Policy Spaces as Organizational Schemes: A Semantic Theory of Electoral Competition

Robert J. Carroll*

March 18, 2026

Abstract

This paper develops a semantic theory of electoral competition in which voters and candidates hold *understandings*—epistemic-semantic pairs ordered by a complete lattice—rather than positions in a policy space. Policy space is not a primitive of the model but an *organizational scheme*: recoverable from the coalition structure of the electorate when that structure supports it, absent otherwise. I characterize recoverability via the *coalition mountain condition* (CMC) and show that a symmetric equilibrium is Downsifiable—interpretable as a median-voter outcome under some ordering of platforms—if and only if the CMC holds. The conditions jointly sufficient to ensure the CMC require policy-dominant credibility filters, single-dimensional policy semantics, single-peaked welfare, full policy receptivity, and trigger-free voters; the empirical record suggests all five are generically violated. When the CMC fails, the semantic lattice provides three results impossible to state in spatial theory: (i) the welfare-relevant McKelvey majority-preference orbit is structurally confined to the lattice interval between the meet and join of voter ideal understandings; (ii) when credibility filters are pairwise disjoint across voter types, vote-share optimization separates by type and the optimal broadcast is polysemic—message ambiguity is a necessary equilibrium property, not a communication failure; (iii) when voter ideals partition into incomparable blocs with one constituting a strict majority, competition forces equilibrium toward the majority bloc's ideal and the minority bloc's welfare shortfall is determined by the lattice-theoretic incomparability between the blocs, providing a structural account of majoritarian polarization without exogenous spatial geometry.

*Department of Political Science, University of Illinois at Urbana-Champaign. `rjc@illinois.edu`

Contents

1 Introduction.

The spatial model of electoral competition rests on an assumption so natural it is rarely examined: that voters and candidates share a common policy space in which political alternatives have fixed, publicly known positions. Downs (1957) built his account of democratic competition on this foundation. The logic is elegant and the predictions are clean.

Sixty years of empirical work on political behavior has documented, from multiple directions, that the shared policy space does not exist in the way the model requires. Converse (1964) showed that most citizens lack the ideological constraint spatial theory presupposes: beliefs on different policy dimensions are largely unrelated, and many citizens hold what amount to non-attitudes—responses to survey questions that bear no stable connection to any underlying policy preference. Zaller (1992) showed that citizens do not receive political communication as information to be weighed against prior beliefs but as material to be accepted or resisted depending on its consistency with prior dispositions. Taber and Lodge (2006) showed that political sophistication predicts entrenchment rather than accuracy: subjects exposed to information challenging their prior attitudes emerged more committed to those attitudes than before exposure. And the framing literature has established that the same political vocabulary carries categorically different meanings for citizens reasoning through different conceptual systems: "freedom," "responsibility," and "government" are not shared terms with shared meanings but contested representations whose content varies systematically across epistemic communities (Entman, 1993; Lakoff, 1996).

Methods such as DW-NOMINATE (Poole and Rosenthal, 1985) describe elite legislative voting reasonably well in formal institutional settings—though the fit varies across historical periods and leaves unexamined whether the recovered dimensions reflect legislators' policy positions or the institutional incentives that produce disciplined voting even absent a shared semantic space. As the spatial framework is transported to actual electoral competition—where candidates communicate with mass publics rather than trading votes with colleagues—the distortions Converse, Zaller, and Taber and Lodge document impose their full weight: the question is no longer whether elite voting fits a low-dimensional model but whether any shared space exists between candidates and the voters they address. The puzzle has sharpened in the decades since: affective polarization—the mutual animosity between party identifiers—has risen steadily since 1980 even as the extent of mass ideological divergence remains actively contested, indicating that partisan division is organized around something other than distance on a shared policy axis (Iyengar et al., 2019).

These findings are not peripheral anomalies to be accommodated by adding dimensions or introducing valence. They concern the semantic foundation of the spatial model: the assumption that political terms travel between candidates and voters with common meaning, that a voter's "position" is a stable object in a shared space, and that the policy axis organizes political competition rather than being organized by it. This paper takes that last phrase seriously. Rather than assuming a policy space, I ask when one exists—and what electoral competition looks like when it does not.

The framework models voters and candidates as holding *understandings*: pairs combining an epistemic stance (which political representations an agent can access and discriminate) with a semantics (what those representations mean to them). Understandings are ordered by a complete lattice whose meet and join operations formalize the notions of semantic common ground and pooled representational capacity. A *semantic update operator* governs how understandings change when agents encounter each other's framings, incorporating credibility filtering—the formal analog of Zaller's resistance mechanism—and a trigger-and-reversion component that formalizes Taber and Lodge's disconfirmation bias. Electoral competition takes place over this lattice, not over an exogenous policy axis.

The paper develops the *Downsification theorem* as its central analytical tool: a symmetric semantic equilibrium is interpretable as a median-voter outcome under some ordering of platforms if and only if a behavioral condition on the electorate's coalition structure—the *coalition mountain condition*—holds. The theorem functions as a diagnostic: it identifies when the Downsian spatial map applies, recovers the classical model as a limiting case, and precisely characterizes the conditions required—policy-dominant credibility filters, single-dimensional policy semantics, single-peaked welfare functions, full policy receptivity, and trigger-free voters. Each is individually demanding; their joint satisfaction is the exception rather than the rule.

The paper's three main results characterize what electoral competition looks like when the diagnostic returns a negative—the generic case in which the spatial map does not apply. When the CMC fails, the semantic lattice delivers findings that cannot be stated in the spatial framework. First, majority cycling does not produce unconstrained chaos: the welfare-relevant majority-preference orbit is confined to an interval in the lattice of understandings determined entirely by the distribution of voter ideal understandings, providing a structural bound that spatial theory cannot supply (Proposition 5.8). Second, when voters' credibility filters are heterogeneous enough that different types absorb disjoint components of a broadcast, vote-share optimization decomposes by type and the optimal campaign platform is polysemic—message ambiguity is a necessary equilibrium property, not a failure of political communication (Corollary 5.10). Third, when voter ideals partition into incomparable epistemic blocs with one constituting a strict majority, competition forces equilibrium platforms toward the majority bloc's ideal; the minority bloc is structurally excluded to a degree determined by the lattice-theoretic incomparability between the blocs—an account of majoritarian polarization that requires no exogenous spatial geometry (Corollary 5.12).

The paper connects to three literatures. In relation to *spatial models of voting* (Downs, 1957; Black, 1948; McKelvey, 1976), the framework endogenizes the policy space, recovers Downs as a limiting case, and extends McKelvey's chaos result with a structural bound. In relation to *behavioral models of political cognition* (Converse, 1964; Zaller, 1992; Taber and Lodge, 2006; Lodge and Taber, 2013), the connection is formal rather than analogical: the credibility filter, the trigger mechanism, and the baseline reversion are formal correlates of resistance, disconfirmation bias, and motivated reasoning, not loose metaphors. In relation to *formal models of belief revision* (Alchourrón, Gärdenfors and Makinson, 1985; Baltag and Renne, 2016; van Benthem, 2011), the update operator generalizes both AGM contraction and dynamic epistemic logic, with the critical departure that it permits downward moves in the lattice of understandings—epistemic regression, reframing-induced coarsening—that world-elimination mechanisms cannot represent.

Section 2 develops the representational framework. Section 3 develops the semantic update operator. Section 4 derives the Downsification theorem. Section 5 establishes the communication equilibrium and the three main results. Section 6 draws out the empirical predictions. Section 7 concludes.

2 The representational field.

The electoral results in Sections 4 and 5 require modeling political agents as reasoning *within* a shared field of representational resources, rather than *about* a fixed space of possible worlds. This section develops that model. The framework has three interlocking layers—an ontic layer of representations, an epistemic layer of agents' stances toward those representations, and a semantic layer of meanings agents attach to them. Together these generate a complete lattice of *understandings* whose meet and join operations formalize semantic common ground and pooled representational capacity.

2.1 What representations are—and aren't.

The standard starting point for models of political knowledge is the possible-worlds framework: fix a set Ω of mutually exclusive and exhaustive states of the world, exactly one of which is true, and represent agents' epistemic states as probability distributions or partitions over Ω. Disagreement, on this picture, concerns which state obtains—agents share the same representational space but locate themselves differently within it.

This starting point is unsuitable for the present project. Two difficulties arise. First, agents routinely hold representations that are false, simplified, or schematic, and these representations shape political behavior regardless of their accuracy. A voter who models Congress as a hierarchy in which a decisive leader controls legislative outcomes differs from one who understands committee dynamics and agenda-setting—but neither model need be true, and the difference between them is not captured by conditioning on the same state space. Second, the politically interesting variation is often not about *which state obtains* but about *how the domain is modeled*: one agent models congressional procedure as a Petri net of transitions and token flows, another as a principal-agent hierarchy, another as a narrative of heroes and villains contesting legitimacy. These are not different probability assignments over a shared space; they are representations drawn from categorically different schemas.

In place of the state space, we introduce a *representational field*.

Definition 2.1. The *representational field* $\mathcal{R}$ is a nonempty set whose elements are *representational resources*: models, schemas, symbols, narratives, formal structures, and other devices through which domains of political life can be conceived, described, or contested.[1]

Elements of $\mathcal{R}$ are not claims about how things actually are. They are the *tools available for making such claims*—or for framing, contesting, or refusing them. A Petri net model of congressional procedure, a conspiracy theory about legislative capture, and a spatial model of ideological competition are all elements of $\mathcal{R}$, regardless of their epistemic status. Nothing in the field is privileged as the correct or canonical representation.

The field $\mathcal{R}$ is *shared* in a specific and limited sense: it is the common ambient space within which all agents operate. Shared means neither

[1] The assumption that $\mathcal{R}$ is a set—rather than a proper class—is an ontological commitment: it ensures that the space of understandings $\mathscr{U}_{\mathcal{R}}$ (Definition 2.6) is a well-defined object within standard set theory, and that the complete lattice $(\mathscr{U}_{\mathcal{R}}, \preceq_{\mathscr{U}})$ can be constructed without impredicativity. Modeling $\mathcal{R}$ as itself subject to change—the endogenous field evolution extension—requires lifting this assumption; see section 7.

uniformly used nor uniformly interpreted. Different agents access different portions of $\mathcal{R}$, and even where two agents access the same element, they may assign it entirely different meanings. The word "immigration" is a shared representational token; what it means—which policy implications, evaluative associations, and social identifications it carries—varies across agents and is supplied by the semantic layer, not by membership in $\mathcal{R}$ (cf. Lakoff, 1996). What the field provides is a common formal reference space: a single universe of representational possibility within which differences in access *and* in interpretation can be located and compared.

2.2 Epistemic stances.

An agent's *epistemic stance* captures which parts of $\mathcal{R}$ are available to them and how finely they can discriminate among those parts.

Definition 2.2. An *epistemic stance* is a pair $e = (\mathcal{R}_e, \sim_e)$ where $\mathcal{R}_e \subseteq \mathcal{R}$ is the agent's *visible scope* and $\sim_e$ is an equivalence relation on *all of* $\mathcal{R}$ with a distinguished *invisible class* $I_e = \mathcal{R} \setminus \mathcal{R}_e$: all elements of I_e are mutually $\sim_e$-equivalent, forming a single equivalence class (empty when $\mathcal{R}_e = \mathcal{R}$). The *perceived representations* $\mathcal{R}_e/{\sim_e}$ are the non-invisible equivalence classes of the partition $\mathcal{R}/{\sim_e}$—the finest distinctions the agent can draw within their visible scope. The set of all epistemic stances is denoted $E_{\mathcal{R}}$.

Representational richness. It is natural to compare epistemic stances by the representational capacity they confer. Say that e_2 is *representationally richer* than e_1, written $e_1 \preceq_E e_2$, when $\sim_{e_2}$ refines $\sim_{e_1}$ as partitions of $\mathcal{R}$:

$$e_1 \preceq_E e_2 \iff \sim_{e_2} \subseteq \sim_{e_1} \text{ as binary relations on } \mathcal{R},$$

i.e., every $\sim_{e_2}$-class is contained in a $\sim_{e_1}$-class. This is the standard partition refinement order on the fixed set $\mathcal{R}$. The condition implies $\mathcal{R}_{e_1} \subseteq \mathcal{R}_{e_2}$: the invisible class of e_2 must be inside the invisible class of e_1 (it is still one of e_1's classes), so anything invisible to e_1 is also invisible to e_2, but not conversely.

A terminological remark is necessary. The relation $\preceq_E$ is *not* an order of epistemic accuracy, calibration, or rationality. Being higher in $\preceq_E$ means having more representational resources and finer discriminations—nothing more. A more elaborated representational toolkit may be epistemically harmful (finer but misleading distinctions), epistemically neutral (finer distinctions in domains irrelevant to the decision at hand), or both. For this reason we use the phrase *representationally richer* rather than "more informative," which carries connotations of Blackwell dominance and truth-localization that are inappropriate here.

Lemma 2.3. *$(E_{\mathcal{R}}, \preceq_E)$ is a partially ordered set.*

Incomparability. Two epistemic stances e_1 and e_2 are *incomparable* in $\preceq_E$ when neither $e_1 \preceq_E e_2$ nor $e_2 \preceq_E e_1$: each agent can represent something the other cannot.

Incomparability is the generic condition for agents who model the same political domain through categorically different schemas. As an illustration, consider two agents reasoning about the legislative process. Agent i draws on a set $\mathcal{R}_P \subset \mathcal{R}$ of Petri net representations, modeling procedure as a directed flow of tokens through transitions. Agent j draws on a disjoint set $\mathcal{R}_N \subset \mathcal{R}$ of narrative representations, understanding the same process through accounts of protagonists, obstacles, and institutional drama. Since $\mathcal{R}_P \cap \mathcal{R}_N = \emptyset$, the stances $e_i = (\mathcal{R}_P, \sim_P)$ and $e_j = (\mathcal{R}_N, \sim_N)$ are incomparable: each can represent things the other cannot, and neither is richer than the other in the sense of $\preceq_E$.

This is not a defect of the framework. It reflects the genuine representational pluralism of political life, in which agents can model the same domain through incommensurable schemas and neither is simply "less informed" than the other in any classical sense. The incomparability that obtains between a Petri-net reasoner and a narrative reasoner is an abstract version of a familiar empirical pattern. Lakoff's (1996) distinction between Strict Father and Nurturant Parent moral worldviews describes two conceptual systems through which citizens reason about national policy: the Strict Father frame organizes policy around discipline, self-reliance, and competitive order; the Nurturant Parent frame organizes the same domain through care, empathy, and communal support. Neither system is simply less informative than the other; each can represent things the other cannot, and their epistemic stances are incomparable in $\preceq_E$.

The fully disjoint case is, however, an idealization. In practice, incomparable stances often share some representational common ground: elements of $\mathcal{R}$ accessible to both agents that can serve as rough translations between schemas—*commensurability devices*. Entman's (1993) notion of a frame—a selection and salience operation on a shared field of facts—is the discourse-level analog of a commensurability device in the present sense: a publicly available representational element that can be accessed and used, differently, by agents with otherwise incomparable stances. Such shared elements come with an inherent cost: translation through them is lossy. A word-thinker can get a rough idea of what a given Petri net is doing without being able to distinguish it from a nearby one that the Petri-net-thinker regards as entirely different. How agents with incomparable or partially overlapping stances actually communicate—whether shared ground is found, built, or foreclosed—is a central question for the updating section to follow.

The epistemic lattice. Despite the prevalence of incomparability, the set $E_{\mathcal{R}}$ admits well-defined meet and join operations. For a family $\{e_j\}_{j\in J}$ with $e_j = (\mathcal{R}_j, \sim_j)$:

1. *Meet* $\bigwedge_{j\in J} e_j$: the most representationally rich stance lying below every e_j. Its partition of $\mathcal{R}$ is $\sim^{\wedge} = \mathrm{EqCl}\big(\bigcup_j \sim_j\big)$—merge any two resources that some e_j treats as indistinguishable, then close transitively. The visible scope is $\bigcap_j \mathcal{R}_{e_j}$: any resource invisible to even one stance ends up in the merged invisible class. The meet is the representational *common ground*.

2. *Join* $\bigvee_{j\in J} e_j$: the least representationally rich stance lying above every e_j. Its partition of $\mathcal{R}$ is $\sim^{\vee} = \bigcap_j \sim_j$—two resources are identified only when every stance already treats them as indistinguishable. The visible scope is $\bigcup_j \mathcal{R}_{e_j}$. The join is the representational *horizon*: what an agent fluent in all the schemas could access.

In the legislative example, the meet of e_i and e_j is $e_\perp$—they share no representational ground at all. Their join is $(\mathcal{R}_P \cup \mathcal{R}_N, \sim_{P\cup N})$—the stance of an agent at home in both Petri nets and narrative. An empty meet is politically meaningful: it signals that i and j have no common representational basis for communication, so that mutual understanding would require building new shared representations from scratch.

Lemma 2.4. $(E_{\mathcal{R}}, \preceq_E)$ *is a complete lattice with bottom element* $e_\perp$ *(empty visible scope:* $\mathcal{R}_{e_\perp} = \emptyset$*,* $I_{e_\perp} = \mathcal{R}$*) and top element* $e_\top = \big(\mathcal{R}, \mathrm{id}_{\mathcal{R}}\big)$ *(full visible scope, each resource its own class).*

The bottom $e_\perp$ represents total representational blindness; the top $e_\top$ represents perfect coverage of the entire field with maximal discrimination—a theoretical limiting case that no actual agent attains.

2.3 Semantics.

The epistemic layer specifies which representations an agent can access and distinguish. It says nothing about what those representations *mean*. Two agents with identical epistemic stances may attach entirely different meanings to the same representations: one reads congressional procedure as a mechanism for collective deliberation; another reads it as a site of elite capture and manufactured consent. This variation—independent of representational scope and discrimination—is the province of the semantic layer.

The meaning space. Let M be a nonempty set of *primitive meanings*. Like $\mathcal{R}$, the space M is left unstructured for generality.[2] In applications, its elements might include:

- *Policy implications*: "entails higher taxes," "requires redistribution," "commits troops"
- *Evaluative content*: "just," "threatening," "wasteful," "legitimate"
- *Social associations*: "for people like me," "an elite agenda," "a working-class cause"
- *Factual associations*: "correlated with crime," "associated with economic growth"
- *Affective content*: dread, resentment, solidarity, hope

The diversity of this list is intentional. Political representations rarely arrive with a single meaning; they come as bundles. The Affordable Care Act is not one thing in political discourse—it carries implications about health-care access, government intervention, personal responsibility, and partisan identity simultaneously, and different agents weight these differently. The set-valued range in the definition below reflects this directly.

[2] The generality of M is a deliberate choice that sets the framework's predictive resolution: without further structure on M, the conflict relation, trigger behavior, and welfare comparisons all depend on how the analyst specifies the parameters rather than being derivable from the model itself. The natural enrichment is a *consistency graph* on M: a reflexive symmetric relation $\sim_M \subseteq M \times M$ identifying which primitive meanings are mutually compatible and which are in tension. Under this enrichment the conflict set $\mathsf{Conf}(u_i, u_j)$ of Definition 3.1 becomes a derived quantity—determined jointly by μ_i, μ_j, and the consistency graph—rather than a free parameter. A complementary enrichment is a sender-side translation function $\gamma_{ji} : \mathcal{R}_{e_j}/{\sim_{e_j}} \to \mathcal{R}_{e_i}/{\sim_{e_i}} \cup \{\star\}$, dual to the receiver-side credibility filter α_{ij}: where α_{ij} encodes what i is willing to absorb, γ_{ji} encodes what j can express in vocabulary that i recognizes, with $\star$ marking untranslatable content. Together, a consistency graph on M and a sender-side translation parameter yield a considerably denser network of semantic constraints with tighter comparative-static predictions. We defer both enrichments to preserve the base framework's generality and to maintain a clean connection to the Downsian benchmark in section 5.

Semantic mappings. We now link epistemic stances to meanings via semantics.

Definition 2.5. For an epistemic stance $e = (\mathcal{R}_e, \sim_e)$, a *semantics* is a function

$$\mu_e : \mathcal{R}/{\sim_e} \longrightarrow 2^M,$$

assigning to each equivalence class of the full partition $\mathcal{R}/{\sim_e}$ a set of meanings, subject to $\mu_e(I_e) = \emptyset$ for the invisible class. (When $\mathcal{R}_e = \mathcal{R}$ there is no invisible class and this condition is vacuous.) A semantics μ_e is *e-consistent* when it satisfies this condition. Let M_e denote the set of all e-consistent semantics.

Invisible resources carry no meaning not because they are intrinsically meaningless, but because the agent cannot access them. If an agent later encounters such a resource—through a politician's platform or a conversation—it may enter scope and acquire meaning; but until then $\mu_e(I_e) = \emptyset$ by definition.

Two agents sharing the same epistemic stance e but holding different semantics $\mu_e \neq \mu'_e$ exhibit *pure semantic disagreement*: they access the same representations and draw the same distinctions, but interpret those representations differently. This is a distinct kind of political disagreement from epistemic incomparability. In practice, both forms of disagreement coexist and interact—which is the subject of the reframing discussion below.

2.4 Understandings.

An agent's *understanding* combines both layers.

Definition 2.6. An *understanding* is a pair $u = (e, \mu_e)$ consisting of an epistemic stance $e \in E_{\mathcal{R}}$ and an e-consistent semantics $\mu_e \in M_e$. The *space of understandings* is

$$\mathscr{U}_{\mathcal{R}} = \bigsqcup_{e \in E_{\mathcal{R}}} M_e = \{(e, \mu_e) \mid e \in E_{\mathcal{R}},\ \mu_e \in M_e\},$$

where $\bigsqcup$ denotes disjoint union: each epistemic stance defines a fiber of compatible semantics, and $\mathscr{U}_{\mathcal{R}}$ assembles all fibers.

Semantic transport. To compare understandings at different epistemic levels, meanings must be translatable between levels. When $e_1 \preceq_E e_2$, each fine perceived representation $D \in \mathcal{R}_{e_2}/{\sim_{e_2}}$ is contained in a unique coarse perceived representation $C \in \mathcal{R}_{e_1}/{\sim_{e_1}}$. This containment supports two canonical transport operations:

$$\begin{aligned}\mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1})(D) &= \mu_{e_1}(C),\\ \mathrm{Crs}_{e_2 \to e_1}(\mu_{e_2})(C) &= \bigcup_{D \subseteq C} \mu_{e_2}(D).\end{aligned}$$

Refinement $\mathrm{Ref}_{e_1 \to e_2}$ extends a coarse semantics to a finer epistemic level by distributing each meaning uniformly across the newly distinguished subclasses. When $e_1 \prec_E e_2$ and D is a newly-visible e_2-class that was inside the invisible class I_{e_1}, the containing e_1-class is I_{e_1}, and $\mu_{e_1}(I_{e_1}) = \emptyset$ by definition, so newly-visible resources inherit empty meaning automatically. *Coarsening* $\mathrm{Crs}_{e_2 \to e_1}$ projects a fine semantics back to a coarser level by pooling the meanings of all subclasses within each coarse class. These operations satisfy the Galois connection $\mathrm{Crs} \dashv \mathrm{Ref}$: for any $\mu_{e_1} \in M_{e_1}$ and $\mu_{e_2} \in M_{e_2}$,

$$\mathrm{Crs}_{e_2 \to e_1}(\mu_{e_2}) \preceq_M \mu_{e_1} \quad \Longleftrightarrow \quad \mu_{e_2} \preceq_M \mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1}).$$

Both directions reduce to the pointwise equivalence: $\bigcup_{D \subseteq C} \mu_{e_2}(D) \subseteq \mu_{e_1}(C)$ if and only if $\mu_{e_2}(D) \subseteq \mu_{e_1}(C)$ for every $D \subseteq C$.

The order on understandings. Using the refinement operator to compare meanings across epistemic levels, we extend the order on epistemic stances to the full space of understandings:

$$(e_1, \mu_{e_1}) \preceq_{\mathscr{U}} (e_2, \mu_{e_2}) \quad \Longleftrightarrow \quad e_1 \preceq_E e_2 \text{ and } \mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1}) \preceq_M \mu_{e_2},$$

where $\preceq_M$ is the pointwise inclusion order on M_{e_2}: $\nu \preceq_M \nu'$ iff $\nu(D) \subseteq \nu'(D)$ for all $D \in \mathcal{R}_{e_2}/{\sim_{e_2}}$. Unpacking the refinement: the condition requires $\mu_{e_1}(C) \subseteq \mu_{e_2}(D)$ for every fine class D at e_2's level that is contained in the coarse class C at e_1's level. An understanding (e_2, μ_{e_2}) lies above (e_1, μ_{e_1}) when it is epistemically richer *and* each fine distinction D in e_2's partition carries at least all the meanings e_1 attributes to the encompassing coarse class C.

The same terminological caution applies: the order $\preceq_{\mathscr{U}}$ measures *representational richness*, not epistemic virtue. A higher understanding has more representational resources and richer meanings. It does not have more accurate, more rational, or more politically appropriate representations.

Theorem 2.7. *$(\mathscr{U}_{\mathcal{R}}, \preceq_{\mathscr{U}})$ is a complete lattice. For any family $\{(e_j, \mu_{e_j})\}_{j \in J} \subseteq \mathscr{U}_{\mathcal{R}}$, writing $e^{\wedge} = \bigwedge_j e_j$ and $e^{\vee} = \bigvee_j e_j$, the meet and join are*

$$\bigwedge_{j \in J} (e_j, \mu_{e_j}) = \big(e^{\wedge},\ \mu_{\wedge}\big), \qquad \bigvee_{j \in J} (e_j, \mu_{e_j}) = \big(e^{\vee},\ \mu_{\vee}\big),$$

where

$$\mu_\wedge(C) = \bigcap_{j\in J} \bigcap_{\substack{D\in\mathcal{R}_{e_j}/\sim_{e_j} \\ D\subseteq C}} \mu_{e_j}(D), \qquad \mu_\vee(D) = \bigcup_{j\in J} \mathrm{Ref}_{e_j\to e^\vee}(\mu_{e_j})(D),$$

for $C \in \mathcal{R}_{e^\wedge}/\sim_{e^\wedge}$ *and* $D \in \mathcal{R}_{e^\vee}/\sim_{e^\vee}$ *respectively.*

The lattice operations have direct political interpretations. The meet $u_i \wedge u_j$ is the *semantic common ground* between two understandings—what they share after all representational differences are resolved by taking the most conservative position on each. The join $u_i \vee u_j$ is their *pooled representational toolkit*—the understanding available to an agent fluent in both schemas. Both operations recur throughout the analysis: the vote function in Section 4 uses the meet to measure how much common ground a candidate's broadcast shares with a voter's ideal; the communication equilibrium in Section 5 uses the join to bound how far updating can proceed.

Epistemic regression. A feature of the framework that has no counterpart in classical information-theoretic models is that understanding can become *less* rich over time. An agent who adopts a coarser epistemic frame—replacing fine-grained distinctions with a crude binary partition—moves strictly *downward* in $\preceq_{\mathscr{U}}$ even if new representational content is introduced along the way. This is *epistemic regression*: the loss of representational distinctions.

Classical Bayesian conditioning can only add information—it cannot produce an agent who draws fewer distinctions about the world than they did before. Epistemic regression is therefore invisible to the Bayesian apparatus. In the representational framework it corresponds naturally to a downward move in the lattice, and as the next subsection shows, it is not just a formal possibility but a politically observable outcome.

2.5 Reframing: the paradigm interaction.

The payoff of maintaining both epistemic and semantic layers appears when we consider reframing. Reframing is the process by which a political actor presents familiar phenomena through an unfamiliar representational lens, reshaping how they are understood. It is the paradigm case of a joint epistemic-semantic transformation, and it motivates much of the machinery to follow. Entman (1993) characterized a frame as selecting and making salient aspects of a perceived reality; the present formalism gives that characterization explicit algebraic structure, identifying the epistemic and semantic moves that together constitute a successful reframe.

A successful reframe typically involves two simultaneous moves:

1. *Semantic transformation*: representations the agent already holds are assigned new or shifted meanings.
2. *Epistemic transformation*: new representations may enter scope, or previously fine-grained distinctions may be collapsed into cruder equivalence classes.

Neither move alone constitutes reframing. A purely epistemic update—new representations, meanings unchanged—adds content without changing how the world is interpreted. A purely semantic update—same representations, meanings reassigned—changes interpretation without introducing new representational elements. Genuine reframing involves both, typically in a coordinated way.

The immigration case. Suppose a politician seeks to shift how voters understand immigration. Voters begin with an epistemic stance e_0 distinguishing several immigration categories—high-skill workers, low-skill workers, asylum seekers, temporary labor visa holders, undocumented entrants—and a semantics μ_{e_0} that assigns meanings including "economic contributor," "humanitarian obligation," and "procedural status."

The politician introduces crime statistics disaggregated by immigration status. This is an epistemic act: a new representation $r_{\text{crime}} \notin \mathcal{R}_{e_0}$ enters scope. It is simultaneously a semantic act: the politician assigns r_{crime} the meaning "threat," and leverages this new element to reassign meanings to existing representations. Previously, "undocumented entry" meant something like "procedural noncompliance"; it now acquires the meaning "physical threat." Previously, "low-skill immigration" was primarily an economic category; it now carries security connotations.

The updated understanding (e_1, μ_{e_1}) may in fact be *less* representationally rich than the original (e_0, μ_{e_0}). If voters, following the politician's lead, collapse their distinctions among immigration categories into a single binary—"legal" versus "illegal"—then their epistemic stance has regressed: $e_1 \prec_E e_0$ despite the introduction of new content. The result is an understanding that is simultaneously epistemically cruder and semantically reorganized around a threat frame. Both the regression and the reorganization are coherent in the lattice; neither is captured by conditioning. This is the formal analog of what Lakoff (1996) identifies as the Nation-as-Family conceptual metaphor applied to immigration: the same referential resources are processed through a security frame or a humanitarian frame, with systematically different meanings attached to each category—and the security frame achieves its effect precisely by collapsing fine distinctions into a cruder binary.

Contrast with Bayesian updating. In a Bayesian model, a voter receiving signal ι updates from prior π to posterior $\pi(\cdot \mid \iota)$. The state space Ω is fixed throughout. Bayesian updating cannot:

- Introduce representational categories not already in Ω
- Collapse previously distinct categories into coarser equivalence classes (conditioning can only add information, never subtract it)
- Shift the *meaning* of a category without changing its probability—that is, it cannot represent the difference between "undocumented entry as procedural noncompliance" and "undocumented entry as physical threat" if these are assigned the same probability

All three operations occur in the immigration case. The representational framework handles them within a single formal apparatus—as transformations of the understanding (e, μ_e) in the lattice $\mathscr{U}_{\mathcal{R}}$. Section 3 develops the updating formalism that makes this precise.

3 Semantic updating.

The framework developed in section 2 describes the static geometry of understandings: how representational resources are accessed, how meanings are assigned, and how understandings are ordered. This section introduces dynamics. We specify how an agent's understanding changes upon encountering another's.

Standard belief revision cannot do this job. Bayesian conditioning operates within a fixed state space and can only refine or preserve the agent's information partition: it eliminates states from consideration but cannot collapse previously distinct states into a single equivalence class, introduce representational categories not already in the state space, or shift the meaning of a category without changing its probability. AGM contraction and revision operate on propositional belief sets and inherit the same fixed-ontology assumption. What is needed is an update operator that acts on the full epistemic-semantic pair (e, μ_e), that can move an agent in any direction in $\mathscr{U}_{\mathcal{R}}$—upward, downward, or sideways—and whose behavior is governed by a richer set of parameters than a single credence function.

The need is not merely formal. Taber and Lodge (2006) present experimental evidence of a phenomenon—disconfirmation bias—that Bayesian updating cannot describe. When subjects encountered information challenging their prior political attitudes, they did not simply fail to absorb it; they generated counterarguments, dismissed the source, and emerged from the encounter with attitudes measurably more entrenched than before exposure. The update, in the disconfirmation case, is not a smaller probability mass on some proposition. It is an *impoverishment of semantic structure*: the agent's representational resources for entertaining the challenging content are actively suppressed in the course of resisting it. This is what the operator developed here must capture: an asymmetric process in which encounters with incongruent framings can leave an agent representationally poorer, not merely unchanged.

3.1 Conflict structure.

We first equip the meaning space M with a notion of incompatibility.

Definition 3.1. A *conflict relation* is a symmetric, irreflexive relation $\mathsf{Cf} \subseteq M \times M$. A set of meanings $A \subseteq M$ is *conflict-free* if no pair $(m, m') \in \mathsf{Cf}$ has both $m, m' \in A$. A semantics μ_e is *coherent* if $\mu_e([r])$ is conflict-free for every $[r] \in \mathcal{R}_e/{\sim_e}$.

The conflict relation is left as a parameter of the model. Its natural content comes from two sources. For truth-apt meanings $M_T \subseteq M$ equipped with a negation map $\neg : M_T \to M_T$, the pairs $(m, \neg m)$ are the canonical conflicts: an agent should not simultaneously hold a truth-apt claim and its negation for the same perceived representation. For evaluative, affective, and social-association meanings, conflicts may be present or absent depending on the application; the framework imposes no requirement. Incoherence of the latter kind—holding incommensurable values, carrying tensions between social associations—is representationally permissible.

Given two understandings $u_i = (e_i, \mu_i)$ and $u_j = (e_j, \mu_j)$, write

$$F_{ij}([r]) := \big\{ m \in M : \exists\, m' \in \mathrm{Crs}_{e_j \to e_i}(\mu_j)([r]),\ (m, m') \in \mathsf{Cf} \big\}$$

for the set of meanings whose adoption by i at $[r]$ would conflict with j's translated content. The *conflict set* of u_j against u_i is then

$$\mathsf{Conf}(u_i, u_j) = \big\{ [r] \in \mathcal{R}_{e_i}/{\sim_{e_i}} : \mu_i([r]) \cap F_{ij}([r]) \neq \emptyset \big\}.$$

This is i's perspective: it asks whether what j is saying, translated into i's representational language, would require i to hold conflicting meanings simultaneously. The conflict set is empty when the two understandings are semantically compatible at i's epistemic level.

3.2 Source credibility.

The degree to which agent i absorbs meanings from agent j is governed by a *source credibility* structure. Credibility is representation-specific: what matters is which *types* of meaning an agent is willing to absorb from a given source—policy implications, evaluative associations, affective signals, social-group content. The credibility function encodes this type-specific receptivity as a *meaning filter*: a subset $\alpha \subseteq M$ specifying which categories of meaning i is willing to receive. The filtering operation is intersection: i absorbs from j exactly $\alpha \cap A$, the portion of j's content A that falls within i's receptivity set.[3] Zaller's (1992) Resistance Axiom (A2)—that citizens resist messages inconsistent with their prior dispositions—corresponds to the extreme case $\alpha_{ij} = \emptyset$ for sources j whose framing conflicts with i's prior commitments, blocking transmission entirely.

Definition 3.2. A *source credibility function* for agent i toward agent j is a map

$$\alpha_{ij} : \mathcal{R}_{e_j}/\!\sim_{e_j} \;\longrightarrow\; 2^M,$$

assigning to each of j's perceived representations a meaning filter $\alpha_{ij}([r]) \subseteq M$.

Credibility is representation-specific: agent i may be highly receptive to j's affective associations while entirely filtering out j's policy implications, or vice versa. Full receptivity at $[r]$ is $\alpha_{ij}([r]) = M$; complete rejection is $\alpha_{ij}([r]) = \emptyset$.

3.3 The update operator.

Let $u_i = (e_i, \mu_i)$ and $u_j = (e_j, \mu_j)$. The update of u_i upon encountering u_j is governed by a *trigger threshold* $\tau_i \in \mathbb{N} \cup \{\infty\}$ and source credibility α_{ij}. The trigger threshold is a tolerance for conflict: the forward update fires when the number of conflicting representation classes is at most τ_i; when it exceeds τ_i, the agent reverts. $\tau_i = 0$ means any conflict triggers reversion; $\tau_i = \infty$ means the agent never reverts regardless of conflict size.

[3] This filtering structure is an instance of a quantale action on 2^M; the general abstraction is available if compositionality across chains of sources is needed, but nothing in this paper requires it.

Definition 3.3. The *update of u_i by u_j*, parameterized by (α_{ij}, τ_i), is

$$T_{\alpha_{ij},\tau_i}(u_i, u_j) = \begin{cases} T_+(u_i, u_j) & \text{if } \left|\mathsf{Conf}(u_i, u_j)\right| \le \tau_i, \\ \beta(u_i, u_j) & \text{otherwise,} \end{cases}$$

where T_+ is the *forward update* and β is the *baseline*, defined below.

The forward update. The forward update operates in two stages: a semantic absorption at i's current epistemic level, followed by an optional epistemic expansion.

The semantic stage takes the meanings of j's representations, coarsened to the *epistemic common ground* $e^\wedge = e_i \wedge e_j$, filters them through i's credibility function at that level, and adds the result to i's existing semantics. For each $[r] \in \mathcal{R}_{e_i}/\sim_{e_i}$, let $[R]$ denote the unique $e^\wedge$-class containing $[r]$ (well-defined since $e^\wedge \preceq_E e_i$):

$$\mu_i^+([r]) \;=\; \mu_i([r]) \;\cup\; \Big(\alpha_{ij}^{e^\wedge}([R]) \;\cap\; \mathrm{Crs}_{e_j\to e^\wedge}(\mu_j)([R])\Big).$$

Here $e^\wedge \preceq_E e_j$ by definition of the meet, so $\mathrm{Crs}_{e_j\to e^\wedge}$ is well-defined. The meet-level credibility filter is

$$\alpha_{ij}^{e^\wedge}([R]) \;=\; \bigcap_{\substack{D\in\mathcal{R}_{e_j}/\sim_{e_j} \\ D\subseteq[R]}} \alpha_{ij}(D),$$

the most conservative filter consistent with α_{ij}: agent i accepts meaning from $[R]$ only if it would accept it from every e_j-subclass of $[R]$. When $[R]$ is the invisible class of $e^\wedge$ (i.e., $[r]$ lies outside the shared visible scope $\mathcal{R}_{e_i} \cap \mathcal{R}_{e_j}$), the intersection over an empty index set gives $\alpha_{ij}^{e^\wedge}([R]) = M$ (using the convention that $\bigcap\emptyset = M$, the full universe of meanings), but $\mathrm{Crs}_{e_j\to e^\wedge}(\mu_j)([R]) = \emptyset$ since $\mu_j(I_{e_j}) = \emptyset$, so the absorption term vanishes: $\mu_i^+([r]) = \mu_i([r])$. Agent i can thus only absorb meanings from j on the shared representational scope; classes outside j's reach are unaffected by the encounter.

When $e_i \preceq_E e_j$ (receiver at least as coarse as source), $e^\wedge = e_i$ and the formula reduces to $\mu_i^+([r]) = \mu_i([r]) \cup (\alpha_{ij}([r]) \cap \mathrm{Crs}_{e_j\to e_i}(\mu_j)([r]))$—the natural case for campaign communication, in which voters have less elaborated stances than candidates on many dimensions. The updated semantics is μ_i^+, with understanding (e_i, μ_i^+).

The epistemic stage allows i's epistemic stance to expand, subject to the constraint that no single interaction can take i beyond the pooled representational resources of both parties:

$$e_i \preceq_E e_i' \preceq_E e_i \vee e_j.$$

The mechanism by which e_i' is selected within this range is left as a parameter; the bound $e_i \vee e_j$ is the only constraint the framework imposes. The forward update is then $T_+(u_i, u_j) = (e_i', \mathrm{Ref}_{e_i \to e_i'}(\mu_i^+))$.

The baseline. When the conflict in $\mathsf{Conf}(u_i, u_j)$ exceeds the threshold τ_i, the update reverts to the *baseline*—the largest understanding below u_i that avoids the conflicting content:

$$\beta(u_i, u_j) = \bigvee S_{ij}, \quad \text{where}$$
$$S_{ij} = \{u \preceq_{\mathscr{U}} u_i : \forall [r] \in \mathsf{Conf}(u_i, u_j),\ \mathrm{Ref}_{e_u \to e_i}(\mu_u)([r]) \cap F_{ij}([r]) = \emptyset\}.$$

This is AGM contraction, restated in $\mathscr{U}_{\mathcal{R}}$[4]: among all understandings weakly below u_i, take the richest one that has removed whatever meanings would conflict with u_j's content. Several features are immediate. First, $\beta(u_i, u_j) \preceq_{\mathscr{U}} u_i$: the baseline is always weakly below the current understanding. Second, the baseline depends on both arguments: what you revert *to* depends on what triggered the reversion. Third, (e_i, μ_i) itself is in the set being joined iff u_i is already coherent with respect to u_j—which is exactly the condition that the trigger does not fire. Fourth, and most consequentially:

Proposition 3.4. *If* $\mathsf{Conf}(u_i, u_j) \neq \emptyset$*, then* $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$*.*

The inequality is strict because every element of S_{ij} has its refinement to e_i's level disjoint from $F_{ij}([r])$ at each conflicting class $[r]$; their join therefore also avoids $F_{ij}([r])$; but $\mu_i([r]) \cap F_{ij}([r]) \neq \emptyset$ by definition of conflict, so the join cannot recover to u_i.

3.4 Parameter dependence.

The update operator T_{α_{ij}, τ_i} is parameterized by the pair (α_{ij}, τ_i). The behavior of the model as these vary exhibits clean monotone structure.

[4] More precisely, $\beta(u_i, u_j) = \bigvee S_{ij}$ is in the spirit of maxichoice contraction (Alchourrón, Gärdenfors and Makinson, 1985): it selects a maximally informative belief state that avoids the conflicting content. Full partial-meet contraction takes an intersection of selected maximal subsets; the present operator takes the join of selected understandings in $\mathscr{U}_{\mathcal{R}}$. The structural parallel—retain as much as possible while removing the offending content—is the key point.

Credibility. For fixed τ_i, higher credibility α_{ij} (larger filter sets in the 2^M case) produces richer forward updates: more of j's meanings are absorbed. Formally, if $\alpha_{ij}([r]) \subseteq \alpha'_{ij}([r])$ for all $[r]$ (pointwise inclusion), then $\mu_i^+([r]) \subseteq \mu_i^{+\prime}([r])$ for all $[r]$, and hence $T_+(u_i, u_j) \preceq_{\mathscr{U}} T'_+(u_i, u_j)$ (when the trigger does not fire under either).

Threshold. For fixed α_{ij}, a lower threshold τ_i makes reversion more likely. Since $\beta(u_i, u_j) \preceq_{\mathscr{U}} u_i \preceq_{\mathscr{U}} T_+(u_i, u_j)$ whenever the forward update is nontrivial, a more hair-trigger agent ends up at a weakly lower understanding than a more tolerant one.

Epistemic cost of reversion. Reversion is not neutral. Proposition 3.4 says that when the trigger fires, the agent does not merely fail to update—they revert to a strictly impoverished understanding. The content removed in the contraction $u_i \mapsto \beta(u_i, u_j)$ is the cost of the interaction, paid even though the triggering content is rejected. An agent who successfully defends against a coherence-violating presentation ends up representationally poorer than before the encounter.

This is the formal correlate of the disconfirmation bias documented by Taber and Lodge (2006). Their experimental subjects, exposed to information challenging prior political attitudes, did not simply fail to update; they generated counterarguments, dismissed the source, and showed increased attitude entrenchment relative to unexposed controls. In the present model, this is Proposition 3.4: the encounter does not leave u_i unchanged but strictly impoverishes it.

Remark 3.5 (Sophistication paradox). The most politically sophisticated subjects in Taber and Lodge's experiments—those with the richest prior understandings—showed the strongest disconfirmation effect. The formal explanation is direct: a richer understanding provides more meanings that can conflict with an incoming frame, so the trigger fires more readily and the baseline drops further. The Downsification conditions of section 4.7 are therefore most severely violated precisely by the voters who most resemble the rational actor of spatial theory. Sophistication, in the semantic framework, is not a route to Downsian behavior but an obstacle to it.

3.5 Special cases.

Warm glow. Suppose $\alpha_{ij}([r]) \subseteq M_{\text{affect}} \cup M_{\text{social}}$ for all $[r]$—agent i filters j's meanings down to affective and social-association content only—and suppose the epistemic expansion is minimal ($e'_i = e_i$). The forward update then modifies μ_i exclusively in the affective and social dimensions, regardless of whatever policy, factual, or evaluative content μ_j carries. The voter absorbs the emotional register and group-identity signals of the candidate's understanding while remaining unaffected by its policy content. This is the warm glow case: the campaign operates entirely through the credibility filter, and the filter admits only affect (cf. Zaller, 1992; Lodge and Taber, 2013). Lodge and Taber's (2013) primacy-of-affect postulate—that affective reactions precede and dominate deliberative response—implies that for most voters α_{vc} is closer to M_{affect} than to M_{policy}, making warm glow the generic mode of campaign reception rather than a special case.

Policy campaign. Suppose $\alpha_{ij} = M$ (full receptivity across all meaning types) and the epistemic expansion is maximal ($e'_i = e_i \vee e_j$). The forward update then expands i's epistemic scope up to the pooled representational resources and absorbs the full semantic content of j's understanding. Downsification—the recovery of classical electoral competition as a limiting case—requires something close to this: voters whose credibility filters admit policy content and whose epistemic stances respond to candidates' platforms. This is a demanding set of conditions, and their joint satisfaction is the exception rather than the rule. Converse (1964) showed that most citizens lack the ideological constraint that would allow policy content to travel coherently from candidate to voter; for the majority of the electorate, the policy campaign case is a theoretical limiting case that actual campaigns do not approximate.

Bayesian conditioning. Suppose the epistemic stance is fixed (no expansion: $e'_i = e_i$), each semantics assigns a single number to each perceived class—a probability $p_i([r]) \in [0, 1]$ satisfying $\sum_{[r]} p_i([r]) = 1$—credibility is full ($\alpha_{ij}([r]) = M$ for all $[r]$), and the trigger is deactivated ($\tau_i = \infty$). Under these conditions the conflict set is empty: single-valued probability assignments cannot conflict with each other in the sense of Definition 3.1, because no pair (m, m') with $m = p_i([r])$ and $m' = p_j([r])$ can be simultaneously in Cf (the conflict relation is defined on meanings, and a probability value does not negate another probability value). The forward update then adds j's probability at each class to i's: $\mu_i^+([r]) = \{p_i([r])\} \cup \{p_j([r])\}$, expanding i's meaning set to a two-element set. To recover Bayesian conditioning as a special case—rather than a meaning-set expansion—one further restriction is needed: interpret j's broadcast as an *event* $E = \{[r] : p_j([r]) > 0\}$ (the support of j's distribution), and require that absorption replaces i's distribution rather than augmenting it: $p_i^+([r]) = p_i([r] \mid E) = p_i([r])/p_i(E)$ for $[r] \in E$ and 0 otherwise. The full update then recovers standard Bayesian conditioning: i observes j's epistemic stance as an event in i's prior, and updates by conditioning. The semantic framework thus contains classical Bayesian updating as the special case: fixed epistemic level, no conflict, single-valued semantics, full credibility, and a replacement (rather than augmentation) interpretation of absorption. The augmentation interpretation—the base case of Definition 3.3—is the strictly more general operation, allowing i to hold multiple meanings simultaneously rather than collapsing to a posterior.

4 Semantic electoral competition and Downsification.

The framework in Section 2 and the update operator in Section 3 make it possible to ask: what does electoral competition look like when voters and candidates are modeled as holding *understandings* rather than probability distributions over policy positions? This section provides the answer in two stages. The first defines the objects of competition and the notion of equilibrium. The second—and the section's central contribution—asks when a symmetric semantic equilibrium can be *Downsified*: assigned a total order on the strategy space that makes it interpretable as a classical median-voter outcome. The answer to that question, we argue, is the correct way to understand the policy dimension itself: not as the primitive space in which competition occurs, but as an organizational scheme that a semantic electorate either supports or does not.

4.1 The electoral setup.

The strategy space X is a nonempty finite set. Elements of X are *platforms*: abstract objects with no pre-given structure—no order, no topology, no metric. Any ordering on X will be recovered from the electoral dynamics, not imposed on them.

Definition 4.1. A *semantic election* is a tuple

$$\mathcal{E} = \big(C,\ V,\ X,\ \phi,\ \{u_v^0,\ u_v^*\}_{v \in V},\ \{\alpha_{vc}\}_{(v,c) \in V \times C},\ \{\tau_v\}_{v \in V}\big),$$

where C is a finite set of *candidates*, V is a finite set of *voters*, X is the *platform space*, $\phi : X \to \mathscr{U}_{\mathcal{R}}$ is the *broadcast map* assigning to each platform an understanding that the candidate broadcasts, u_v^0, $u_v^* \in \mathscr{U}_{\mathcal{R}}$ are the initial and *ideal* understanding of voter v, $\alpha_{vc} : \mathcal{R}_{e_c}/{\sim_{e_c}} \to 2^M$ is voter v's source credibility function toward candidate c (Definition 3.2), and $\tau_v \in \mathbb{N} \cup \{\infty\}$ is voter v's trigger threshold.

The ideal understanding u_v^* is the analogue of the ideal point in classical spatial voting. It is not a point in a policy space but an understanding in $\mathscr{U}_{\mathcal{R}}$: the most semantically rich state the voter would prefer to occupy after the campaign. Unlike an ideal policy, u_v^* can incorporate evaluative, affective, and social-association meanings alongside policy content—capturing the multidimensional character of political preferences documented in the empirical literature on belief systems (cf. Converse, 1964; Lodge and Taber, 2013). On Lodge and Taber's (2013) account, political preferences are constituted primarily by affective associations and on-line tallies formed from social identity signals, not deliberative policy evaluation; the ideal understanding u_v^* accommodates this structure directly, without forcing non-policy preferences into the geometry of a policy space.

4.2 Campaign updating and voter response.

A *campaign* is a profile of platforms $(x_c)_{c \in C}$ chosen by the candidates. Upon observing the campaign, voter v forms a *post-campaign understanding* for each candidate separately, by applying the update operator of Definition 3.3 to the candidate's broadcast $\phi(x_c)$:

$$\hat{u}_v^c \;=\; T_{\alpha_{vc},\,\tau_v}\big(u_v^0,\; \phi(x_c)\big), \qquad c \in C.$$

This is the understanding voter v would hold after updating from candidate c's broadcast. The full range of updating behavior from section 3—absorption filtered by credibility, forward expansion of epistemic scope, reversion to baseline under excessive conflict—is inherited without modification.

Two features are worth making explicit. First, post-campaign understandings are voter-specific: voters with different initial understandings, credibility functions, or thresholds respond differently to the same campaign. A candidate's broadcast is not received uniformly. Second, post-campaign understandings are candidate-specific: $\hat{u}_v^{c_1}$ and $\hat{u}_v^{c_2}$ can differ substantially, since the update operator depends on the source understanding $\phi(x_c)$. A campaign that brings one voter closer to their ideal may move another voter further from theirs.

Remark 4.2 (Campaign priming). The base formula treats each candidate's broadcast as processed from the fixed initial state u_v^0. A richer formulation accounts for *opponent priming*: voter v, having absorbed candidate c''s broadcast $\phi(x_{c'})$ before encountering candidate c's, arrives at c's message in the state $T_{\alpha_{vc'},\tau_v}(u_v^0, \phi(x_{c'}))$ rather than u_v^0. At a symmetric equilibrium $x^* = x_A = x_B$, the relevant voter state when evaluating a deviation by A is the primed state $T_{\alpha_{vB},\tau_v}(u_v^0, \phi(x^*))$, not u_v^0—candidates must throw the ball to where the receiver *will be*, not where they started. The Downsification conditions of section 4.7 should be understood as holding at these primed voter states; the base formula is the limiting case where priming is negligible.

4.3 Electoral welfare and voting.

To rank candidates, voters compare their post-campaign understandings against their ideals. The comparison is mediated by an *electoral welfare function*.

Definition 4.3. An *electoral welfare function* is a map $\mathcal{W} : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$ satisfying

$$u \prec_{\mathscr{U}} u' \implies \mathcal{W}(u) < \mathcal{W}(u').$$

Voter v evaluates candidate c by the welfare generated at the *meet* of the post-campaign understanding and their ideal:

$$w_v\big(c,\, (x_c)_{c\in C}\big) \;=\; \mathcal{W}\big(\hat{u}_v^c \,\wedge\, u_v^*\big).$$

The meet $\hat{u}_v^c \wedge u_v^*$ is the semantic common ground between what candidate c's campaign produces in voter v and what voter v ideally wants to understand.[5] By strict monotonicity of $\mathcal{W}$, richer common ground translates directly into higher welfare. The meet exists for all pairs by the complete lattice theorem (theorem 2.7).

Remark 4.4 (Meet as canonical welfare aggregator). The meet is the canonical choice for this role. Any voter welfare function satisfying two conditions—(a) it depends only on what the broadcast and the voter's ideal *share*, not on either taken alone, and (b) it is monotone in that shared content—must be

[5] This formula has a structural analog in the on-line processing model of Lodge and Taber (2013): the vote-relevant summary evaluation is formed incrementally from each campaign encounter, with each encounter contributing whatever affective charge is activated and passes a relevance threshold. In the present model, the meet captures the same intuition—the voter takes credit only for the semantic content that both the campaign and their prior ideal contain. The OL model weights encounters by cognitive accessibility; the present model weights them by the meet, which identifies what the two understandings actually share.

a function of $\hat{u}_v^c \wedge u_v^*$. The meet is the largest understanding weakly below both arguments: it solves $\max\{u \in \mathscr{U}_{\mathcal{R}} : u \preceq_{\mathscr{U}} \hat{u}_v^c \text{ and } u \preceq_{\mathscr{U}} u_v^*\}$, so any $\mathcal{W}$-maximizing agent choosing how much of a candidate's broadcast to credit will land at the meet. Condition (a) rules out $\mathcal{W}(\hat{u}_v^c)$ directly, which ignores the voter's ideal entirely. Condition (b) rules out distance-based welfare functions, which require a metric structure on $\mathscr{U}_{\mathcal{R}}$ that the present framework does not impose.

Definition 4.5. Voter v's *vote function* is

$$c^*\big(v,\, (x_c)_{c\in C}\big) \;\in\; \arg\max_{c\in C}\, \mathcal{W}\big(\hat{u}_v^c \,\wedge\, u_v^*\big),$$

with ties broken by a fixed rule. The *vote share* of candidate c is $\mathcal{V}_c\big((x_c)_{c\in C}\big) = \big|\{v \in V : c^*(v,\cdot) = c\}\big|/|V|$.

4.4 Candidate strategy and symmetric equilibrium.

Each candidate $c \in C$ chooses $x_c \in X$ to maximize vote share, taking the other candidates' platforms as given.

Definition 4.6. A campaign profile $(x_c^*)_{c\in C}$ is a *symmetric semantic equilibrium* if all candidates choose the same platform $x^* \in X$—$x_c^* = x^*$ for all $c \in C$—and no candidate can gain vote share by deviating: for every $c \in C$ and every $x_c \in X$,

$$\mathcal{V}_c\big(x^*,\, (x_{c'}^*)_{c'\neq c}\big) \;\geq\; \mathcal{V}_c\big(x_c,\, (x_{c'}^*)_{c'\neq c}\big).$$

The symmetric equilibrium is a Nash equilibrium in which every candidate's best response to a field of identical opponents is to broadcast the same platform. Whether this equilibrium admits a classical median-voter interpretation—and under what conditions—is the question the remainder of this section addresses.

4.5 Policy as organizational scheme.

In the classical Downsian model, X is a totally ordered policy space given exogenously; voters and candidates locate on it, and the competition is about who is closest to the median. The semantic framework inverts this picture. The primitive is the semantic election—the pattern of meanings, credibility structures, and trigger behavior. A total order $\leq_X$ on X is a *valid organizational scheme* for an election if and only if it correctly summarizes the coalition structure of the electorate: if the electorate's majority-preference tournament, when arranged along $\leq_X$, exhibits the mountain shape described below. Policy position is not the space within which competition takes place; it is a derived label we assign to the semantic landscape when that landscape is organized enough to support a single dimension.

For each platform $x \in X$, define the *majority-wins set*

$$W(x) \;=\; \{x' \in X : x \text{ defeats } x' \text{ in pairwise majority voting}\},$$

where x defeats x' if a strict majority of voters have $w_v(x, \cdot) > w_v(x', \cdot)$.

Finite illustration. Take $X = \{1, 2, 3, 4, 5\}$ with five voters whose ideal understandings correspond to interior positions ($u_v^* \in \phi(\{2, 3, 4\})$). Positions 1 and 5 are beaten by everything: no voter's ideal lies at or beyond the extremes, so every voter prefers any interior platform to an extreme one. Hence $W(1) = W(5) = \emptyset$. Moving inward, $W(2)$ is nonempty but smaller than $W(3)$; $W(3)$—the median—is largest; $W(4)$ shrinks back; $W(5) = \emptyset$ again. The mountain shape of $|W(\cdot)|$ is the behavioral signature of a recoverable ordering. The labels $1, \ldots, 5$ on platforms are meaningful precisely because they track this coalition structure—not the other way around.

Trivial direction. If X is already totally ordered and the broadcast map ϕ embeds X monotonically into a chain in $\mathscr{U}_{\mathcal{R}}$, and the Downsification conditions of section 4.7 hold on $\phi(X)$, then the welfare formula reduces to classical spatial proximity and the median voter result obtains immediately. The semantic model contains the classical one as the special case in which the organizational scheme happens to be given in advance.

Heterogeneous credibility. When credibility filters differ across voters—$\alpha_{v_1c} \neq \alpha_{v_2c}$—the same platform $\phi(x_c)$ generates different post-campaign understandings across voter types, as if different voters were receiving broadcasts from different positions on different scales. Voter v_1 admitting only policy content from $\phi(x_c)$ locates the candidate at a policy position; voter v_2 admitting only affective content locates the same candidate at an affective register. Downsification requires a single $\leq_X$ that organizes *all* voters' responses simultaneously. With sufficiently incompatible filter types, no such ordering exists. Since credibility heterogeneity is the generic empirical condition (cf. Lodge and Taber, 2013), this alone makes Downsification exceptional.

4.6 The coalition mountain condition.

Definition 4.7. The *coalition mountain condition* (CMC) holds at $x^* \in X$ if there exists a total order $\leq_X$ on X such that the majority-wins function $x \mapsto |W(x)|$ is single-peaked on $\leq_X$ with maximum at x^*, and $|W(x^*)| = |X| - 1$ (i.e., x^* defeats every other platform in pairwise majority voting).

Theorem 4.8. *A symmetric equilibrium x^* is* Downsifiable*—admits a total-order interpretation making it the majority median—if and only if the CMC holds at x^*. When the CMC holds, the recovered ordering $\leq_X$ is unique up to reversal, and x^* is the majority median of the voter ideal distribution projected onto $\leq_X$.*

Proof. Forward (CMC ⇒ Downsifiable). By the CMC, $|W(x^*)| = |X| - 1$: x^* defeats every alternative in pairwise majority voting. The single-peaked structure of $|W(\cdot)|$ on $\leq_X$—increasing to x^*, decreasing thereafter—identifies $\leq_X$ as the ordering under which majority support is organized around x^*; a Condorcet winner on such a tournament is the majority median by construction. The recovered ordering $\leq_X$ is unique up to reversal: any permutation of X on which $|W(\cdot)|$ is strictly single-peaked with maximum at x^* must arrange alternatives in non-decreasing then non-increasing order of majority support, and reversing the order preserves this property.

Backward (Downsifiable ⇒ CMC). Suppose x^* is the majority median under $\leq_X$ with single-peaked voter welfare functions. Then $W(x^*) = X \setminus \{x^*\}$ (the median beats everything). For any $x <_X x^*$, x loses to x^* and to everything between x and x^*, so $|W(x)| < |W(x^*)|$; similarly for $x >_X x^*$. Under single-peakedness, $|W(\cdot)|$ is non-decreasing up to x^* and non-increasing thereafter. Extremes get $|W| = 0$ when no voter's ideal lies at or beyond them (the finite case with interior ideals). Hence $|W(\cdot)|$ is single-peaked at x^* and the CMC holds. □

The recovered $\leq_X$ is not an exogenous input; it is a property of the coalition structure, readable off the behavioral data of the election. Two elections with the same abstract X may admit different valid orderings or none at all, depending entirely on their semantic parameters. When the CMC holds, $\leq_X$ *is* the policy dimension; when it fails, no single dimension organizes the competition, and the semantic lattice provides a richer but less compressed account—the subject of section 5.[6]

[6] Nehring and Puppe (2007) characterize strategy-proof social choice on generalized single-peaked domains and show that the strongest possibility results obtain exactly on

4.7 Downsification conditions.

The CMC is a behavioral criterion. The Downsification conditions are the semantic-parameter conditions sufficient to ensure it. Together they characterize when the organizational scheme works: when the pattern of meanings, credibility structures, and trigger behavior is structured enough that a total order on X faithfully summarizes the coalition landscape.

A semantic election satisfies the *Downsification conditions* if:

1. *Policy dominance*: $\alpha_{vc}([r]) \subseteq M_{\text{policy}}$ for all $v, c, [r]$.

2. *Single dimensionality*: there exists a total order on M_{policy} and a well-defined ideal-policy projection $\pi : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$.

3. *Single-peaked welfare*: $\mathcal{W}(\hat{u}_v^c \wedge u_v^*)$ is single-peaked in $\pi(\phi(x_c))$ for every voter v, with peak at $\pi(u_v^*)$.

4. *Policy campaign*: $\alpha_{vc}([r]) = M_{\text{policy}}$ and epistemic expansion is maximal ($e'_v = e_v \vee e_c$) for all v, c.

5. *Trigger-free*: $\tau_v = \infty$ for all v.

Under these conditions, the post-campaign understanding $\hat{u}_v^c$ is determined entirely by the policy content of $\phi(x_c)$, projected to the single-dimensional policy space. The vote function reduces to proximity:

$$c^*\big(v,\, (x_c)_c\big) \;\in\; \arg\min_{c \in C} \big|\pi(\phi(x_c)) - \pi(u_v^*)\big|,$$

recovering the standard spatial model and establishing the CMC. In the two-candidate case, the unique symmetric equilibrium has both candidates broadcasting at the semantic median: the platform $x^{\text{med}} \in X$ satisfying $\pi(\phi(x^{\text{med}})) = p^*_{\text{med}}$ where p^*_{med} is the median of $\{\pi(u_v^*)\}_{v \in V}$.

Rarity of Downsification. The Downsification conditions are individually demanding and jointly restrictive. Policy dominance fails whenever voters' credibility filters admit affective or social meanings—the warm glow case

---

median spaces—property spaces in which every triple of alternatives has a "between" element. Median spaces are the structural analog of the CMC: when the CMC holds, the recovered ordering $\leq_X$ generates precisely the betweenness geometry that makes X a median space; when the CMC fails, no such geometry is recoverable. The frameworks are orthogonal in orientation—Nehring and Puppe ask what social choice functions are strategy-proof on a fixed domain; the present framework asks when a domain with single-peaked structure exists at all—but the structural duality is exact.

of section 3.5: campaign affect shifts voter understandings independently of policy content, and candidates compete on emotional register rather than policy position. Policy campaign fails whenever candidate epistemic stances are incomparable to voter stances (section 2.2): coarsening across incommensurable schemas loses information, and voters absorb distorted or vacuous policy signals. Single dimensionality fails whenever the contested domain is genuinely multidimensional, which is the generic case. Trigger-free fails whenever voters' coherence thresholds are finite: sufficiently challenging platforms trigger reversion rather than absorption, and voter understandings may end up strictly below where they started (Proposition 3.4).

The demanding character of these conditions is not a formal artifact. Converse (1964) demonstrated empirically that the mass public lacks the ideological constraint that spatial models presuppose: most citizens do not hold coherent, constrained sets of policy positions that map onto a single ordering. The present framework shows precisely why: satisfying the Downsification conditions would require policy-dominant credibility, single-dimensional semantics, and trigger-free updating—conditions that jointly describe a mode of political reasoning that is empirically rare. Downs's model is correct on its own assumptions; the present framework characterizes those assumptions as a knife-edge.

Warm glow domination. When credibility filters are restricted to affective and social meanings, the electoral game is fought entirely on non-policy terrain. The semantic model predicts what empirical work on campaign effects has long documented: that much of what moves voters is the affective register and group-identity signals of political communication, not its policy content (cf. Zaller, 1992; Lodge and Taber, 2013). In the formal terms of section 3.5, the warm glow case is the generic equilibrium of a campaign in which neither candidate's epistemic stance is sufficiently close to voter stances for policy content to pass through credibility filters.

Epistemic incomparability. When the voter population is epistemically heterogeneous—some voters reasoning through one schema, others through an incomparable one—no single understanding can serve as an effective common broadcast. The Strict Father and Nurturant Parent moral worldviews identified by Lakoff (1996) are an instance of precisely this incomparability: both conservative and progressive citizens deploy the same political vocabulary, but the meanings assigned to shared terms—"freedom," "responsibility," "government"—are drawn from categorically different conceptual systems, making their epistemic stances incomparable in $\preceq_E$. Candidates face a dilemma: a broadcast optimized for one epistemic constituency is incoherent or invisible to another. The optimal broadcast is a compromise that is heard by many but heard differently by each—a structural source of message ambiguity that is strategically rational, not a communication failure.

When the Downsification conditions fail generically, the coalition mountain collapses and the election is not Downsifiable; the structure of what remains is the subject of section 5.

5 Semantic equilibrium beyond Downsification.

Section 4 characterized when a symmetric semantic equilibrium is Downsifiable—when the coalition structure admits a one-dimensional organizational scheme. The generic case is that it does not. This section asks what can be said about electoral and legislative competition when Downsification fails.

The section proceeds in two parts. The first establishes that even without Downsification, the mutual updating process among election participants always reaches a stable point: the communication equilibrium exists by Tarski's theorem and is independent of whether the coalition mountain condition holds. The second examines competition itself when the CMC fails. Failure takes two forms—acyclicity without single-peakedness, and full McKelvey cycling—and the semantic framework says different things about each. The section closes with a remark on the legislative extension.

5.1 Communication equilibrium.

Fix a semantic election $\mathcal{E}$ and suppose all $N = |C| + |V|$ participants—candidates and voters alike—are engaged in a mutual updating process. Restrict to the case in which all participants share a common epistemic level $e \in E_{\mathcal{R}}$, and assume the trigger is deactivated ($\tau_i = \infty$ for all i) and epistemic expansion is minimal ($e'_i = e$ for all i). Under these restrictions, the update operator reduces to semantic absorption: agent i adds to their current meanings whatever they are credibly open to absorbing from the others.

For a profile $\boldsymbol{\mu} = (\mu_i)_{i=1}^N \in M_e^N$ and an agent i, define the *available content* at class $[r]$ as

$$A_i(\boldsymbol{\mu})([r]) \;=\; \alpha_i([r]) \,\cap\, \bigcup_{j \neq i} \mu_j([r]),$$

where $\alpha_i([r]) \subseteq M$ is agent i's credibility filter at $[r]$ (Definition 3.2, coarsened to e). The *mutual update map* $\mathbf{T} : M_e^N \to M_e^N$ is then

$$\mathbf{T}(\boldsymbol{\mu})_i([r]) \ = \ \mu_i([r]) \ \cup \ A_i(\boldsymbol{\mu})([r]).$$

Each agent adds to their current meanings whatever is available from the pool of others and passes their filter. A fixed point $\boldsymbol{\mu}^* = \mathbf{T}(\boldsymbol{\mu}^*)$ is a *semantic communication equilibrium*: a profile from which no agent can absorb further meaning from any other.

Proposition 5.1. $\mathbf{T}$ *is a monotone self-map on* $(M_e^N, \preceq_M^N)$.

The proof is immediate: both $\mu_i([r])$ and $A_i(\boldsymbol{\mu})([r])$ are monotone in $\boldsymbol{\mu}$ under pointwise inclusion, so their union is.

Theorem 5.2. *There exists a greatest communication equilibrium* $\boldsymbol{\mu}^*$ *and a least communication equilibrium* $\boldsymbol{\mu}^\dagger$ *in* M_e^N. *Every communication equilibrium lies between them, and the set of all communication equilibria is itself a complete lattice.*

Proof. $(M_e^N, \preceq_M^N)$ is a complete lattice: $M_e \cong (2^M)^{\mathcal{R}_e/\sim_e}$ is complete under pointwise inclusion, and M_e^N is a product of complete lattices. $\mathbf{T}$ is monotone by Proposition 5.1. Tarski's fixed-point theorem gives a greatest fixed point $\boldsymbol{\mu}^*$, a least fixed point $\boldsymbol{\mu}^\dagger$, and the full fixed-point set is a complete lattice. □

The existence argument is structural rather than surprising: once the domain is recognized as a complete lattice and the map as monotone, Tarski's theorem applies immediately. The substance lies in what the fixed-point set looks like and how active triggers alter the picture.

The greatest communication equilibrium $\boldsymbol{\mu}^*$ is the ceiling of mutual comprehension under the given credibility structure—the semantically richest stable profile. The least fixed point $\boldsymbol{\mu}^\dagger$ is the floor: no stable profile lies below it. The spread of the fixed-point lattice measures the semantic range of stable outcomes consistent with the given credibility structure.

Remark 5.3 (Fixed-point spread and credibility convergence). The fixed-point set collapses to a singleton when the credibility structure is *convergent*: there exists an agent i^* with $\alpha_{i^*}([r]) = M$ for all $[r]$ (full receptivity), and the other credibility filters are nested. In this case iterating $\mathbf{T}$ from any starting profile reaches $\boldsymbol{\mu}^*$ in finitely many steps—the fully receptive agent pulls the communication process to its semantic ceiling. At the other extreme, when $\alpha_i([r]) \cap \alpha_j([r]) = \emptyset$ for all $i \neq j$, no agent can absorb anything from any other; $\boldsymbol{\mu}^* = \boldsymbol{\mu}^\dagger$; and every profile is already a fixed point. The spread thus measures how much semantic potential the credibility structure permits to be realized.

Active triggers and persistent polarization. The result holds under deactivated triggers ($\tau_i = \infty$). When triggers are active, the mutual update map is no longer monotone: a richer profile in $\boldsymbol{\mu}$ can increase conflict exposure for some agents, firing their triggers and sending them to a strictly lower baseline rather than forward. The trigger introduces a non-monotone discontinuity that breaks the Tarski argument.

Remark 5.4 (Active-trigger limit cycles). Active triggers can produce mutual impoverishment rather than convergence. Schematically: suppose agents i and j each have a finite threshold $\tau < \infty$. If $|\mathsf{Conf}(u_i, u_j)| > \tau$, i's trigger fires and i reverts to $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$. The reduced u_i now presents a different profile: if $|\mathsf{Conf}(u_j, \beta(u_i, u_j))| > \tau$, j's trigger fires and j reverts to $\beta(u_j, \beta(u_i, u_j)) \prec_{\mathscr{U}} u_j$. If the cycle repeats, both agents drift strictly downward in $\mathscr{U}_{\mathcal{R}}$: each encounter leaves both representationally poorer than before. This is active mutual impoverishment—not merely failure to converge but ongoing deterioration of both parties' representational resources through sustained contact.

5.2 Non-Downsifiable elections: semantic dimension and McKelvey.

The failure of the CMC takes two distinct forms with different electoral implications.

The first is *acyclicity without single-peakedness*: the majority tournament is transitive and a Condorcet winner exists, but voter welfare functions are not single-peaked on any total order—so the winner admits no median interpretation. Elections dominated by warm glow or epistemic incomparability (section 4.7) are of this type: stable symmetric equilibria may exist but carry no recoverable policy dimension. The equilibrium is real; the policy label is not.

The second form is *cycling*: the majority tournament is intransitive, no Condorcet winner exists, and no stable symmetric equilibrium obtains. The operative notion of dimensionality for understanding when this occurs is not the dimension of any exogenous policy space but an intrinsic property of the voter ideal distribution.

Definition 5.5. The *semantic dimension* of a voter population V is the width of the poset $\big(\{u_v^*\}_{v \in V}, \preceq_{\mathscr{U}}\big)$—the cardinality of the largest set of pairwise $\preceq_{\mathscr{U}}$-incomparable ideal understandings.

Semantic dimension 1 means all voter ideals are comparable: they lie on a chain in $\preceq_{\mathscr{U}}$, and Downsification is possible. Semantic dimension exceeding

1 means some voter ideals are incomparable: no single representational direction covers all of them, and the election is genuinely multidimensional in the lattice-theoretic sense.

Definition 5.6. A broadcast map $\phi : X \to \mathscr{U}_{\mathcal{R}}$ is *voter-distinguishing* if for every pair of voters v_1, v_2 with $u^*_{v_1} \not\preceq_{\mathscr{U}} u^*_{v_2}$ and $u^*_{v_2} \not\preceq_{\mathscr{U}} u^*_{v_1}$, there exist $x, y \in X$ such that

$$\mathcal{W}(\phi(x) \wedge u^*_{v_1}) > \mathcal{W}(\phi(x) \wedge u^*_{v_2}),$$
$$\mathcal{W}(\phi(y) \wedge u^*_{v_2}) > \mathcal{W}(\phi(y) \wedge u^*_{v_1}).$$

A voter-distinguishing broadcast map has enough range to place the two voter types on opposite sides of a majority-preference divide: for each incomparable pair, there exist platforms each voter strictly prefers to the other's preferred platform. This is the minimal condition for the majority tournament to exhibit the cycling structure McKelvey's theorem requires.

Proposition 5.7. *If the semantic dimension of V exceeds 1 and the broadcast map $\phi : X \to \mathscr{U}_{\mathcal{R}}$ is voter-distinguishing, then the majority tournament on X generically has no Condorcet winner. When voter welfare functions in the policy projection of ϕ are Euclidean, the majority orbit is path-connected: for any two platforms $x, y \in X$ there exists a finite majority-preference path $x = z_0, z_1, \ldots, z_k = y$ with z_{i+1} defeating z_i in pairwise majority voting for each i.*

Proof. When semantic dimension exceeds 1, there exist voters v_1, v_2 with $u^*_{v_1} \not\preceq_{\mathscr{U}} u^*_{v_2}$ and $u^*_{v_2} \not\preceq_{\mathscr{U}} u^*_{v_1}$. Their welfare functions $x \mapsto \mathcal{W}(\phi(x) \wedge u^*_{v_i})$ are driven by incompatible semantic content: maximizing common ground with $u^*_{v_1}$ requires meanings absent from $u^*_{v_2}$, and conversely. The voter-distinguishing condition ensures that platforms generating these competing preference rankings exist in X: there are platforms the two voter types strictly rank in opposite orders, so the induced majority tournament has no stable Condorcet winner. The voter-distinguishing condition establishes that the induced majority tournament has the competitive structure of a multidimensional preference profile: cycling exists. When voter welfare functions in the policy projection are Euclidean, McKelvey's theorem (McKelvey, 1976) delivers the stronger conclusion: the orbit is path-connected, and any platform can be reached from any other by a majority path. □

Proposition 5.8. *When the semantic dimension of* V *exceeds* 1*, the welfare-relevant McKelvey majority-preference orbit is confined to the* semantic interval

$$\Big[\bigwedge_{v\in V} u_v^*,\ \ \bigvee_{v\in V} u_v^*\Big]$$

in $\mathscr{U}_{\mathcal{R}}$.

Proof. Any platform $\phi(x)$ strictly above $\bigvee_v u_v^*$ gives every voter the same welfare—each voter's common ground with their ideal is u_v^* itself, the maximum achievable—so all such platforms are welfare-indistinguishable and generate no majority preferences between them. Any platform strictly below $\bigwedge_v u_v^*$ contributes nothing beyond what is already present in every voter's initial understanding: it adds no semantic common ground with any voter's ideal. Neither class participates in the majority tournament; the welfare-relevant orbit is therefore contained in the semantic interval. □

The semantic interval $[\bigwedge_v u_v^*, \bigvee_v u_v^*]$ is an intrinsic property of the voter ideal distribution, defined without reference to any exogenous policy space. Classical spatial theory has no structural bound on the chaos orbit; the semantic lattice imposes one automatically.

Strategic polysemy. Partition the voter population into *credibility types*: $V = T_1 \cup \cdots \cup T_k$ where voters within type T_i share credibility filter α_i. Filters are *pairwise disjoint* if $\alpha_i([r]) \cap \alpha_j([r]) = \emptyset$ for all $i \neq j$ and all $[r] \in \mathcal{R}_e/{\sim_e}$. A broadcast $\phi(x)$ is *polysemic* with respect to this partition if its meaning assignment contains content in multiple filter domains: each voter type absorbs a component of the broadcast that is disjoint from what any other type absorbs.

Proposition 5.9. *When credibility filters are pairwise disjoint, the vote-share optimization for candidate* c *decomposes by type: the component* $\alpha_i([r]) \cap \mu_c([r])$ *of* $\phi(x_c)$ *absorbed by type* T_i *can be chosen to maximize* T_i*'s welfare independently of what the broadcast does for any other type.*

Proof. With disjoint filters, $\hat{u}_v^c = T_{\alpha_{vc},\tau_v}(u_v^0, \phi(x_c))$ for $v \in T_i$ depends only on the component $\alpha_i \cap \phi(x_c)$. Since $\alpha_i([r]) \cap \alpha_j([r]) = \emptyset$ for $i \neq j$, modifying the T_i-component of $\phi(x_c)$ leaves $\hat{u}_v^c$ unchanged for all $v \notin T_i$. Vote share from T_i is therefore a function of $\alpha_i \cap \phi(x_c)$ alone, and the optimization over $\phi(x_c)$ separates across types. □

Corollary 5.10. *In a two-candidate race with pairwise disjoint voter types, the equilibrium broadcast is polysemic: the winning platform simultaneously addresses distinct voter constituencies through separate semantic channels, with each channel optimized for its target type. Message ambiguity is strategically rational—not a failure of communication—and is an equilibrium property of any election in which credibility filters are sufficiently heterogeneous.*

This is the strategic content of the epistemic incomparability result of section 4.7: the candidate's optimal response to an epistemically fragmented electorate is a broadcast heard differently by each fragment, not one that tries to speak a common language. The mechanism differs from the context-dependent voting model of Callander and Wilson (2008), in which ambiguity is a lottery over positions in a fixed one-dimensional policy space and voters develop a preference for uncertainty through relative evaluation; here polysemy is a structural decomposition of a single broadcast arising from disjoint credibility filters, with no policy space assumed.

Majoritarian convergence from antichain geometry. Consider an election in which voter ideals form a 2-element antichain $\{u_A^*, u_B^*\} \subset \mathscr{U}_{\mathcal{R}}$—two incomparable ideal understandings. This partitions the electorate into blocs $V_A = \{v : u_v^* \preceq_{\mathscr{U}} u_A^*\}$ and $V_B = \{v : u_v^* \preceq_{\mathscr{U}} u_B^*\}$ of sizes n_A and n_B. Since u_A^* and u_B^* are incomparable, $u_A^* \wedge u_B^* \prec_{\mathscr{U}} u_A^*$ and $u_A^* \wedge u_B^* \prec_{\mathscr{U}} u_B^*$: the ideals are genuinely distinct in semantic content, with neither bloc's ideal subsumed in the other's.

Proposition 5.11. *Suppose voter ideals form a 2-element antichain, $n_A > |V|/2$, and there exists $y \in X$ with $\phi(y) \succeq_{\mathscr{U}} u_A^*$. Then every symmetric equilibrium x^* satisfies $\phi(x^*) \succeq_{\mathscr{U}} u_A^*$: any symmetric platform with $\phi(x^*) \not\succeq_{\mathscr{U}} u_A^*$ is unstable.*

Proof. Let $\phi(x^*)$ be any symmetric platform with $\phi(x^*) \not\succeq_{\mathscr{U}} u_A^*$. Then $\phi(x^*) \wedge u_A^* \prec_{\mathscr{U}} u_A^*$ and $\mathcal{W}(\phi(x^*) \wedge u_A^*) < \mathcal{W}(u_A^*)$ for all $v \in V_A$. By hypothesis there exists $y \in X$ with $\phi(y) \succeq_{\mathscr{U}} u_A^*$; the deviation to y gives every $v \in V_A$ welfare $\mathcal{W}(\phi(y) \wedge u_A^*) = \mathcal{W}(u_A^*)$ (the maximum achievable). Since $|V_A| > |V|/2$, the deviating candidate wins a strict majority, so x^* is not a Nash equilibrium. □

Corollary 5.12. *When voter ideals form a 2-element antichain and $n_A > |V|/2$, every symmetric equilibrium satisfies $\phi(x^*) \succeq_{\mathscr{U}} u_A^*$ (Proposition 5.11): competition forces both candidates toward the majority bloc's ideal. Since u_A^* and u_B^* are incomparable, any symmetric equilibrium with $\phi(x^*) \not\succeq_{\mathscr{U}} u_A^* \vee u_B^*$*

satisfies $\phi(x^*) \not\succeq_{\mathscr{U}} u_B^*$, giving bloc B strictly less than their ideal welfare: $\mathcal{W}(\phi(x^*) \wedge u_B^*) < \mathcal{W}(u_B^*)$. The degree of this exclusion is determined by the incomparability of u_A^* and u_B^*: the further $u_A^* \wedge u_B^*$ lies below u_B^* in $\preceq_{\mathscr{U}}$, the greater the welfare deficit of the minority bloc at any symmetric equilibrium.

This gives a lattice-based account of majoritarian convergence and minority exclusion as equilibrium outcomes, without any exogenous spatial model. The systematic exclusion of the minority bloc is not a spatial artifact but a structural consequence of semantic incomparability in the voter ideal distribution: when the majority bloc's ideal is not subsumed in the minority's, competition selects for the majority ideal and the minority welfare shortfall is determined by how far the two bloc ideals diverge in $\preceq_{\mathscr{U}}$.[7]

The meet as semantic fallback. Even in full McKelvey chaos, $\bigwedge_{v \in V} u_v^*$ is a specific, well-defined element of $\mathscr{U}_{\mathcal{R}}$: it is the most representationally conservative platform compatible with every voter's ideal as a floor—the understanding whose content is contained in every voter's ideal. It is the platform that offends no one while fully satisfying no one. Where classical spatial theory has no natural focal point in the absence of a Condorcet winner, the complete lattice structure of $\mathscr{U}_{\mathcal{R}}$ supplies one: the semantic meet of voter ideals is always well-defined and always represents a minimal common semantic ground.

Remark 5.13 (Legislative extension)*.* The framework extends to multi-actor legislative bargaining in the natural way. Call a coalition $S \subseteq [N]$ *semantically viable* if $\bigwedge_{i \in S} u_i \succeq_{\mathscr{U}} u_{\min}$ for some threshold $u_{\min}$—its members share enough representational common ground for joint action. The *semantic core* is the set of understanding profiles from which no majority coalition has a semantically viable, welfare-improving deviation. By the same argument as theorem 4.8 and Proposition 5.7, the semantic core is non-empty when semantic dimension does not exceed 1 and generically empty when it does. The fixed-point spread of theorem 5.2 measures the range of communicatively stable profiles before strategic deviation: when credibility is convergent the legislature settles near a unique communication equilibrium and the core contains it; when

[7] Existing formal accounts of partisan polarization typically invoke either spatial forces (primary constraints, policy-motivated candidates) or cognitive bias: Orteleva and Snowberg (2015) derive ideological extremeness from *correlational neglect*, in which citizens underestimate how correlated their experiences are and thereby become overconfident about their beliefs. The present account requires neither: minority exclusion is a structural consequence of semantic incomparability in the voter ideal distribution, derived without a policy space and without any assumption of belief miscalibration.

credibility is divergent the legislature is simultaneously informationally at rest and politically volatile.

6 Empirical implications.

The three main results are structural claims, but each generates a prediction that is empirically distinguishable from what the spatial model would say. This section makes those predictions explicit. Each follows directly from the formal results and inverts at least one prediction that spatial theory makes.

The sophistication-proximity inversion. Spatial theory predicts that politically sophisticated voters—those with richly constrained, ideologically coherent belief systems—should be the most likely to vote on the basis of spatial proximity. The credibility structure and trigger mechanism generate the opposite prediction.

Sophisticated voters occupy higher epistemic stances and possess denser networks of semantically interconnected meanings. When a candidate broadcasts content that conflicts with this prior structure, the conflict set $|\mathsf{Conf}(u_v, \hat{u}_v^c)|$ is larger precisely because there is more prior content that can conflict. Under a finite trigger threshold τ_v, the trigger fires more readily for the voter with the richest prior. The post-encounter understanding is driven toward the baseline $\beta(u_v, \hat{u}_v^c) \prec_{\mathscr{U}} u_v$: the sophisticated voter ends up semantically poorer and more entrenched, not updated.

The testable prediction is a negative interaction: among voters with strong ideological priors, the fit of the spatial proximity model should *decrease* with political sophistication. The most ideologically engaged citizens should be the *least* well-described by proximity voting. This is consistent with the experimental findings of Taber and Lodge (2006) and Lodge and Taber (2013)—that political sophistication predicts motivated skepticism rather than evenhanded updating—but the semantic framework supplies a structural account of why: sophistication increases the size of the conflict surface, which makes the trigger fire more readily and the baseline reversion deeper. The prediction is not that sophisticated voters fail to vote; it is that their voting behavior is less accurately modeled by proximity to a policy ideal and more accurately modeled by resistance to update.

Credibility heterogeneity and polysemic campaigns. Spatial theory predicts that candidates communicate coherent policy positions, with strategic ambiguity arising only under specific conditions—multidimensionality, preference for flexibility, or the kind of context-dependence modeled by Callander and Wilson (2008). The polysemy result generates a structurally different prediction: as credibility heterogeneity increases across the electorate, the optimal campaign message becomes less semantically unified as an equilibrium property, not as a failure.

The mechanism is specific. When voter types have pairwise disjoint credibility filters, the candidate's vote-share optimization separates by type (Proposition 5.9): each component of the broadcast is independently optimized for the audience that can absorb it, and the aggregate message is polysemic—simultaneously coherent to each type along its own semantic channel and incoherent as a single policy statement. This is structurally different from a lottery over positions: polysemy is not about offering different voters different expected positions but about a single broadcast being genuinely heard differently, with each hearing optimized.

The empirical prediction: as the credibility structure of the electorate becomes more fragmented—as different voter segments are less likely to absorb the same semantic content—campaign messaging should become less semantically unified at the aggregate level while becoming more precisely targeted at the type level. The rise of disaggregated campaign communication over the past two decades is consistent with this prediction, but the mechanism distinguishes it from simpler explanations based on voter heterogeneity in policy positions. The polysemy prediction is specifically about the *semantic decomposition* of a broadcast: a polysemic campaign message is not ambiguous because the candidate is uncertain about voter preferences or hiding a position; it is polysemic because the electorate's credibility structure makes a unified message suboptimal.

Semantic incomparability and the structure of polarization. The spatial account of ideological polarization is a story about distance: voters who are far apart on a shared dimension pull candidates toward the extremes. The majoritarian convergence result (Corollary 5.12) identifies a structural mechanism of minority exclusion entirely distinct from spatial distance: when voter ideal understandings are lattice-theoretically incomparable—when no single representational direction covers both blocs—and one bloc constitutes a strict majority, competition forces equilibrium toward the majority bloc's ideal, regardless of the spatial distance between voter ideal points.

This generates a prediction about the *type* of polarization, not just its level. Two electorates with identical spatial dispersion but different lattice geometry should exhibit different degrees of minority bloc exclusion at equilibrium. An electorate in which voters disagree about their positions on a shared policy dimension is different from one in which voters hold beliefs organized around incommensurable conceptual frameworks—not merely far apart but genuinely unable to place each other's positions on a common scale. The former can, in principle, be Downsified; the latter cannot.

The relevant measure is not the distance between voter positions but the *constraint* of voter belief systems in the sense of Converse (1964): the degree to which different policy attitudes cohere within a single conceptual system. Low between-voter constraint—voters' belief systems organized around different, non-overlapping dimensions—corresponds to the incomparability structure that generates the majority-convergence result. High constraint within each bloc, combined with low constraint across blocs, is the distributional profile that produces the 2-element antichain in the voter ideal space.

The prediction distinguishes the semantic account from behavioral accounts of polarization based on overconfidence (Orteleva and Snowberg, 2015): those accounts predict that polarization tracks overconfidence as an individual characteristic; the semantic account predicts that minority exclusion tracks the incomparability of the voter ideal distribution as a structural property of the electorate, independently of any belief miscalibration. Both predictions can be empirically separated: overconfidence-based polarization should be reducible by improving citizens' epistemic calibration; incomparability-based minority exclusion is structural and persists regardless of calibration.

7 Conclusion.

The three main results of this paper — the semantic chaos bound, the polysemy decomposition, and the lattice account of equilibrium polarization — share a common structure: each is a statement about what the lattice geometry of the voter ideal distribution implies for competition, without reference to any exogenous policy space. The semantic chaos bound says that even when majority cycling is unavoidable, it is confined to the interval between the meet and join of voter ideals — a bound the lattice imposes automatically. The polysemy result says that when credibility filters are sufficiently heterogeneous, the optimal campaign is not a coherent policy message but a structured ambiguity whose components are independently optimized for disjoint audiences. The polarization result says that when voter ideals are incomparably structured and one bloc is a strict majority, competition forces equilibrium toward the majority bloc's ideal; the minority bloc's welfare shortfall is determined by the incomparability structure of voter ideals, not by the geometry of a pre-given policy axis.

The deeper point these results share is the paper's central conceptual inversion. In the Downsian model, the policy space is given and competition is the process of locating on it. In the semantic framework, the policy space is derived: it exists when and only when the coalition structure of the electorate admits a one-dimensional organizational scheme, and the dimension — when recoverable — is readable off the majority-preference tournament rather than imposed from outside. When the coalition mountain condition holds, the recovered ordering *is* the policy dimension. When it does not, there is no policy dimension to recover, and the semantic lattice provides a richer but less compressed account of what competition is about.

Several extensions remain open. Dynamic campaigns — in which candidates' broadcasts alter the semantic landscape across multiple rounds of communication, shifting voter credibility structures and epistemic stances before the next cycle — introduce path dependence that the static framework does not capture. Endogenous credibility — in which candidates invest in signal quality or schema fluency with specific voter types, actively shaping their own α_{vc} parameters — turns credibility from a fixed parameter into a strategic variable. Endogenous field evolution — in which the representational field $\mathcal{R}$ itself changes as an outcome of political and technological processes — is the deepest open extension.[8] And the question of measurement presses: the semantic dimension of a voter population is defined as the width of a poset of ideal understandings, but the empirical tools of belief systems research — survey measures of ideological constraint, cross-cutting issue positions, schema complexity — suggest natural operationalizations that could connect the formal object to data.

[8]Two mechanisms press in this direction. First, algorithmic curation reshapes the available distinctions: as information environments fragment, the equivalence relations $\sim_e$ diverge across voter populations, generating the disjoint credibility filter structure that the polysemy result requires as a condition rather than an assumption. Second, political entrepreneurs invest not just in credibility on the existing field but in introducing new representational resources and retiring old ones—constituting rather than merely locating on the semantic landscape. At this level, political power operates through field-making rather than position-taking, and the systematic divergence of available distinctions provides a structural explanation for why the conditions for Downsian competition are harder to satisfy in a fragmented media environment than in the broadcast-media environment Downs described. A formal treatment requires modeling $\mathcal{R}$ as an object within a higher-order understanding lattice. Since the understanding lattice $\mathscr{U}_{\mathcal{R}}$ at level n can represent frameworks at level $< n$ as semantic objects but cannot represent itself without impredicativity, the appropriate infrastructure is a hierarchy of Grothendieck universes: each universe U_n contains $\mathscr{U}_{\mathcal{R}_{n-1}}$ as a set-sized object, and field evolution is modeled as a map between levels of the hierarchy rather than as a process within a fixed level.

Downs's model is correct on its own assumptions. The present framework characterizes those assumptions precisely, shows when they are satisfied and when they are not, and provides a set of qualitatively distinct predictions for the cases where they fail. The spatial model's parsimony is a virtue when the policy space is well-defined; the semantic framework shows that whether it is well-defined is itself an empirical and structural question, one that the coalition geometry of the electorate either answers or refuses to answer.

A Proofs.

Lemma 2.3. $(E_{\mathcal{R}}, \preceq_E)$ *is a partially ordered set.*

Proof. We verify the three properties of a partial order. Recall that $e_1 \preceq_E e_2$ iff $\sim_{e_2}$ refines $\sim_{e_1}$ as partitions of $\mathcal{R}$, i.e., every $\sim_{e_2}$-class is contained in a $\sim_{e_1}$-class, equivalently $\sim_{e_2} \subseteq \sim_{e_1}$ as binary relations.

Reflexivity. Every equivalence class of $\sim_e$ is contained in itself, so $\sim_e$ refines $\sim_e$.

Antisymmetry. Suppose $e_1 \preceq_E e_2$ and $e_2 \preceq_E e_1$: $\sim_{e_2} \subseteq \sim_{e_1}$ and $\sim_{e_1} \subseteq \sim_{e_2}$. Then $\sim_{e_1} = \sim_{e_2}$, so the two partitions are identical and $e_1 = e_2$.

Transitivity. Suppose $e_1 \preceq_E e_2$ and $e_2 \preceq_E e_3$: $\sim_{e_2} \subseteq \sim_{e_1}$ and $\sim_{e_3} \subseteq \sim_{e_2}$. Then $\sim_{e_3} \subseteq \sim_{e_2} \subseteq \sim_{e_1}$, giving $e_1 \preceq_E e_3$. □

Lemma 2.4. $(E_{\mathcal{R}}, \preceq_E)$ *is a complete lattice with bottom element* $e_\perp$ *(empty visible scope:* $\mathcal{R}_{e_\perp} = \emptyset$*,* $I_{e_\perp} = \mathcal{R}$*) and top element* $e_\top = \big(\mathcal{R}, \mathrm{id}_{\mathcal{R}}\big)$ *(full visible scope, each resource its own class).*

Proof. Let $\{e_j\}_{j \in J}$ be a family of epistemic stances, each a partition of $\mathcal{R}$. We exhibit the meet and join and verify the universal properties. Throughout, $e_1 \preceq_E e_2$ iff $\sim_{e_2} \subseteq \sim_{e_1}$ as binary relations on $\mathcal{R}$ (finer partition = richer stance = smaller equivalence relation).

Meet. Set $\sim^\wedge = \mathrm{EqCl}\big(\bigcup_{j \in J} \sim_j\big)$, the smallest equivalence relation on $\mathcal{R}$ containing every $\sim_j$. Write $e^\wedge$ for the corresponding partition; its visible scope is $\bigcap_j \mathcal{R}_{e_j}$.

Lower bound. We need $\sim_{e_j} \subseteq \sim^\wedge$ for all j. This holds by construction: $\sim^\wedge$ is defined to contain every $\sim_j$.

Greatest lower bound. Let e' satisfy $e' \preceq_E e_j$ for all j, i.e., $\sim_j \subseteq \sim_{e'}$ for all j. Since $\sim_{e'}$ is an equivalence relation containing every $\sim_j$, it contains $\mathrm{EqCl}(\bigcup_j \sim_j) = \sim^\wedge$. Hence $\sim^\wedge \subseteq \sim_{e'}$, i.e., $e' \preceq_E e^\wedge$.

Join. Set $\sim^\vee = \bigcap_{j \in J} \sim_j$, the intersection of all $\sim_j$ as binary relations on $\mathcal{R}$. The intersection of equivalence relations on a fixed set is an equivalence relation, so $\sim^\vee$ is well defined. Write $e^\vee$ for the corresponding partition; its visible scope is $\bigcup_j \mathcal{R}_{e_j}$.

Upper bound. We need $\sim^\vee \subseteq \sim_j$ for all j. This holds by definition of intersection.

Least upper bound. Let e' satisfy $e_j \preceq_E e'$ for all j, i.e., $\sim_{e'} \subseteq \sim_j$ for all j. Then $\sim_{e'} \subseteq \bigcap_j \sim_j = \sim^\vee$, i.e., $e^\vee \preceq_E e'$.

Extrema. $e_\perp = (\emptyset, \cdot)$ corresponds to the single-class partition $\{\mathcal{R}\}$ (entire field is the invisible class). This is the coarsest possible partition: $\sim_{e_\perp} = \mathcal{R} \times \mathcal{R} \supseteq \sim_e$ for every e, so $e_\perp \preceq_E e$ for every e. $e_\top = (\mathcal{R}, \mathrm{id}_\mathcal{R})$ corresponds to the discrete partition $\{\{r\} : r \in \mathcal{R}\}$ (empty invisible class, maximal discrimination). The discrete partition refines every partition: $\sim_{e_\top} = \Delta_\mathcal{R} \subseteq \sim_e$ for every e, so $e \preceq_E e_\top$ for every e. □

Theorem 2.7. *$(\mathscr{U}_\mathcal{R}, \preceq_\mathscr{U})$ is a complete lattice. For any family $\{(e_j, \mu_{e_j})\}_{j\in J} \subseteq \mathscr{U}_\mathcal{R}$, writing $e^\wedge = \bigwedge_j e_j$ and $e^\vee = \bigvee_j e_j$, the meet and join are*

$$\bigwedge_{j\in J} (e_j, \mu_{e_j}) = \big(e^\wedge,\ \mu_\wedge\big), \qquad\qquad \bigvee_{j\in J} (e_j, \mu_{e_j}) = \big(e^\vee,\ \mu_\vee\big),$$

where

$$\mu_\wedge(C) = \bigcap_{j\in J} \bigcap_{\substack{D\in\mathcal{R}_{e_j}/\sim_{e_j} \\ D\subseteq C}} \mu_{e_j}(D), \qquad \mu_\vee(D) = \bigcup_{j\in J} \mathrm{Ref}_{e_j\to e^\vee}(\mu_{e_j})(D),$$

for $C \in \mathcal{R}_{e^\wedge}/\sim_{e^\wedge}$ and $D \in \mathcal{R}_{e^\vee}/\sim_{e^\vee}$ respectively.

Proof. Let $\{(e_j, \mu_{e_j})\}_{j\in J} \subseteq \mathscr{U}_\mathcal{R}$. Write $e^\wedge = \bigwedge_j e_j$ and $e^\vee = \bigvee_j e_j$, which exist by Lemma 2.4. Define $\mu_\wedge$ and $\mu_\vee$ as in the theorem statement. Both are well-defined elements of $M_{e^\wedge}$ and $M_{e^\vee}$ respectively, since each value is a subset of M.

The meet is a lower bound. Fix $j \in J$. We have $e^\wedge \preceq_E e_j$ by construction. For the semantic condition: fix $C \in \mathcal{R}_{e^\wedge}/\sim_{e^\wedge}$ and $D \in \mathcal{R}_{e_j}/\sim_{e_j}$ with $D \subseteq C$. Then

$$\mu_\wedge(C) = \bigcap_{k\in J} \bigcap_{D'\subseteq C} \mu_{e_k}(D') \ \subseteq\ \mu_{e_j}(D),$$

since $D \subseteq C$ and $j \in J$, so $\mu_{e_j}(D)$ appears as one of the terms in the intersection. Hence $\mathrm{Ref}_{e^\wedge\to e_j}(\mu_\wedge)(D) = \mu_\wedge(C) \subseteq \mu_{e_j}(D)$, giving $(e^\wedge, \mu_\wedge) \preceq_\mathscr{U} (e_j, \mu_{e_j})$.

The meet is the greatest lower bound. Let (e', μ') be any lower bound, so $e' \preceq_E e_j$ and $\mu'(C') \subseteq \mu_{e_j}(D)$ for all j, all $C' \in \mathcal{R}_{e'}/{\sim_{e'}}$, and all $D \subseteq C'$ with $D \in \mathcal{R}_{e_j}/{\sim_{e_j}}$.

Since $e' \preceq_E e_j$ for all j, we have $e' \preceq_E e^\wedge$. We must show $(e', \mu') \preceq_{\mathscr{U}} (e^\wedge, \mu_\wedge)$, i.e., $\mu'(C') \subseteq \mu_\wedge(C)$ for all $C \subseteq C'$ with $C \in \mathcal{R}_{e^\wedge}/{\sim_{e^\wedge}}$.

Fix such $C' \supseteq C$. For any $j \in J$ and any $D \in \mathcal{R}_{e_j}/{\sim_{e_j}}$ with $D \subseteq C$: since $D \subseteq C \subseteq C'$, the lower bound condition gives $\mu'(C') \subseteq \mu_{e_j}(D)$. Taking the intersection over all j and all $D \subseteq C$:

$$\mu'(C') \;\subseteq\; \bigcap_{j\in J} \; \bigcap_{D\subseteq C} \mu_{e_j}(D) = \mu_\wedge(C).$$

Thus $\mathrm{Ref}_{e'\to e^\wedge}(\mu')(C) = \mu'(C') \subseteq \mu_\wedge(C)$, giving $(e', \mu') \preceq_{\mathscr{U}} (e^\wedge, \mu_\wedge)$.

The join is an upper bound. Fix $j \in J$. We have $e_j \preceq_E e^\vee$ by construction. For any $D' \in \mathcal{R}_{e^\vee}/{\sim_{e^\vee}}$, let C_j be the e_j-class containing D'. Then

$$\mu_{e_j}(C_j) = \mathrm{Ref}_{e_j\to e^\vee}(\mu_{e_j})(D') \;\subseteq\; \bigcup_{k\in J} \mathrm{Ref}_{e_k\to e^\vee}(\mu_{e_k})(D') = \mu_\vee(D').$$

Since $\mathrm{Ref}_{e_j\to e^\vee}(\mu_{e_j})(D') = \mu_{e_j}(C_j) \subseteq \mu_\vee(D')$, we have $(e_j, \mu_{e_j}) \preceq_{\mathscr{U}} (e^\vee, \mu_\vee)$.

The join is the least upper bound. Let (e', μ') be any upper bound, so $e_j \preceq_E e'$ and $\mu_{e_j}(C_j) \subseteq \mu'(E)$ for all j, all $E \in \mathcal{R}_{e'}/{\sim_{e'}}$, and all $C_j \supseteq E$ with $C_j \in \mathcal{R}_{e_j}/{\sim_{e_j}}$.

Since $e_j \preceq_E e'$ for all j, we have $e^\vee \preceq_E e'$. We must show $(e^\vee, \mu_\vee) \preceq_{\mathscr{U}} (e', \mu')$, i.e., $\mu_\vee(C^\vee) \subseteq \mu'(E)$ for all $C^\vee \in \mathcal{R}_{e^\vee}/{\sim_{e^\vee}}$ and $E \in \mathcal{R}_{e'}/{\sim_{e'}}$ with $E \subseteq C^\vee$.

Fix such $C^\vee$ and $E \subseteq C^\vee$. For any $j \in J$, let $C_j^{(C^\vee)}$ be the e_j-class containing $C^\vee$. Since $E \subseteq C^\vee \subseteq C_j^{(C^\vee)}$, the upper bound condition gives $\mu_{e_j}(C_j^{(C^\vee)}) \subseteq \mu'(E)$. Therefore

$$\mu_\vee(C^\vee) = \bigcup_{j\in J} \mathrm{Ref}_{e_j\to e^\vee}(\mu_{e_j})(C^\vee) = \bigcup_{j\in J} \mu_{e_j}(C_j^{(C^\vee)}) \;\subseteq\; \mu'(E).$$

Hence $(e^\vee, \mu_\vee) \preceq_{\mathscr{U}} (e', \mu')$.

Extrema. $u_\perp = (e_\perp, \mu_{e_\perp})$ with $\mu_{e_\perp}(?) = \emptyset$ lies below every understanding: $e_\perp \preceq_E e$ for all e, and the refinement of $\mu_{e_\perp}$ is the zero map, which is $\preceq_M$ any semantics. $u_\top = (e_\top, \mu_{e_\top})$ with $\mu_{e_\top}(\{r\}) = M$ for every singleton class lies above every understanding: $e \preceq_E e_\top$ for all e, and the refinement of any μ_e at $e_\top$'s level assigns $\mu_e(C) \subseteq M = \mu_{e_\top}(\{r\})$ for every singleton $\{r\} \subseteq C$. $\square$

Proposition 3.4. *If* $\mathsf{Conf}(u_i, u_j) \neq \emptyset$*, then* $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$*.*

Proof. Let $[r^*] \in \mathsf{Conf}(u_i, u_j)$. By definition of the conflict set, $\mu_i([r^*]) \cap F_{ij}([r^*]) \neq \emptyset$. Since u_i is an upper bound for S_{ij}, the join $\beta = \bigvee S_{ij}$ satisfies $\beta \preceq_{\mathscr{U}} u_i$. It remains to show $\beta \neq u_i$.

S_{ij} is nonempty. Define $\mu_i^*(C) = \mu_i(C) \setminus F_{ij}(C)$ for $C \in \mathsf{Conf}(u_i, u_j)$ and $\mu_i^*(C) = \mu_i(C)$ otherwise. Then $(e_i, \mu_i^*) \preceq_{\mathscr{U}} u_i$ since $\mu_i^*(C) \subseteq \mu_i(C)$ for all C, and $\mu_i^*([r]) \cap F_{ij}([r]) = \emptyset$ for all $[r] \in \mathsf{Conf}$. So $(e_i, \mu_i^*) \in S_{ij}$.

The join β has epistemic component e_i. Since $(e_i, \mu_i^*) \in S_{ij}$ and all elements of S_{ij} have $e_u \preceq_E e_i$, the epistemic join is e_i. Write $\beta = (e_i, \mu_\beta)$.

$\mu_\beta([r^*])$ avoids $F_{ij}([r^*])$. By the join formula, $\mu_\beta([r^*]) = \bigcup_{u \in S_{ij}} \mathrm{Ref}_{e_u \to e_i}(\mu_u)([r^*])$. Every $u \in S_{ij}$ satisfies $\mathrm{Ref}_{e_u \to e_i}(\mu_u)([r^*]) \cap F_{ij}([r^*]) = \emptyset$, so $\mu_\beta([r^*]) \cap F_{ij}([r^*]) = \emptyset$.

$u_i \not\preceq_{\mathscr{U}} \beta$. If $(e_i, \mu_i) \preceq_{\mathscr{U}} (e_i, \mu_\beta)$, then $\mu_i([r^*]) \subseteq \mu_\beta([r^*])$. But $\mu_\beta([r^*]) \cap F_{ij}([r^*]) = \emptyset$ while $\mu_i([r^*]) \cap F_{ij}([r^*]) \neq \emptyset$, a contradiction. So $u_i \not\preceq_{\mathscr{U}} \beta$, and since $\beta \preceq_{\mathscr{U}} u_i$, antisymmetry gives $\beta \neq u_i$. Therefore $\beta \prec_{\mathscr{U}} u_i$. □

Proposition 5.1. **T** *is a monotone self-map on* $(M_e^N, \preceq_M^N)$.

Proof. Let $\boldsymbol{\mu} \preceq_M^N \boldsymbol{\mu}'$, i.e., $\mu_i([r]) \subseteq \mu_i'([r])$ for all i and all $[r] \in \mathcal{R}_e/{\sim_e}$. Fix i and $[r]$. Since $\mu_j([r]) \subseteq \mu_j'([r])$ for all $j \neq i$,

$$A_i(\boldsymbol{\mu})([r]) = \alpha_i([r]) \cap \bigcup_{j \neq i} \mu_j([r]) \;\subseteq\; \alpha_i([r]) \cap \bigcup_{j \neq i} \mu_j'([r]) = A_i(\boldsymbol{\mu}')([r]).$$

Therefore

$$\mathbf{T}(\boldsymbol{\mu})_i([r]) = \mu_i([r]) \cup A_i(\boldsymbol{\mu})([r]) \;\subseteq\; \mu_i'([r]) \cup A_i(\boldsymbol{\mu}')([r]) = \mathbf{T}(\boldsymbol{\mu}')_i([r]).$$

Hence $\mathbf{T}(\boldsymbol{\mu}) \preceq_M^N \mathbf{T}(\boldsymbol{\mu}')$. □

References

Alchourrón, Carlos E., Peter Gärdenfors and David Makinson. 1985. "On the Logic of Theory Change: Partial Meet Contraction and Revision Functions." *Journal of Symbolic Logic* 50(2):510–530.

Baltag, Alexandru and Bryan Renne. 2016. Dynamic Epistemic Logic. In *The Stanford Encyclopedia of Philosophy*, ed. Edward N. Zalta. Winter 2016 ed. Metaphysics Research Lab, Stanford University.

Black, Duncan. 1948. "On the Rationale of Group Decision-making." *Journal of Political Economy* 56(1):23–34.

Callander, Steven and Catherine H. Wilson. 2008. "Context-Dependent Voting and Political Ambiguity." *Journal of Public Economics* 92(3–4):565–581.

Converse, Philip E. 1964. "The Nature of Belief Systems in Mass Publics." pp. 206–261.

Downs, Anthony. 1957. *An Economic Theory of Democracy.* New York: Harper & Row.

Entman, Robert M. 1993. "Framing: Toward Clarification of a Fractured Paradigm." *Journal of Communication* 43(4):51–58.

Iyengar, Shanto, Yphtach Lelkes, Matthew Levendusky, Neil Malhotra and Sean J. Westwood. 2019. "The Origins and Consequences of Affective Polarization in the United States." *Annual Review of Political Science* 22:129–46.

Lakoff, George. 1996. *Moral Politics: What Conservatives Know That Liberals Don't.* Chicago: University of Chicago Press.

Lodge, Milton and Charles S. Taber. 2013. *The Rationalizing Voter.* Cambridge: Cambridge University Press.

McKelvey, Richard D. 1976. "Intransitivities in Multidimensional Voting Models and Some Implications for Agenda Control." *Journal of Economic Theory* 12(3):472–482.

Nehring, Klaus and Clemens Puppe. 2007. "The Structure of Strategy-Proof Social Choice — Part I: General Characterization and Possibility Results on Median Spaces." *Journal of Economic Theory* 135(1):269–305.

Orteleva, Pietro and Erik Snowberg. 2015. “Overconfidence in Political Behavior.” *American Economic Review* 105(2):504–535.

Poole, Keith T. and Howard Rosenthal. 1985. “A Spatial Model for Legislative Roll Call Analysis.” *American Journal of Political Science* 29(2):357–384.

Taber, Charles S. and Milton Lodge. 2006. “Motivated Skepticism in the Evaluation of Political Beliefs.” *American Journal of Political Science* 50(3):755–769.

van Benthem, Johan. 2011. *Logical Dynamics of Information and Interaction*. Cambridge: Cambridge University Press.

Zaller, John R. 1992. *The Nature and Origins of Mass Opinion*. Cambridge: Cambridge University Press.