforcement operates on capacity, not on motivation. Arendt (1970) separates power, resting on acting in concert, from violence, which can destroy power while never creating it; enforcement in the present sense operates on the violence side of that line, reaching conduct while leaving the preference untouched.

Enforcement is sometimes necessary, since an immediate threat to life cannot wait for persuasion. But enforcement alone cannot produce peace. If all disputes were resolved by enforcement, the only "peaceful" society would be one in which all dissenters have been killed: peace by extermination. Keeping dissenters alive through imprisonment or forced silence is not peace in the deeper sense, because the preference for conflict remains.

The post-1945 occupation of Germany and Japan involved massive enforcement: the military apparatus was dismantled, leaders were tried, and the capacity for aggressive action was physically removed. The enforcement was necessary, since it stopped the immediate harm; it did not by itself produce peace. What followed — reconstruction, institutional reform, cultural exchange, and decades of diplomatic engagement — was the slow work of persuasion.

The German case audited. The case can carry more weight than illustration, because it is concrete enough to run the informed-registration audit on, and running it shows the test operating. Ask of the post-war German transformation the audit's questions.

Were the preference shifts produced under concealment — of the conduct's nature, of the victims, of the record? The opposite: occupation-era policy forced exposure through documentation, trials, curricula, and the publicised evidence of the camps, so the shift was produced by registration where the prior regime had operated by its suppression.

Did the shifts revert when the sustaining pressure was removed? The occupation's enforcement apparatus wound down within a decade, though other pressures persisted — Western integration and its incentives in the Federal Republic, Soviet control in the East — so the axis returns a provisional verdict rather than a clean one. What it can record is that the preference shift deepened across the following two generations, surviving reunification and political stresses that would re-expose a merely strategic accommodation. Did the transformed preferences generalise beyond the enforced domain? They did: into constitutional culture, memorial practice, and a politics of historical responsibility no occupier demanded.

On all three axes the audit returns the deep verdict, and the contrast case sits within the same history: where exposure was curtailed and accommodation strategic, in parts of the post-Soviet space, the pattern re-expressed when the pressure lifted (Chapter 22, §7.4). The audit is thereby shown to be a usable instrument, with the taxonomy's discriminator — durability under removed sustaining conditions — an observable one.

The instrument should be stated as method and carried forward, because live cases cannot wait for the retrospect in which durability settles. Classification of a current transformation proceeds on the three axes prospectively: whether the shifts are produced under concealment or under exposure; whether they revert wherever the sustaining pressure locally relaxes; and whether they generalise beyond the enforced

domain. The axes are measurable while the case runs, the verdict they return is provisional and revisable as the record lengthens, and durability remains the criterion of which they are the running proxies.

3.2 Shallow persuasion

Shallow persuasion occurs when an agent refrains from conflicting action not because her preference has disappeared but because another emotional response temporarily dominates it. A bribe may make compliance more attractive than defiance. A threat of punishment may make the cost of acting on the preference outweigh the benefit. In each case the conflicting preference survives, overridden by a stronger emotional response.

Institutions on both sides of the line. Institutional arrangements occupy both sides of the deep/shallow line, which runs where §2.1 drew it. Constitutional frameworks, federal arrangements, negotiated compromises, and modus vivendi agreements divide into two cases.

Where acceptance is strategic — each party tolerating the arrangement because the alternative is worse for it, given present power — the arrangement is shallow persuasion at a high level of institutional sophistication: real, valuable, and reversion-prone, since the acceptance lasts while the balance that sustains it lasts. Where acceptance is real, each party sharing a preference for the arrangement as fair, the conflict has been resolved by deep persuasion at the framework level (§2.1): the level-zero disagreement persists, nothing at level one is suppressed, and no shift of circumstance finds a conflicting framework-preference waiting to re-express itself.

What the chapter notes of the strategic case is that the underlying disagreement remains and can resurface when enforcement weakens, when the incentive structure shifts, or when circumstances embolden the dissenter. Cold War deterrence (Schelling, 1960) provides a further illustration: the nuclear powers were shallowly persuaded to avoid direct confrontation by the overriding emotional response of fear. The compliance was real and lasted decades; it was management rather than peace, since the confrontation was suspended while the preferences

behind it were untouched. Cooperation sustained by reciprocity (Axelrod, 1984) sits on the same side of the line, since the strategy answers the incentive structure while the preferences behind the rivalry persist.

A second case from the same period shows the other mode. Within the Soviet sphere, national and ethnic conflicts were held down by enforcement, and when the enforcing capacity collapsed after 1991 the suppressed disagreements re-expressed themselves — which is what the taxonomy predicts of a preference that was never transformed.

3.3 Deep persuasion

Deep persuasion occurs when the conflicting preference itself is changed or removed. If you show the attacker convincing evidence that you were not involved in the wrongdoing he suspected, his motivation to harm you vanishes. The conflict is resolved at its root — not by removing his capacity to act, not by overriding his desire with a stronger one, but by transforming the desire itself. After deep persuasion, the agent no longer has the conflicting preference. There is nothing to suppress, nothing to override, and nothing to re-emerge when circumstances change.

Deep persuasion works by restructuring the agent's evaluative hierarchy, by introducing emotionally meaningful insights, perspectives, or experiences that cause the agent to re-evaluate the significance of the conduct in question. It is not merely the provision of new information, which might leave preferences unchanged, but the provision of information that engages the evaluative modality and produces a real shift in emotional response.

Stated in the vocabulary of Chapter 15, deep persuasion is field-enrichment (Chapter 15, §4.4) operating between agents rather than within one. The perspectives communicated are new secondary objects entering the other's evaluative field, each carrying added value grounded in real registrations, until the composite governing the conflicting preference tips. The mechanism is inherited rather than constructed, since what Chapter 15 derived for the deliberative deadlock

transfers to the interpersonal case, the field being enriched having the same architecture whoever supplies the enrichment. The discipline transfers with it: legitimate enrichment bottoms out in real registrations, a constraint whose work appears in §7.6.

Lederach (2005) develops the "moral imagination" in peacebuilding: the capacity to imagine oneself in a web of relationships that includes one's adversary and to envision a shared future neither party has yet experienced. On the present analysis that is the modality operating in its empathic deployment at the social scale. Bar-Tal (2013) provides the complementary social-psychological analysis: intractable conflicts are sustained by "conflict-supporting narratives" — in the present vocabulary, hierarchies in which the derived value of group identity overrides the primary empathic response to the adversary's suffering. Deep persuasion, in his terms, is then the transformation of these narratives through sustained engagement with the adversary's actual evaluative experience.

The South African Truth and Reconciliation Commission provides a concrete, if imperfect, illustration. It exposed perpetrators and victims to each other's experiences, creating conditions in which the modality could be engaged by the suffering the system had caused. There is no better illustration of the deep/shallow difference than the TRC's mixed record. Where the process succeeded it produced real shifts in evaluative response — perpetrators who came to feel the weight of what they had done, victims who came to see perpetrators as persons capable of remorse. Where it failed it produced the appearance of reconciliation: shallow persuasion dressed as deep, the forms of engagement without the substance.

3.4 The exhaustiveness of the taxonomy

The exhaustiveness derived at the head of §3 can be restated at the level of the architecture. A preference can be addressed at exactly three points: the capacity to act on it (enforcement), the relative strength of competing responses in the field (shallow persuasion), or the preference itself (deep persuasion). There is no fourth, because a preference either

persists unchanged or is transformed, and persistence admits only the capacity and calculation handles.

One putative fourth mode must be addressed directly, because it appears to resolve conflict without operating at any of the three levels. The mode is exit, in Hirschman's (1970) sense of the option a party keeps beside voice: separation, secession, avoidance, or the division of a contested domain so that each party governs its own. Voice is persuasion in the present vocabulary; exit reduces as follows. Exit seems to dissolve the conflict by removing the co-presence that the definition of conflict requires, rather than by acting on capacity, on the relative strength of competing responses, or on the preference itself.

It is not a fourth mode, since it reduces to the existing three. Where exit divides a contested object — two communities that disagree over a single rule agreeing that each locality will set its own — the original preference is not addressed at all. What is removed is the shared object over which it was opposed. This is real resolution only if the division is itself a shared preference, both parties preferring local determination to the alternatives. In that case the conflict has been resolved at the third level with respect to the arrangement: deep persuasion about the framework rather than a new mode. If the division is accepted only strategically, the first-order preference survives and re-expresses itself whenever the division is contested, which is shallow persuasion.

Where exit instead removes co-presence over an object that cannot be divided, the parties simply ceasing to interact, the opposed preferences are suspended rather than transformed, dormant until renewed contact revives them. This is management by avoidance, a limiting case of enforcement in which the capacity to act on the preference against the other is removed by removing the other from reach. On neither construal does separation resolve the conflict at a level the taxonomy does not already contain. The three modes remain exhaustive.

Where the argument stands: conflict and peace have been defined structurally, with alignment levelled and the ceiling principle fixed (§2); the three modes of resolution have been derived as an exhaustive

partition (§3). What follows orders them and shows why the ordering is one of dependency rather than of practice.

4 The Hierarchical Ordering

The three modes stand in an ordering of need, and the ordering is structural in one respect only: where survival or security is not secured, deep persuasion cannot begin. The ordering has parallels with Maslow's (1943) hierarchy of needs, since both identify a sequence in which higher-level achievements presuppose the satisfaction of lower-level conditions. The grounding differs. Maslow's hierarchy is a psychological generalisation about the order in which human beings typically pursue needs; the present hierarchy is a claim about dependency.

Deep persuasion requires the prior satisfaction of enforcement and shallow persuasion conditions, regardless of the empirical order in which individuals happen to pursue them. The claim is not that all human beings or all societies satisfy survival needs before security needs. The claim is that deep persuasion structurally requires that the conditions the other modes address, survival and security, obtain by whatever means: you cannot transform someone's evaluative hierarchy if she is dead or too afraid to engage. Where the conditions obtain untended, as between parties who were never at war, deep persuasion proceeds with neither mode having operated. Where they do not, enforcement and shallow persuasion are what secures them — which is why the modes are ordered as they are even though they need not all have run.

Enforcement addresses survival. When someone is pointing a gun at your head, the immediate evaluative priority is survival. You cannot deeply persuade someone who is actively killing you.

Shallow persuasion addresses security. Once survival is secured, the evaluative priority shifts to security — the confidence that the threat will not return tomorrow. A person cannot begin the slow work of understanding her adversary's evaluative perspective if she is still afraid he will attack her.

Deep persuasion addresses alignment. Once security is established, a deeper evaluative priority can emerge: understanding, shared values, the elimination of the underlying disagreement. Deep persuasion is the terminal mode, the one that produces actual resolution rather than management, and it depends on the prior modes having done their work.

Each instance of deep persuasion expands the territory of shared understanding, and as shared understanding expands, fewer conflicts require enforcement or shallow persuasion. The trajectory is progressive without being guaranteed. New conflicts can arise through new circumstances, new generations with differently constituted evaluative hierarchies, or new situations that trigger previously absent emotional responses. The institutional apparatus of conflict resolution must therefore be preserved as a standing capacity, not because it is currently needed but because it may be needed again. The apparatus is like a fire station in a town with no fires: the absence of fires does not justify dismantling the station.

The progression can be visualised as a clock of peace, each moment of deep persuasion turning the hand forward and expanding the territory of shared understanding. Over time, enforcement and shallow persuasion fade as they are needed less, until a full rotation is complete and peace is achieved.

The clock is a regulative ideal and makes no prediction. A society may never complete the rotation. But the direction of travel matters: a society whose clock is advancing, whose instances of deep persuasion are increasing and whose reliance on enforcement and shallow persuasion is decreasing, is moving toward structural peace even if it never arrives. The German trajectory examined in §3.1 supplies the one dated illustration the metaphor otherwise lacks (an illustration, since a regulative ideal admits no calibration) — enforcement dominant in 1945, institutional and educational persuasion carrying the load within a decade, and preference-level transformation measurable across two generations: one dated case of the hand advancing, with the post-Soviet contrast (§3.1) as the clock run in place.

5 Lawmaking and Law Enforcement

5.1 Two operations

Lawmaking and law enforcement are different operations, and their conflation produces characteristic pathologies.

Lawmaking is the process by which private moral preferences are transformed into shared public standards. Its ideal mechanism is deep persuasion: agents communicate their evaluative hierarchies, engage with the hierarchies of others, and arrive at shared preferences that are then codified as rules. The moral authority of a law derives from the quality of the persuasion that produced it. A law codifying a shared preference has deeper moral authority than a law imposed by a majority over a dissenting minority, which in turn has deeper authority than a law imposed by force.

Law enforcement operates across all three modes: it enforces through arrest, fines, and imprisonment; it shallowly persuades through deterrence; and it can contribute indirectly to deep persuasion through education. The distinction matters because lawmaking and law enforcement have different normative requirements. Lawmaking requires persuasion, because its product, shared standards, derives its moral authority from evaluative convergence. Law enforcement may legitimately use enforcement and shallow persuasion, because its function is to maintain compliance with standards already established.

Law enforcement is also a more technical act than lawmaking, since it does not require emotional involvement. An officer is not tasked with agreeing or disagreeing with legislation; she compares observed behaviour to what is written in the legal code. Even if she personally opposes a specific rule, she can still recognise whether an action violates it. The written law translates subjective moral preferences into objective behavioural standards that can be applied without evaluative engagement.

A society that uses enforcement in its lawmaking, imposing laws by force rather than establishing them through persuasion, is authoritarian. A society that uses only deep persuasion in its law enforcement, relying

entirely on transforming preferences without any coercive apparatus, cannot protect its members from imminent harm. The two operations require different mixes of the three modes, and conflating them produces these characteristic pathologies.

In practice the operation runs on mutual adjustment: each person works on the structure of the others' evaluations until a shared preference emerges, an adjustment that, done well, deepens each party's empathic grasp of the others' primary and secondary objects of evaluation (Chapter 15). A participant who comes to weigh environmental protection more heavily once she has felt the health consequences it carries for identifiable others has been persuaded in the deep sense, since her modality has responded to differences it had not previously been shown, and nothing has overridden her. Chapter 22 (§5.1) works the mechanism through in detail.

5.2 Emotion and the vulnerability of law

Because moral preferences are rooted in emotional responses, law — being built upon such preferences — is vulnerable to the fluctuations of emotion.

Consider a person who regards the death penalty as immoral and campaigns for its abolition. After years of advocacy, the government changes the law. But the day after the reform, her child is abducted and harmed by a serial offender. Overcome by grief and rage, the same person who fought to ban capital punishment now demands its reinstatement. Her evaluative hierarchy has shifted, not through new evidence or deeper empathic understanding but through the overwhelming intensity of a personal emotional response. This is not a failure of character; it is the evaluative modality operating as it always operates — constituting oriented responses whose intensity is determined by the proximity and vividness of the relevant evaluative differences.

The scenario reveals the danger of placing lawmaking in the hands of a single individual, whose emotional state may shift quickly and radically: fear or hatred toward a group one day, a law to punish them

the next. Most legal systems therefore require that several people, each with a different emotional context, agree before laws change, so laws are revised only after multiple evaluative hierarchies have been engaged, reducing the probability that a single intense response will distort the process. The more fundamental the legislation, the broader the agreement needed to amend it — as in constitutional law, where supermajority requirements and multi-stage ratification are designed to buffer the evaluative volatility the modality's emotional constitution makes possible.

H. L. A. Hart (1961) provides the jurisprudential foundation. The "rule of recognition", the master rule that identifies which other rules count as law, is itself a social practice, sustained by the convergence of officials' evaluative attitudes. On the structural account, Hart's rule of recognition is the institutional expression of a shared evaluative hierarchy among officials: a hierarchy produced by persuasion, deep or shallow, and maintained by the evaluative modality's continued oriented response to the practice.

The vulnerability is that the rule of recognition, being evaluatively constituted, can shift when the evaluative hierarchies of officials shift. The institutional safeguards against evaluative volatility — constitutional entrenchment, judicial independence, separation of powers — are therefore structurally necessary rather than merely convenient.

6 Habermas and the Ideal Speech Situation

6.1 Discourse ethics and deep persuasion

Deep persuasion overlaps with Habermas's (1990a, 1990b, 1996) ideal speech situation and discourse ethics. Both hold that consensus requires free and open communication, that a norm's legitimacy depends on the quality of the discursive process that produced it, and that compliance achieved through force differs from agreement achieved through the uncoerced exchange of reasons.

The convergence is substantial while the grounding differs. Habermas derives the conditions for legitimate discourse from the pragmatics of communication: anyone who enters an argument implicitly commits to norms of sincerity, truth-telling, and equal participation. The present framework derives them from the priority of deep persuasion, itself derived from the evaluative architecture — deep persuasion is the only mode that produces structural peace, it requires communication, and communication requires free and open speech. The conditions are therefore derived from the metaphysics of evaluative difference rather than the pragmatics of argumentation.

The difference in grounding has a consequence. Habermas's transcendental-pragmatic argument is vulnerable to the objection that one can simply refuse to argue, so that the presuppositions of argumentation bind only those who choose to engage. The present framework's derivation avoids this objection, because the conditions are derived from the requirements for peace rather than from the act of arguing. A person who declines to argue is free to do so, so her declining limits the available modes of conflict resolution to enforcement and shallow persuasion, neither of which produces peace.

The symmetry should be stated as plainly against the framework as for it: the derivation binds no one either. Its conditions are hypothetical imperatives, binding on whoever wants structural peace, and a person who declines the end escapes them as completely as Habermas's refuser escapes the pragmatics of argumentation. The advantage claimed is therefore not that the framework obligates where Habermas cannot, but that its conditional survives the refusal intact: the refuser is not caught in a performative contradiction, and needs to be caught in none, because the fact that only enforcement and shallow persuasion remain to her stands whether or not she ever argues.

6.2 Deliberative democracy: Fishkin and Dryzek

The concept of deep persuasion also converges with the deliberative democracy tradition, which holds that legitimate political decisions are those produced through inclusive, reason-giving deliberation among affected parties. Fishkin (2009) has developed deliberative polling: a

representative sample of citizens deliberates on a policy question under conditions of balanced information and moderated discussion, with the resulting opinion shifts measured.

Fishkin's empirical findings bear directly on the account. Deliberative polls consistently show significant opinion change after sustained engagement with opposing perspectives, particularly where participants are exposed to the experiences and arguments of those who disagree with them. On the structural analysis, deliberative polling is an institutional mechanism for facilitating deep persuasion at scale. The structured exposure to others' evaluative hierarchies provides the emotionally meaningful input that can produce real shifts in evaluative response.

The durability evidence. A more pointed piece of evidence bears directly on the deep/shallow distinction, because it isolates the one variable the taxonomy makes load-bearing: durability. Broockman and Kalla's (2016) field experiment found that a single canvassing conversation inviting the resident to recall and dwell on an analogous experience of her own — active perspective-taking — produced reductions in prejudice toward transgender people. The reductions persisted for at least three months and survived exposure to subsequent attack advertising.

What the intervention supplied was neither information nor incentive but the emotionally meaningful engagement of the resident's own modality with a perspective she had not previously registered. The persistence is the diagnostic feature. A shift produced by incentive or framing reverts when the sustaining factor is removed (§2); a shift produced by evaluative engagement leaves nothing to revert. That the effect outlasted a deliberate counter-message is the empirical signature the taxonomy predicts for the deep mode and denies to the shallow one. Kalla and Broockman's (2020) replications, identifying non-judgemental exchange of narratives as the operative ingredient, reinforce the reading: what transforms a preference durably is the very input §3.3 identifies as constitutive of deep persuasion.

Three qualifications keep the evidence weighed at its size. The finding's own history enforces the first: the 2016 study was undertaken as the corrective replication of an earlier canvassing result that collapsed under fraud. The subsequent experiments, while confirming durable effects, find them modest in size and uneven across issues. The magnitudes belong in the text: the 2016 study, on 501 voters, reported a reduction comparable to the fourteen-year national decline in homophobia; the three 2020 experiments, on 6,869 voters, found d = 0.08 for exclusionary attitudes toward immigrants and 0.08 and 0.04 in the transphobia replications, durable to four months. That is support for the mechanism's reality at the scale of a conversation, short of a conversion; what the chapter rests on it is durability, which the designs measure, with magnitude left where the studies leave it.

The second sets the single paradigm inside its base. The meta-analytic literature on intergroup contact (Pettigrew & Tropp, 2006) finds prejudice reliably reduced by contact across hundreds of studies, with the largest effects where the contact is personal and sustained. That is the population-level trace of the same mechanism, and the reason the chapter treats canvassing as one instrument of deep persuasion rather than its proof.

The third concerns the base itself. Paluck, Green, and Green's (2019) re-analysis of the contact literature confines the estimate to experimental designs and finds the effects smaller there than the correlational record suggests. Few studies measure outcomes at a delay, and fewer still among adults on the prejudices that matter most politically.

The framework's stake accordingly rests where the durable measurements are: on the canvassing experiments above, whose delayed outcomes were the point of their design, and on whatever the contact base retains once the re-analysis has done its work. A base that thinned further under such scrutiny would narrow the evidence for the mechanism without touching the derivation. The Register's stake on Chapter 15 (§4.4) is to be read at the size this paragraph fixes.

Dryzek (2000) develops a broader conception of deliberative democracy, extending beyond formal polling to include informal public deliberation, social movements, and the discursive contestation of norms in civil society. Dryzek argues that deliberative democracy is not a single institutional form but a family of practices unified by the commitment to decision-making through the exchange of reasons rather than through aggregation of fixed preferences.

The contribution to Dryzek's project is the evaluative architecture. The "reasons" Dryzek identifies as the currency of deliberation are, on the present account, emotionally meaningful inputs that engage the evaluative modality and can produce real shifts in evaluative response. A reason that engages no evaluative response is inert — the field with no body in it of Chapter 14 (§A.5). It may be logically compelling and motivationally empty.

7 Objections and Replies

7.1 The utopian objection

Objection. The strict definition of peace is unattainable. Deep persuasion is slow, difficult, and frequently impossible.

Reply. The chapter does not claim that deep persuasion is always achievable or that enforcement and shallow persuasion are never appropriate. It claims that deep persuasion is the only mode that produces peace. Real societies operate with a mix of all three modes on different timescales. The strict definition serves a diagnostic function: it tells us what peace actually requires, which in turn tells us what the alternatives can and cannot achieve.

7.2 The agonistic democracy objection

Objection. Mouffe (2000, 2013) argues that political life requires ongoing disagreement — that the attempt to eliminate all conflict is depoliticisation.

Reply. Mouffe (2005) identifies a real feature of democratic life. She argues that the liberal attempt to achieve rational consensus,

paradigmatically Habermas's discourse ethics, is not merely utopian but dangerous, depoliticising fundamental disagreements that reflect irreconcilable differences in values and identity. The task of democratic politics is instead to transform antagonism, enmity, into agonism, adversarial respect, in which opponents are adversaries whose right to disagree is acknowledged rather than enemies to be destroyed.

The response runs along the line already drawn: Mouffe's agonistic transformation divides at §2.1 and §3.2. Strategic acceptance of the democratic framework is shallow persuasion at a high level of sophistication, reverting as the balance shifts. A shared preference for the adversary's right to disagree is deep persuasion at the framework level — level-zero contest housed inside level-one alignment — with durability the empirical discriminator, as everywhere in the taxonomy. Either reading records a real achievement, divided by the same line as the constitutional pluralism of §3.2.

The analysis disagrees with Mouffe on two points. First, on the claim that all fundamental political disagreements are irreconcilable. Some are, and agonistic management is the best available response; others are susceptible to deep persuasion through the kind of sustained evaluative engagement the chapter describes. The analysis does not claim that deep persuasion is always possible. It claims that deep persuasion is the only mode producing structural peace, and that Mouffe's agonism is the correct fallback when deep persuasion fails or is not yet achievable.

Second, on the claim that the preservation of disagreement is intrinsically valuable. Mouffe holds that democratic vitality requires ongoing conflict; the analysis holds that conflict is a condition to be resolved where possible and managed where not, and that the institutional apparatus should be preserved because disagreement is always liable to recur rather than because it is valuable.

7.3 The pluralism objection

Objection. Stable pluralism — a society in which groups with fundamentally different values coexist under a shared constitutional

framework — is a real political achievement that the chapter's binary treatment misses (Rawls, 1993; I. Berlin, 1969).

Reply. The chapter agrees that stable pluralism is a real and important achievement, and the account's diagnostic value is that it distinguishes the two achievements "stable pluralism" names together, along the line of §§2.1 and 3.2. A shared framework-preference is peace in the strict sense at level one; strategic acceptance is shallow persuasion at its most sophisticated, its instability latent. The binary treatment the objection alleges is what the levelled analysis replaces.

7.4 The moral limits objection

Objection. Should we deeply persuade a Nazi rather than enforce against him? Some preferences are morally monstrous.

Reply. The hierarchical ordering makes clear that enforcement addresses the immediate threat (survival) and shallow persuasion contains the ongoing danger (security). Deep persuasion, if possible, eliminates the preference that generated the danger (alignment). A Nazi who is imprisoned but not persuaded will resume his project if released. A Nazi who is deeply persuaded — who comes to feel the suffering his ideology causes — is no longer a Nazi. Enforcement manages the threat; deep persuasion alone eliminates it. Whether deep persuasion is possible in a given case is empirical; that it is the only mode achieving structural resolution is structural.

The empirical record counsels the sobriety the reply's conditional already encodes. The literature on exit from violent extremism finds disengagement — behavioural departure, a shallow-persuasion outcome sustained by changed incentives and circumstances — far more commonly than deradicalisation proper, the transformation of the underlying commitments, and finds the latter slow, relational, and comparatively rare (J. Horgan, 2009; Bjørgo & Horgan, 2008).

The taxonomy predicts that asymmetry, since transformation requires the emotionally meaningful engagement of the subject's own evaluative modality, which cannot be administered as a programme, while behaviour yields to management. That the deep mode is rare

where it is most needed is not an objection to the hierarchy but the hierarchy's own forecast — and the reason enforcement and containment remain, for the Nazi as for anyone dangerous, the modes the analysis itself prescribes in the meantime.

7.5 The incommensurability objection

Objection. Some disagreements involve incommensurable values (I. Berlin, 1969), so deep persuasion may be impossible where evaluative hierarchies are fundamentally different in kind.

Reply. The objection identifies a real structural limit. Where incommensurability obtains at level zero, the first resort is ascent rather than management: framework alignment, deep persuasion at level one about the terms under which the incommensurable values coexist (§2.1), remains available where value-alignment is not. Where even that is unavailable, shallow persuasion at its most sophisticated — strategically accepted constitutional arrangements, federal compromises, prudential agonism — is the best remaining outcome. The analysis still serves a diagnostic function, since it identifies which conflicts are resolvable at level zero, which at a higher level, and which require permanent management.

7.6 The manipulation objection

Objection. If deep persuasion transforms evaluative hierarchies, how is it distinguished from manipulation?

Reply. Deep persuasion works through the engagement of the modality with emotionally meaningful content; manipulation works through its bypassing or exploitation, with responses triggered by deception, selective framing, or cognitive bias, without engagement with the content. In the terms of Chapter 14 this is the line between a fitting response and an unfitting one: deep persuasion brings the agent to register the relevant features more accurately, above all the states of those her conduct affects, whereas manipulation engineers a response that misreads them (Chapter 14, §6.4). The line is close to Noggle's (1996), on which manipulation is the attempt to lead another astray, inducing beliefs, desires, or emotions that fall short of the ideals

governing them; the account takes that shape and fixes the ideal by the anchor, the accurate and complete registration of the affected parties, where Noggle leaves it to the manipulator's own standards.

The boundary can be unclear in specific cases while the distinction is sharp — and inherited rather than invented here. Chapter 15 (§4.4) states it as the discipline on enrichment: added value must bottom out in real registrations. The time-indexed transformation test (Chapter 15, §3.2) supplies the audit, and the standard it applies is the informed-registration anchor (Chapter 14, §6.5). A preference transformed by deep persuasion carries its history — the encounters and registrations that restructured it, traceable in principle — whereas a manipulated shift pleads a history keyed to registrations that were concealed, fabricated, or misread. This is why the fixed-point asymmetry below holds: the manipulation's sustaining misreading is itself the real registration its history lacks. Sunstein's (2015) nudging occupies the boundary, since a nudge that engages the agent's evaluative capacity by presenting consequences vividly is closer to deep persuasion, while one that bypasses it by exploiting default biases is closer to manipulation.

A deeper structural difference lies behind this difference in mechanism: manipulation has no stable fixed point. It works by inducing a response keyed to an incomplete or inaccurate picture of the relevant evaluative differences, so its product survives only while the gap between that picture and reality holds — the same contingent overrider §2 identified in every state on the conflict side of the ledger. Once reality manifests the features the manipulation concealed, the induced response loses its support, so the original preference, still untransformed, re-expresses itself. Manipulation is in this respect a species of shallow persuasion, reverting whenever its sustaining misreading is exposed.

Deep persuasion leaves nothing of the kind to expose, because the agent has already registered the relevant differences accurately, so no concealed reality remains whose arrival could undo the shift. None of this guarantees that a given manipulation will unravel, since reality may stay hidden, or the misreading may be renewed faster than it decays, and

it locates the asymmetry. Deep persuasion is stable because it is fitting; manipulation is unstable because it fits nothing that reality will confirm.

8 Conclusion

Deep persuasion is the only mode that produces peace. This is a strict finding. Enforcement stops the harm. Shallow persuasion prevents its recurrence. Deep persuasion eliminates the preference that generated it. The three modes are not alternatives among which a society chooses; they are layers through which a society progresses, each creating the conditions for the next.

The chapter yields three results, all constructed.

Three modes of conflict resolution. Conflict resolution admits of exactly three modes: enforcement (capacity-removal), shallow persuasion (preference-override), and deep persuasion (preference-transformation). The taxonomy is exhaustive because a preference can be addressed at exactly three levels of the evaluative architecture, and the third level's mechanism is inherited — deep persuasion is field-enrichment (Chapter 15, §4.4) between agents, carrying that result's discipline and its time-indexed audit. There is no fourth level.

The hierarchy of modes. The three modes stand in a hierarchical relationship grounded in structural dependency: enforcement addresses survival, shallow persuasion addresses security, and deep persuasion addresses alignment. The trajectory is progressive, since each instance of deep persuasion expands shared understanding, without being guaranteed, and the institutional apparatus must be preserved as a standing capacity.

Deep persuasion and structural peace. Deep persuasion is the only mode that produces structural peace, the state of maximal social freedom in which alignment dissolves the opposed preferences that are the sole social constraint on action. Lawmaking and law enforcement are distinct operations with different normative requirements. The derivation converges with Habermas's ideal speech situation while being grounded in the metaphysics of evaluative difference rather than

the pragmatics of communication, its conditional surviving the refusal to argue intact (§6.1).

If deep persuasion is the terminal mode, what conditions must a society satisfy for deep persuasion to be structurally possible? And once those conditions are met, what is the concrete mechanism by which private hierarchies become shared law? The next chapter derives the conditions, identifies the mechanism, and shows that the institutional form that preserves both is democracy.

Grades and Dependencies

1. Stance and method. All three results are *constructed*, on Chapters 15 and 16. The chapter's stance is on record (§1): descriptive, not prescriptive — its claims are hypothetical imperatives that command nothing of those the persuasion has not reached, a symmetry stated as plainly against the framework as for it (§6.1). Its method is on record too (§3.1): the informed-registration audit runs prospectively on three axes — concealment vs exposure, local reversion, generalisation — with verdicts provisional and durability the criterion they proxy. The canvassing evidence is weighed at its size: one instrument of deep persuasion, not its proof (§6.2). The Premise Register carries the chain.

References

Arendt, H. (1970). *On violence*. Harcourt, Brace & World.

Axelrod, R. (1984). *The evolution of cooperation*. Basic Books.

Bar-Tal, D. (2013). *Intractable conflicts: Socio-psychological foundations and dynamics*. Cambridge University Press.

Berlin, I. (1969). *Four essays on liberty*. Oxford University Press.

Bjørgo, T., & Horgan, J. (Eds.). (2008). *Leaving terrorism behind: Individual and collective disengagement*. Routledge.

Broockman, D., & Kalla, J. (2016). Durably reducing transphobia: A field experiment on door-to-door canvassing. *Science*, 352(6282), 220–224.

Burton, J. W. (1990). *Conflict: Resolution and provention*. St. Martin's Press.

Deutsch, M. (1973). *The resolution of conflict: Constructive and destructive processes*. Yale University Press.

Dryzek, J. S. (2000). *Deliberative democracy and beyond*. Oxford University Press.

Fishkin, J. S. (2009). *When the people speak: Deliberative democracy and public consultation*. Oxford University Press.

Galtung, J. (1969). Violence, peace, and peace research. *Journal of Peace Research*, 6(3), 167–191.

Habermas, J. (1990a). Discourse ethics: Notes on a program of philosophical justification. In *Moral consciousness and communicative action* (C. Lenhardt & S. W. Nicholsen, Trans., pp. 43–115). MIT Press.

Habermas, J. (1990b). *Moral consciousness and communicative action*. MIT Press.

Habermas, J. (1996). *Between facts and norms*. MIT Press.

Hart, H. L. A. (1961). *The concept of law*. Clarendon Press.

Hirschman, A. O. (1970). *Exit, voice, and loyalty: Responses to decline in firms, organizations, and states*. Harvard University Press.

Horgan, J. (2009). *Walking away from terrorism: Accounts of disengagement from radical and extremist movements*. Routledge.

Kalla, J. L., & Broockman, D. E. (2020). Reducing exclusionary attitudes through interpersonal conversation: Evidence from three field experiments. *American Political Science Review*, 114(2), 410–425.

Kelman, H. C. (1958). Compliance, identification, and internalization: Three processes of attitude change. *Journal of Conflict Resolution, 2*(1), 51–60.

Lederach, J. P. (2005). *The moral imagination: The art and soul of building peace*. Oxford University Press.

Maslow, A. H. (1943). A theory of human motivation. *Psychological Review*, 50(4), 370–396.

Mouffe, C. (2000). *The democratic paradox*. Verso.

Mouffe, C. (2005). *On the political*. Routledge.

Mouffe, C. (2013). *Agonistics: Thinking the world politically*. Verso.

Noggle, R. (1996). Manipulative actions: A conceptual and moral analysis. *American Philosophical Quarterly*, 33(1), 43–55.

Paluck, E. L., Green, S. A., & Green, D. P. (2019). The contact hypothesis re-evaluated. *Behavioural Public Policy*, 3(2), 129–158.

Pettigrew, T. F., & Tropp, L. R. (2006). A meta-analytic test of intergroup contact theory. *Journal of Personality and Social Psychology*, 90(5), 751–783.

Pettit, P. (1997). *Republicanism: A theory of freedom and government*. Clarendon Press.

Rawls, J. (1993). *Political liberalism*. Columbia University Press.

Schelling, T. C. (1960). *The strategy of conflict*. Harvard University Press.

Sunstein, C. R. (2015). *Choosing not to choose: Understanding the value of choice*. Oxford University Press.