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

Robert J. Carroll*

April 7, 2026

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

Abstract

This manuscript 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 Downsifiability—interpretability as a median-voter outcome under some ordering of platforms—requires the CMC, and that the CMC in turn forces the winning platform to be semantically maximal among all available broadcasts. Four Downsification conditions—full policy credibility, single-dimensional policy semantics, single-peaked welfare, and trigger-free voters—are jointly sufficient to ensure the CMC; the empirical record suggests all four are generically violated.

When the CMC fails, the semantic lattice delivers three results without analog in the spatial framework. First, the welfare-relevant majority-preference orbit is structurally confined to a lattice interval determined by the distribution of voter ideal understandings—an intermediate bound between core and Pareto set that spatial theory identifies as missing. Second, 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. Third, when voter ideals partition into incomparable blocs with one constituting a strict majority, competition forces equilibrium toward the majority bloc's ideal, providing a structural account of majoritarian polarization without exogenous spatial geometry.

The manuscript develops empirical predictions that are testable with existing survey infrastructure and distinguishable from what the spatial model generates, maps the update operator onto three classical epistemic traditions (AGM belief revision, Bayesian conditioning, and public announcement logic), and extends the framework to legislative bargaining, dynamic campaigns, and endogenous field evolution.

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.

But is that what mass politics looks like? Affective polarization between party identifiers has risen steadily since 1980 even as the extent of mass ideological divergence remains actively contested (Iyengar et al., 2019); partisan blocs disagree not only about which policies are good but about which factual claims are true; the same political vocabulary—“freedom,” “responsibility,” “government”—carries categorically different meanings for citizens reasoning through different conceptual systems (Entman, 1993; Lakoff, 1996); and campaigns routinely broadcast meaning content that is, by design, heard differently by different constituencies. The spatial model is not trying to predict any of this. It assumes that a platform is announced and that everybody sees it in the same space, interprets it at the same point, measures distance by the same rule, and decides on the basis of that distance. That is not a prediction of the model—it *is* the model, for better or for worse.

These observations are not new. Writing six years after Downs and sixty-three before your humble author, Stokes (1963) examined the spatial apparatus, identified four hidden axioms—unidimensionality, fixed structure, ordered policy dimensions, and a common frame of reference between the perceived spaces of voters and parties—and called for a successor formalization that retained its analytic ambition while treating the axioms themselves as cognitive variables rather than as fixed background assumptions. Stokes singled out the axiom of common reference as the most fragile of the four. The empirical literature on political behavior has been carrying out the empirical half of his program ever since.

Subsequent behavioral work has identified specific problem points. 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 non-attitudes that bear no stable connection to any underlying policy preference—so the difficulty is not where to put voters in a single ordered space but whether any such space exists for them. 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, and Bartels (2002) reinforced the point with the strongest panel-data evidence: partisan filtering applies to the perception of objective economic facts and not only to the evaluation of disputed policies, with bias magnitudes that do not diminish at higher levels of political information. Taber and Lodge (2006) added that the resistance is sharpest among the very voters spatial theory should be able to handle most cleanly: political sophistication predicts entrenchment rather than accuracy, with subjects exposed to challenges to their prior attitudes emerging more committed to those attitudes than before exposure.

Taken together, this lineage strains the formal tradition. The empirical record asks the spatial framework two structural questions it cannot answer with internal modifications. The first concerns the nature of the alternative space: Converse and the framing literature suggest there is no obvious shared space to put voters in. The second concerns the nature of agents' movement around that space under informational interactions: Zaller, Bartels, and Taber and Lodge suggest that any rigid model of learning from new information misrepresents how voters absorb and resist political content. A successor formalization needs to address both questions at once: it needs a model of the space, not just a stipulation that one exists, and it needs a model of how agents move through that space under the realistic assumption that information sometimes lands and sometimes is rejected.

Excellent formal work has made headway in this vein. Small modifications to the Downsian starting point have turned out to be insufficient: the problem is structural, not parametric, and adding dimensions or introducing valence does not address it. Hinich and Munger (1994) replace the high-dimensional policy space with a low-dimensional ideological space mediated by a voter-specific Bayesian inference rule, taking voter-relative perception seriously and identifying the operative space of competition as something other than the engineering policy space. Kollman, Miller and Page (1992, 1998) replace the rational-actor party with bounded-rational adaptive parties, show that the empirical regularity of centrist competition survives the relaxation of full information and global optimization, and identify voter heterogeneity in issue intensity as a primary driver of persistent platform separation. Bendor et al. (2011) provide the analytical follow-through on the bounded-rationality program, deriving structural divergence from satisficing dynamics under general assumptions about agent behavior. Each of these projects has delivered real results. The framework developed here resembles each of them in one respect or another. What it adds—and what none of them, for all their accomplishments, contains—is a model in which the meaning of political vocabulary is itself voter-relative; in which updating can resist and reverse rather than only accumulate; and in which two voters can hold political ideals that simply do not line up on the same scale. Each of these is a structural feature the empirical literature documents in detail. None is captured by any of the three frameworks named above.

Consider three voters in the same election, all with opinions about immigration. None of them is talking about the same thing. The first thinks about immigration through enforcement: border infrastructure works, it is the government's job, and the economic costs of unauthorized crossings are real. The second thinks about immigration through integration: a path to citizenship is legally viable, economically productive, and morally right. The third does not engage with the policy specifics at all—enforcement feels cruel, and a welcoming posture feels right. She could not name the relevant statute or estimate the fiscal impact, and she is not trying to.

A spatial model places all three on a left–right axis and assumes they disagree about *where* to stand on it. But they do not share an axis. The first two voters share factual vocabulary—they agree that enforcement exists and has economic consequences—but disagree completely on whether those consequences are good or bad. The third voter does not participate in the factual exchange at all; her engagement is entirely evaluative. A centrist platform that splits the difference on the facts speaks to the first two voters and says nothing to the third—not because it is too moderate, but because it is in a language she does not use.

The framework developed in this manuscript provides the machinery to make these distinctions precise: how agents can hold different *kinds* of meanings for the same political representations, when their semantic systems overlap and when they do not, and what electoral competition looks like when the shared policy space that the spatial model takes for granted does not exist.

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 correlate of Zaller's resistance mechanism, given empirical anchor in Bartels's partisan-filtering result—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 manuscript 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 only when a behavioral condition on the electorate's coalition structure—the *coalition mountain condition* (CMC)—holds, and the CMC in turn forces the winning platform to be semantically maximal among all available broadcasts. The theorem functions as a diagnostic: the CMC identifies when the Downsian spatial map might apply, and the four *Downsification conditions*—full policy credibility, single-dimensional policy semantics, single-peaked welfare functions, and trigger-free voters—characterize when it must. Each condition is individually demanding; their joint satisfaction is the exception rather than the rule. The four Downsification conditions are Stokes's four hidden axioms in formal dress.

We find three results that characterize electoral competition in the regime where the Downsification conditions fail.

The first is a structural bound on majority-preference cycling. McKelvey (1976) and Schofield (1983) showed that under multidimensional Euclidean preferences with no Condorcet winner, the majority-preference orbit is generically path-connected and can reach any platform from any other—the chaos result. Theorem 5.9 shows that the McKelvey-Schofield orbit, when computed in the lattice of understandings rather than in a Euclidean space, is structurally confined to the interval bounded below by the meet of voter ideal understandings and above by the join. Chaos still happens; it is no longer unconstrained. The result complements the computational answer to McKelvey developed by Kollman, Miller and Page (1992, 1998)—bounded-rational parties cannot in practice locate the destabilizing platform—by adding a structural constraint at the alternative-space level rather than at the agent-side level.

The second is a result about strategic communication in heterogeneous electorates. Corollary 5.12 shows that when voter credibility filters are heterogeneous enough that different types absorb disjoint components of a broadcast, the optimal campaign platform is polysemic: a single message that carries content optimized separately for each constituency and heard differently by each. This is the campaign-strategy phenomenon Hinich and Munger (1994) model as variance management of perceived ideological position; the lattice setting identifies a different mechanism for the same phenomenon. The lattice meet of two stances under heterogeneous credibility filters and the convex combination of two ideological positions in a Euclidean space are different objects on different structures, so the polysemy result is not a contradiction of Hinich and Munger's account of ideological combinatorics; it is a result on an object their framework does not contain, and it gives a formal account of how broad-tent coalition broadcasts actually operate.

The third is a result about non-convergence in two-candidate equilibrium. Corollary 5.15 shows that when voter ideal understandings partition into epistemic blocs that cannot be lined up on a shared scale, with one bloc constituting a strict majority, equilibrium platforms are forced toward the majority bloc's ideal and the minority is structurally excluded by an amount determined by the lattice geometry of the partition. Bendor et al. (2011, Prop. 3.4) reach a closely related anti-Downsian prediction—platform divergence as a structural equilibrium property rather than as a deviation—through a different mechanism: satisficing winners and searching losers under retrospective adaptive dynamics. The two mechanisms are independent and can operate simultaneously, so the same persistent divergence is overdetermined.

Section 2 develops the representational framework, with immigration as a running example. Section 3 develops the semantic update operator, grounded in cognitive dissonance theory. Section 4 derives the Downsification theorem and identifies the conditions under which the spatial model is recovered. Section 5 establishes the communication equilibrium and the three main results—the chaos bound, strategic polysemy, and majoritarian convergence. Section 6 develops the empirical predictions and operationalizes the key constructs for empirical research. Section 7 maps the framework onto three classical epistemic traditions: AGM belief revision, Bayesian conditioning, and public announcement logic. Section 8 develops three extensions: legislative bargaining, dynamic campaigns, and endogenous field evolution. Section 9 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.

Throughout, $\mathcal{R}$ and the meaning space M (definition 2.5) are assumed to be finite. This ensures that the understanding lattice $\mathscr{U}_{\mathcal{R}}$ is finite—a standing assumption used in the convergence results of section 5 and the extensions of section 8.

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 2.7) 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 8.3.

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]

To fix ideas, consider immigration as a representational domain. The field $\mathcal{R}$ includes, among many other elements, three representations that will recur as a running example: r_1 (border enforcement mechanisms), r_2 (paths to citizenship), and r_3 (the economic impact of immigration). These are not policy positions; they are representational resources available for making, contesting, or refusing claims about immigration. A voter may engage with all three, with a subset, or—as the epistemic layer will formalize—may be unable to distinguish some of them from one another. Different voters may assign categorically different meanings to the same representation: r_1 can carry the meaning "enforcement works" or the meaning "enforcement is cruel," and a single voter can hold both meanings simultaneously. The next three subsections build the apparatus to make these differences precise.

2.2 Epistemic stances.

The representational field is shared, but agents do not engage with it uniformly. A voter who follows immigration through policy journals, court rulings,

[2] Treating the structure of representational alternatives as a primitive of the model, rather than as a derivative of agents' preferences, has informal precursors in two literatures the framework formalizes. Bendor et al. (2011, ch. 1, §1.1) argue that mental representations are one of the two pillars of bounded rationality (alongside heuristics) and quote the Markman–Dietrich claim that "virtually all theories about cognition are based on hypotheses that posit mental representations as carriers of information"—but their formal models operate on a fixed exogenous policy space and never formalize the representational layer they identify as foundational. Sniderman and Bullock (2004), working in the political-psychology tradition, propose the hypothesis of *menu dependence*: consistency in public opinion is jointly conditional on the citizen and on the structure of the menu of choices presented to them, and political parties do the heavy lifting of producing menus that allow citizens to be coherent. The representational field $\mathcal{R}$ formalizes both moves: it is the structured space of representational resources prior to and shaping the agent's possible understandings, and the lattice $\mathscr{U}_{\mathcal{R}}$ that the next subsections construct on top of it is the formal apparatus that the verbal hypotheses of menu dependence and bounded-rationality representations require.

and statistical reports accesses a different portion of $\mathcal{R}$—and draws finer distinctions within it—than a voter whose engagement is limited to cable news and social media memes. The *epistemic stance* formalizes this variation: it captures which parts of $\mathcal{R}$ are available to an agent 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* $E_{\mathcal{R}}$*.*

Returning to the immigration example, consider how differently three voters engage with the field $\{r_1, r_2, r_3\}$. The enforcement voter and the integration voter both access all three representations and draw fine distinctions among them: their visible scopes are $\mathcal{R}_{e_A} = \mathcal{R}_{e_B} = \{r_1, r_2, r_3\}$, and their equivalence relations place each r_i in its own class. The affect-only voter, by contrast, does not distinguish enforcement mechanisms from economic impact—both register as "the harsh side of the debate"—while engaging fully with the path-to-citizenship representation. Her epistemic stance collapses r_1 and r_3 into a single equivalence class while keeping r_2 distinct: $\mathcal{R}_{e_C} = \{r_1, r_2, r_3\}$ with $r_1 \sim_{e_C} r_3$. She can access every representation the other voters can, but she draws fewer distinctions—her stance is representationally poorer in the sense formalized next.

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 $\mathcal{R}$:

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

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.

Despite the prevalence of incomparability, the space of epistemic stances admits well-defined operations for extracting what agents *do* share (the meet) and for characterizing what a hypothetical agent fluent in all schemas could access (the join). These operations—not the comparability of stances—are the structural engine of the results to follow: the meet will define semantic common ground between voter and candidate, and the join will bound how far communication can extend.

The epistemic lattice. Formally, the set $E_{\mathcal{R}}$ admits the following 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. In the immigration case, the meet of the enforcement voter's stance and an affect-only voter's coarsened stance e_C (which collapses $r_1 \sim r_3$) is not empty—both agents see all three representations—but the meet is strictly coarser than either: the common representational ground preserves only those distinctions both agents make.

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

The bottom $e_\bot$ 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 need for a separate semantic layer is not obvious from the standard framework, where representations and their interpretations are packaged together in a state space. In the present setting, the separation is essential. Recall that epistemic stances are partially ordered by representational richness, with incomparability as the generic condition. A coarsening of this ordering—where two stances are compared only by the *meanings* their holders extract from the same representations—requires a distinct formal object. Two agents with identical stances who read the same representation as "mechanism for collective deliberation" versus "site of elite capture" are not epistemically different; they are *semantically* different, and the updating, welfare, and equilibrium results to follow depend on this distinction.

The meaning space.

Definition 2.5. *The* meaning space M *is a nonempty set of* primitive meanings.

Like $\mathcal{R}$, the space M is left unstructured for generality.[3] In applications, its elements might include:

[3] 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

- *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.

To see the semantic layer in isolation, suppose all three immigration voters share the maximal epistemic stance $e_\top$—full scope, each representation in its own class—so that the action is purely in what meanings they attach. Let $M = \{T, F, G, B\}$, where T and F are factual primitives ("the policy

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 $\mathsf{Conf}(u_i, u_j)$ (defined in Section 3 below) 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. The closest formal-theory ancestor of the receiver-side credibility filter α_{ij} is the voter-specific mapping vector v_i in the predictive-mapping construction of Hinich and Munger (1994, ch. 10), which translates a candidate's ideological position into expected positions on policy issues via a voter-specific linear projection; α_{ij} generalizes this construction from a continuous linear projection between two Euclidean spaces to a discrete admission rule on structured meaning content, and gains the ability to represent incomparable voter understandings that the linear apparatus cannot. In the behavioral-formal tradition, the closest continuous-weight precursor is the strength–ideal-position correlation studied by Kollman, Miller and Page (1998), in which voters' issue-intensity weights covary with their ideal positions; α_{ij} is the discrete meaning-theoretic analog of that covariation structure.

works," "the policy fails") and G and B are evaluative primitives ("desirable," "undesirable"). A representation can carry any subset of M: assigning $\{T, G\}$ to r_1 says "enforcement works and is good"; assigning $\{T, B\}$ says "enforcement works but is bad." The enforcement voter assigns $\mu_A(r_1) = \{T, G\}$, $\mu_A(r_2) = \{F, B\}$, $\mu_A(r_3) = \{T, B\}$: enforcement is effective and good, paths to citizenship are based on false premises and bad, and immigration's economic impact is real but harmful. The integration voter assigns $\mu_B(r_1) = \{T, B\}$, $\mu_B(r_2) = \{T, G\}$, $\mu_B(r_3) = \{T, G\}$: enforcement is real but cruel, paths to citizenship are viable and good, and immigration is economically beneficial. The affect voter assigns $\mu_C(r_1) = \{B\}$, $\mu_C(r_2) = \{G\}$, $\mu_C(r_3) = \emptyset$: enforcement is bad and citizenship is good, with no factual content at all and no engagement with the economic dimension.

These three semantics illustrate pure semantic disagreement. Voters A and B agree on the factual status of enforcement ($T \in \mu_A(r_1) \cap \mu_B(r_1)$) but disagree evaluatively (G versus B). Voter C's engagement is entirely evaluative: she carries no factual content and cannot distinguish a platform that says "enforcement reduces crossings" from one that says "enforcement is ineffective," because neither factual proposition intersects her meaning system. This is not ignorance in the classical sense—voter C is not uncertain about which state of the world obtains. She simply does not engage with the factual dimension of immigration at all; her representational apparatus maps the issue entirely into evaluative terms. A Bayesian model would represent C's stance as a diffuse prior over states, but that framing misidentifies the nature of the difference: C is not epistemically impoverished about the facts; she is semantically disengaged from them.

The semantic layer captures this variation independently of the epistemic layer—all three voters see the same representations and draw the same distinctions.

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

Definition 2.6. *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. (The domain is the full partition $\mathcal{R}/{\sim_e}$, which includes the invisible class I_e; the perceived representations *$\mathcal{R}_e/{\sim_e}$ of Definition 2.2 are the non-invisible classes of this partition. When*

$\mathcal{R}_e = \mathcal{R}$ *there is no invisible class and the distinction is vacuous.) A semantics* μ_e *is* e-consistent *when it satisfies the condition* $\mu_e(I_e) = \emptyset$. *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: which representations the agent can access and distinguish (the epistemic stance) and what meanings the agent attaches to those representations (the semantics). Two agents can differ in either layer or both. The immigration voters illustrate the full range: voters A and B share the same epistemic stance but differ semantically; voter C (in her coarsened version from Section 2.2) differs from both in the epistemic layer as well. The definition below packages both layers into a single object and organizes the space of all such objects into a structure—a complete lattice—that supports the comparison, aggregation, and updating operations that drive the rest of the manuscript.

Definition 2.7. *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. The fiber structure of $\mathscr{U}_{\mathcal{R}}$ poses an immediate problem: how can we compare understandings held at different epistemic levels? If voter A operates at $e_\top$ and voter C at a coarser stance e_C, any comparison requires translating one agent's meanings into the other's representational vocabulary. The transport operations below formalize this translation, and their interplay—captured by a Galois connection[4]—ensures that the comparison is canonical: there is a unique best way to lift meanings up or project them down across epistemic 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.

Proposition 2.8. *The pair* $(\mathrm{Crs}_{e_2 \to e_1}, \mathrm{Ref}_{e_1 \to e_2})$ *is a monotone Galois connection: for any* $\mu_{e_1} \in M_{e_1}$ *and* $\mu_{e_2} \in M_{e_2}$,

$$\mathrm{Crs}_{e_2 \to e_1}(\mu_{e_2})(C) \subseteq \mu_{e_1}(C) \ \text{ for all } C \quad \Longleftrightarrow \quad \mu_{e_2}(D) \subseteq \mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1})(D) \ \text{ for all } D.$$

[4] A *monotone Galois connection* (equivalently, an *adjunction between posets*) between two partially ordered sets $(P, \leq)$ and $(Q, \leq)$ is a pair of monotone maps $f : Q \to P$ and $g : P \to Q$ satisfying $f(q) \leq p \Leftrightarrow q \leq g(p)$ for all $p \in P$, $q \in Q$. The map f is the *lower adjoint* and g is the *upper adjoint*; here Crs is the lower adjoint and Ref is the upper adjoint. The defining equivalence has three important consequences. First, the round-trip compositions satisfy $\mathrm{id} \leq g \circ f$ and $f \circ g \leq \mathrm{id}$: coarsening and then refining can only add meaning (the round-trip is *inflationary*), while refining and then coarsening can only lose it (the round-trip is *deflationary*). Semantically, this says that translating down to a coarser vocabulary and back up again is lossy in a controlled, one-directional way—information is not scrambled but systematically rounded. Second, the lower adjoint Crs preserves all joins (unions of meaning sets) and the upper adjoint Ref preserves all meets (intersections). This is the structural guarantee that pooling meanings across subclasses and distributing meanings across subclasses are mutually coherent operations. Third, the Galois connection determines a *closure operator* $\mathrm{Ref} \circ \mathrm{Crs}$ on the finer semantics and an *interior operator* $\mathrm{Crs} \circ \mathrm{Ref}$ on the coarser semantics, whose fixed points are the semantics that survive round-trip translation without distortion—the *translatable* fragment of each level. For a systematic treatment of Galois connections between complete lattices, see Davey and Priestley (2002, Ch. 7).

Both directions reduce to the same 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. In the immigration example, $u_A\preceq_{\mathscr{U}}u_B$ would require $\mu_A(r_i)\subseteq\mu_B(r_i)$ for all i (since both voters share $e_\top$). This fails at r_1: $\mu_A(r_1)=\{T,G\}$ is not contained in $\mu_B(r_1)=\{T,B\}$, because $G\notin\mu_B(r_1)$. The same check in reverse fails at r_2. The two voters are incomparable: neither understanding contains the other, and no amount of reindexing or relabeling can force them onto a common scale.

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 2.9. *$(\mathscr{U}_{\mathcal{R}},\preceq_{\mathscr{U}})$ 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.

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.

The immigration voters make the lattice operations concrete. Consider the meet of voters A and B (at the shared stance $e_\top$, the meet reduces to pointwise intersection of meaning sets):

$$u_A \wedge u_B : \quad r_1 \mapsto \{T\}, \quad r_2 \mapsto \emptyset, \quad r_3 \mapsto \{T\}.$$

The semantic common ground between the enforcement voter and the integration voter is $(\{T\}, \emptyset, \{T\})$: they share the factual judgment that enforcement has consequences and that immigration has economic impact, but they share *nothing* evaluative and have zero semantic overlap on paths to citizenship. This is the formal content of the familiar observation that "they agree on the facts but not on the values"—the meet strips away everything contested and leaves only what survives intersection.

Now consider the meet of voters A and C:

$$u_A \wedge u_C : \quad r_1 \mapsto \emptyset, \quad r_2 \mapsto \emptyset, \quad r_3 \mapsto \emptyset.$$

The result is $u_\perp$: zero common ground. The enforcement voter's factual-evaluative system and the affect voter's pure-evaluative system share no semantic content at all—not even at the level of which representations carry meaning. This is not a matter of disagreement but of *incommensurability*: the two meaning systems do not overlap.

Incomparability is the key structural feature. Voters A and B are incomparable in $\preceq_{\mathscr{U}}$: A assigns G to r_1 where B does not ($u_B \not\succeq_{\mathscr{U}} u_A$), and B assigns $\{T, G\}$ to r_2 where A assigns $\{F, B\}$ ($u_A \not\succeq_{\mathscr{U}} u_B$). Neither understanding is richer than the other; each carries semantic content the other lacks. Voters A and C are also incomparable, but their incomparability is more severe—their meet is the bottom element, indicating that no communication built from shared semantic content is possible without first building new common ground.

The join $u_A \vee u_B = (\{T, G, B\}, \{T, F, G, B\}, \{T, G, B\})$ pools both meaning systems, producing an understanding that is internally conflicted—at r_2 it simultaneously carries T and F, G and B—but representationally maximal: the stance of a hypothetical agent fluent in both voters' semantic systems.

The tension at r_2 is informative, not paradoxical. It says that the full semantic picture of paths to citizenship includes both "viable and good" and "false and bad"—the pooled understanding does not resolve this conflict but rather *registers* it as a feature of the electorate's semantic landscape.

These computations preview the central structural challenge for electoral competition. In a classical spatial model, two voters who disagree about policy can be located on a common ideological dimension, and a median voter theorem organizes the equilibrium. When voter ideals are incomparable, no such dimension exists: the enforcement voter and the integration voter cannot be ranked on a single axis, because each carries semantic content the other lacks. The *semantic dimension* of the electorate—the width of the largest antichain in the voter ideal poset—measures the severity of this problem. When the semantic dimension exceeds one, the Downsian framework cannot be recovered, and the classical median voter result gives way to the structural results developed in Sections 4 and 5: a chaos bound, strategic polysemy, and majoritarian convergence that have no spatial analog.

Figure 1 summarizes the lattice structure of the three immigration voters.

The field-anchor. The lattice operations acquire their full political force when understandings are anchored to agents. An agent's *ideal understanding* u_v^* is not an aspirational description of how the world should be. It is the understanding that anchors the agent's welfare evaluation: as Section 4 will formalize, a voter's welfare from any candidate broadcast $\phi(x)$ is determined by $\mathcal{W}(\phi(x) \wedge u_v^*)$—the meet captures the semantic common ground between what the candidate says and what the voter values. Changing u_v^* reshapes the entire welfare landscape, not just the peak. In the immigration case, voter C's ideal $u_C^* = (\{B\}, \{G\}, \emptyset)$ generates a welfare field that is entirely flat in the factual dimension: a platform broadcasting only factual content about enforcement produces zero common ground with u_C^*, regardless of what the facts say. The ideal understanding is a *field-anchor*—it generates the voter's evaluative field over the lattice of possible broadcasts—and the lattice structure determines how that field varies with the semantic content of political communication.

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

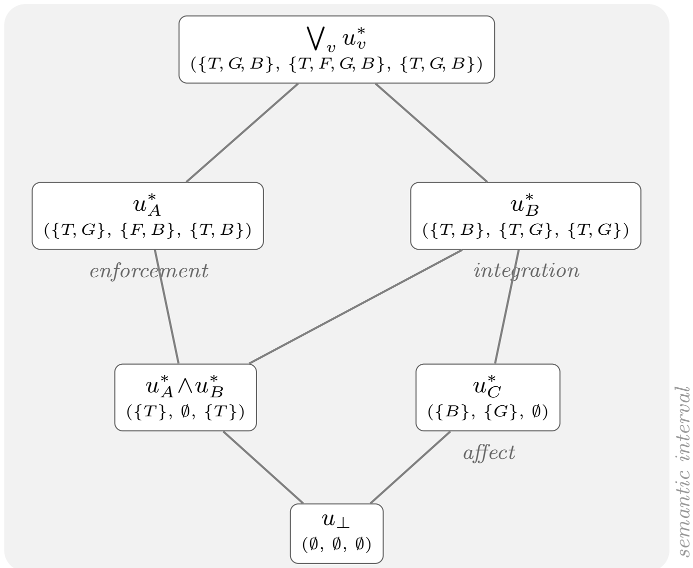

Figure 1: *Hasse diagram of the immigration voters' understanding lattice. Nodes show the voter label and semantic content $(\mu(r_1), \mu(r_2), \mu(r_3))$ at each lattice element. Voter ideals u_A^* and u_B^* are incomparable (no edge connects them); u_C^* lies below u_B^*. The three-way meet $\bigwedge_v u_v^* = u_\bot$ reflects zero semantic common ground across the electorate. The shaded region is the semantic interval $[\bigwedge_v u_v^*, \bigvee_v u_v^*]$, which bounds the welfare-relevant majority-preference orbit (Theorem 5.9).*

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.

The formal comparison with classical updating frameworks—Bayesian conditioning, AGM contraction, and dynamic epistemic logic—and the conditions under which each is recovered as a special case of the semantic update operator, is developed in Section 7.

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.

The theoretical foundation for this asymmetry is older than the motivated reasoning literature. Festinger's (1957) theory of cognitive dissonance postulates that pairs of cognitions can be consonant, dissonant, or irrelevant; that the coexistence of dissonant cognitions produces psychological discomfort; and that this discomfort motivates the agent to reduce the dissonance—by removing dissonant cognitions, adding consonant ones, or reducing the importance of the dissonant elements (for an overview of current perspectives on the theory, see Harmon-Jones and Mills, 2019). The update operator developed below formalizes each of these mechanisms within the lattice $\mathscr{U}_{\mathcal{R}}$. The conflict relation Cf identifies dissonant pairs of meanings; the conflict set $\mathsf{Conf}(u_i, u_j)$ measures the magnitude of the dissonance (Festinger's "dissonance ratio": the proportion of the agent's cognitive system in tension with the incoming content). The forward update corresponds to adding consonant cognitions—absorbing compatible meanings through the credibility filter—while the baseline β corresponds to removing dissonant cognitions—stripping the meanings that generate conflict. The trigger threshold τ parameterizes the agent's tolerance for dissonance: below the threshold, the agent absorbs despite tension; above it, the discomfort is sufficient to motivate restructuring. What Taber and Lodge document experimentally—disconfirmation bias among politically sophisticated subjects—is the belief-disconfirmation paradigm of cognitive dissonance theory (see Harmon-Jones and Mills, 2019, for an overview), applied to political attitudes: the incoming information creates dissonance with the agent's existing commitments, and the agent reduces the dissonance not by updating but by contracting.

The lattice structure of $\mathscr{U}_{\mathcal{R}}$ is what makes this tractable. Because the space of understandings is a complete lattice, the update operator has well-defined monotonicity properties: the forward update is extensive (Lemma 3.5), the baseline is a join in $\mathscr{U}_{\mathcal{R}}$ whose position is determined by the conflict structure, and the dependence on parameters—credibility filter, trigger threshold—is monotone. Without the lattice, the update would be an arbitrary map on an unstructured space, and the equilibrium analysis of Section 5 would have no fixed-point theorems to invoke. The lattice is not a convenience; it is the structural precondition for the results that follow.

3.1 Conflict structure.

Before specifying how understandings change, we need a formal notion of when two understandings are in tension. The conflict relation captures incompatibility at the level of individual meanings; the conflict set aggregates this to the level of representation classes, measuring how broadly one understanding clashes with another.

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.

The choice of conflict relation is an analytical decision with empirical consequences, not a modeling convenience. A sparse conflict relation—few pairs in Cf—produces an electorate in which encounters rarely trigger reversion; the forward update dominates, voters accumulate meanings, and the communication equilibrium of Section 5 converges quickly. A dense conflict relation—many pairs in Cf—produces the opposite: encounters routinely trigger defensive contraction, voters shed content, and the equilibrium involves widespread impoverishment. The empirical question is which regime holds for a given domain. On issues where voters have well-articulated schemas with clear internal tensions—abortion, gun control, immigration—the conflict relation is dense, triggers fire frequently, and the sophistication paradox is severe. On issues where voters have minimal prior commitments and few meanings in play—technical regulatory questions, most aspects of foreign policy for most voters—the conflict relation is effectively empty, the trigger is inert, and the forward update operates unimpeded.

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.

In the immigration example, the natural conflict relation on $M = \{T, F, G, B\}$ is $\mathsf{Cf} = \{(T, F), (F, T), (G, B), (B, G)\}$: factual values conflict with each other and evaluative values conflict with each other. Consider voter A encountering an integration understanding u_L with semantics $\mu_L = (\{T, B\}, \{T, G\}, \{T, G\})$—the integration voter's understanding from Section 2. At the shared epistemic level $e_\top$, the coarsening map is the identity, and the conflict meanings are:

$$F_{AL}(r_1) = \{F, G\}, \qquad F_{AL}(r_2) = \{F, B\}, \qquad F_{AL}(r_3) = \{F, B\}.$$

The meanings in voter A's semantics that fall within these conflict sets are $\mu_A(r_1) \cap F_{AL}(r_1) = \{G\}$, $\mu_A(r_2) \cap F_{AL}(r_2) = \{F, B\}$, and $\mu_A(r_3) \cap F_{AL}(r_3) = \{B\}$—all nonempty. Thus $\mathsf{Conf}(u_A, u_L) = \{r_1, r_2, r_3\}$: every representation class is in conflict. The enforcement voter's entire semantic system is under pressure from the integration understanding.

3.2 Source credibility.

The conflict structure determines when an encounter is contentious; the credibility structure determines what content passes through. Together they govern the two branches of the update operator: forward absorption (filtered by credibility) and defensive reversion (triggered by conflict).

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.[5] 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.

[5]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 manuscript requires it.

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 $\alpha_{ij}([r]) = \emptyset$.

The empirically interesting case is when credibility filters are *heterogeneous across voters*. Iyengar et al. (2019) document that affective polarization—the tendency to distrust and dislike members of the opposing party—has grown dramatically in the American electorate, even as the extent of mass ideological divergence remains actively contested. Bartels (2002) provides the observational complement from NES panel data: strong Democrats and strong Republicans give systematically opposite answers to objective questions about unemployment and inflation rates, with partisan bias of magnitudes $t = 18.3$ and $t = 15.3$ respectively. Two features of Bartels's finding anchor the filter as a structural primitive rather than a low-information shortcut. First, the bias acts on factual content, not only on evaluative content, so the filter cannot be confined to M_{affect}; it acts on M_{policy} as well. Second, the magnitude of the bias is comparable for high- and low-information respondents, ruling out the interpretation of α_{ij} as a heuristic substitute for information processing.[6] In the present framework, affective polarization corresponds to credibility filters that are concentrated on $M_{\text{affect}} \cup M_{\text{social}}$ (the affective and social components of the meaning space; see the formal partition in section 3.5) for in-group sources and attenuated or empty for out-group sources. When different voter types have filters concentrated on disjoint subsets of M—one type absorbs factual content and filters affect, another absorbs affect and filters fact—the candidate's broadcast is heard differently by each type. A single message, passing through heterogeneous filters, produces different post-update understandings for different voters. This is the mechanism underlying the polysemy results of Section 5: strategic ambiguity in equilibrium is not a

[6]The experimental political-cognition literature documents process-side companions to this structural filter—the information-acquisition and decision-strategy heuristics catalogued by Lau and Redlawsk (2006)—but heuristics and credibility filters are distinct objects: heuristics are decision rules that operate *on* the voter's representation, while α_{ij} is part of the representation itself.

failure of the candidate to communicate clearly but a structural consequence of heterogeneous credibility in the electorate.

3.3 The update operator.

With conflict and credibility in hand, we can specify the full update. The operator has two branches: a *forward update* that absorbs filtered content from the source and enriches the receiver's understanding, and a *baseline* that contracts the receiver's understanding by stripping conflicting content. Which branch fires depends on a threshold parameter that measures the receiver's tolerance for conflict.

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.

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

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

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 $e^\wedge \preceq_E e_i$). Since $e^\wedge \preceq_E e_j$ by definition of the meet, $\mathrm{Crs}_{e_j \to e^\wedge}$ is well-defined, and we can define the *meet-level credibility filter*

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

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 $[R]$. The forward update at the semantic stage is then

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

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 forward update in the immigration example. Since voter A and the integration understanding u_L share $e_\top$, the formula reduces to $\mu_A^+(r_k) = \mu_A(r_k) \cup (\alpha_{AL}(r_k) \cap \mu_L(r_k))$. Under *full credibility* ($\alpha_{AL} = M$ everywhere), the filter is transparent:

$$\mu_A^+(r_1) = \{T, G\} \cup \{T, B\} = \{T, G, B\},$$
$$\mu_A^+(r_2) = \{F, B\} \cup \{T, G\} = \{T, F, G, B\},$$
$$\mu_A^+(r_3) = \{T, B\} \cup \{T, G\} = \{T, G, B\}.$$

At r_1 and r_3, voter A now holds evaluative ambivalence: enforcement is simultaneously good (G) and bad (B). At r_2, the full meaning space is active—factual and evaluative tension coexisting. This is the augmentation interpretation made concrete: holding G and B simultaneously is not incoherence (the conflict relation governs updating, not the content of understandings) but the formal representation of attitudinal ambivalence.

Under *factual-only credibility* ($\alpha_{AL}(r_k) = \{T, F\}$ for all k), the filter admits only factual content:

$$\begin{aligned}
\mu_A^+(r_1) &= \{T, G\} \cup (\{T, F\} \cap \{T, B\}) = \{T, G\},\\
\mu_A^+(r_2) &= \{F, B\} \cup (\{T, F\} \cap \{T, G\}) = \{T, F, B\},\\
\mu_A^+(r_3) &= \{T, B\} \cup (\{T, F\} \cap \{T, G\}) = \{T, B\}.
\end{aligned}$$

Voter A absorbs the factual claim that paths to citizenship are viable (T enters at r_2) but does not absorb the evaluative endorsement (G is blocked by the filter). The result at r_2 is $\{T, F, B\}$: factual ambivalence (T and F simultaneously active) alongside the retained negative evaluation B. The credibility filter does the psychological work that motivated reasoning does in the experimental literature: it controls which dimensions of an incoming understanding can penetrate, leaving the voter's evaluative stance undisturbed even as new factual content is absorbed.

The contrast between the two cases illuminates what is at stake in the credibility parameter. Under full credibility, voter A's post-update understanding at r_2 is $\{T, F, G, B\}$—the entire meaning space, carrying both factual claims and both evaluative assessments simultaneously. Under factual-only credibility, r_2 carries $\{T, F, B\}$: the voter has acquired factual ambivalence ("paths to citizenship might work *and* might fail") but the evaluative frame remains undisturbed (B persists, G is excluded). The full-credibility voter is more *informed* in the standard sense—more meanings are active—but is also more internally conflicted and more vulnerable to the trigger mechanism. A voter holding $\{T, F, G, B\}$ at a representation class is a *conflict magnet*: any future encounter whose content includes a factual or evaluative claim will generate a conflict, because the voter already holds the opposing meaning. The factual-only voter, by contrast, is selectively exposed: vulnerable to factual challenges but shielded from evaluative ones. This asymmetry is not an artifact of the example; it is a structural feature of the model, and it drives the sophistication paradox formalized below.

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:

$$\beta(u_i, u_j) = \bigvee S_{ij}, \quad \text{where}$$

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

This is AGM contraction, restated in $\mathscr{U}_{\mathcal{R}}$[7]: 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 S_{ij} iff $\mathsf{Conf}(u_i, u_j) = \emptyset$—that is, iff u_i has no conflicting content with u_j at any representation class. When $\tau_i > 0$, this is strictly stronger than the trigger-passing condition $|\mathsf{Conf}(u_i, u_j)| \leq \tau_i$: the trigger can pass even when $u_i \notin S_{ij}$. 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.

In the immigration example, with $|\mathsf{Conf}(u_A, u_L)| = 3$ and threshold $\tau_A = 2$, the trigger fires. The baseline strips the conflicting meanings at each class: at r_1, removing G (which is in $F_{AL}(r_1) \cap \mu_A(r_1)$) leaves $\{T\}$; at r_2, removing F and B (both in $F_{AL}(r_2) \cap \mu_A(r_2) = \{F, B\}$) leaves $\emptyset$; at r_3, removing B leaves $\{T\}$. The baseline is $\beta(u_A, u_L) = (\{T\}, \emptyset, \{T\})$: bare factual acknowledgments at r_1 and r_3, and total semantic disengagement at r_2. Voter A began with a rich evaluative stance on immigration—enforcement is good, paths to citizenship are wrong, the economy takes a real hit—and ends with nothing but the bare factual skeleton. The cost of defending against the integration understanding is the evacuation of the citizenship dimension entirely and the loss of all evaluative content elsewhere.

Lemma 3.5. *For any understandings* u_i, u_j *with* $|\mathsf{Conf}(u_i, u_j)| \leq \tau_i$ *(the trigger does not fire):*

$$u_i \ \preceq_{\mathscr{U}} \ T_+(u_i, u_j).$$

The forward update is weakly above the input: encountering another understanding can only enrich, never impoverish, when the trigger does not fire.

[7] 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.

3.4 Parameter dependence.

The previous subsection defined the update operator for fixed credibility and trigger threshold. In practice, these parameters vary across voters and across voter-source pairs: some voters are more receptive than others, some more tolerant of conflict. How does the update vary as the parameters change? The answer is that the model exhibits clean monotone structure—higher credibility and higher thresholds produce richer post-update understandings—and this monotonicity has a sharp empirical corollary in the sophistication paradox.

Lemma 3.6. *For fixed epistemic expansion rule and trigger threshold, if* $\alpha_{ij}([r]) \subseteq \alpha'_{ij}([r])$ *for all* $[r]$ *and neither credibility fires the trigger, then* $T_+(u_i, u_j; \alpha_{ij}) \preceq_{\mathscr{U}} T_+(u_i, u_j; \alpha'_{ij})$*.*

Higher credibility produces richer forward updates: more of j's meanings pass the filter and are absorbed.

Lemma 3.7. *For fixed* α_{ij}*, the output* $T_{\alpha_{ij},\tau}(u_i, u_j)$ *is weakly increasing in* τ*.*

A lower threshold makes reversion more likely. Since $\beta(u_i, u_j) \preceq_{\mathscr{U}} u_i \preceq_{\mathscr{U}} T_+(u_i, u_j)$ (Proposition 3.4 and lemma 3.5), 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.

In Festinger's (1957) terms, the baseline implements the most drastic of the dissonance-reduction strategies: the agent removes the dissonant cognitions outright rather than reinterpreting them or adding consonant cognitions to dilute the ratio. The alternative strategies—adding consonant cognitions (the forward update) or reducing the importance of the dissonant elements (the credibility filter)—are available only when the dissonance is below the threshold. Above it, the lattice structure forces a clean contraction: the baseline is the join of all understandings below u_i that avoid the conflicting content, and by Proposition 3.4 this join is strictly below u_i. There is no middle ground in the model between absorption and contraction—a feature that is empirically realistic. Taber and Lodge's (2006) experimental subjects, exposed to information challenging prior political attitudes, did not partially absorb it; they generated counterarguments, dismissed the source, and showed increased attitude entrenchment relative to unexposed controls. The encounter did not leave u_i unchanged but strictly impoverished it—the formal analog of Festinger's observation that dissonance reduction, once triggered, restructures the cognitive system rather than merely dampening the dissonant element.

Lemma 3.8 (Sophistication paradox)**.** *Let $u_i \preceq_{\mathscr{U}} u_i'$ at the same epistemic level ($e_i = e_{i'}$, so $\mu_i([r]) \subseteq \mu_{i'}([r])$ for all $[r]$). Then $\mathsf{Conf}(u_i, u_j) \subseteq \mathsf{Conf}(u_{i'}, u_j)$: richer priors produce weakly larger conflict sets.*

Moreover, if τ is such that the trigger fires for u_i (i.e., $|\mathsf{Conf}(u_i, u_j)| > \tau$), then it also fires for $u_{i'}$. In that case, $\beta(u_{i'}, u_j) \prec_{\mathscr{U}} u_{i'}$ and $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$ by Proposition 3.4, and the richer voter's reversion strips weakly more content: the baseline must remove conflicting meanings at every class in Conf, and since $\mathsf{Conf}(u_{i'}, u_j) \supseteq \mathsf{Conf}(u_i, u_j)$, the richer voter has weakly more classes to evacuate.

This is the formal content of the sophistication paradox documented by Taber and Lodge (2006): their most politically sophisticated subjects—those with the richest prior understandings—showed the strongest disconfirmation effect. The lemma says why. A richer understanding provides more meanings that can conflict with an incoming frame, so the trigger fires more readily and the baseline drops further—this is the lattice-theoretic version of Festinger's prediction that the magnitude of dissonance increases with the number of dissonant cognitions. In the immigration example, voter A holds $\mu_A(r_2) = \{F, B\}$—two meanings, both of which conflict with the integration understanding at r_2. A simpler voter holding only $\{B\}$ at r_2 has one meaning at risk and loses only that; voter A must strip both F and B to reach the baseline, evacuating the citizenship dimension entirely. Representational richness is not a route to Downsian rationality but an obstacle to it: the voters who most resemble the rational actor of spatial theory are precisely those for whom the Downsification conditions of Section 4.7 are most severely violated.

3.5 Special cases.

The general update operator has many moving parts: a conflict relation, a credibility filter, a trigger threshold, an epistemic expansion rule. To build intuition for what each part does—and to connect the operator to familiar models from the literature—it helps to examine what happens when specific parameters take extreme values. The three cases below correspond to three regimes of political communication: affect-only reception, full-policy absorption, and classical Bayesian conditioning. Each is recovered by restricting the credibility filter and trigger threshold; together they span the range from the generic (warm glow) to the knife-edge (Bayesian).

We partition M into semantically typed subsets: M_{policy} (factual and policy-implication meanings), M_{affect} (affective and evaluative meanings), and M_{social} (social-association and identity meanings). The partition is an analyst's choice, not a structural feature of the framework; different partitions yield different special cases below and different Downsification conditions in Section 4.7.

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).[8] 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. Consider a voter whose credibility filter admits only affective content encountering a policy-dense broadcast about net neutrality—bandwidth regulation, interconnection fees, common-carrier obligations. The technical content is filtered out entirely; what passes through the filter is "this is about fairness" or "this is about

[8]The warm-glow case is the formal counterpart of the *valence-issue* regime identified by Stokes (1963, pp. 373–374): political controversies in which voters share a common evaluation (peace, prosperity, honesty) and the electoral question is only which candidate is credibly linked to the valued condition, rather than where candidates locate on an ordered policy dimension. When credibility filters admit only evaluative content, the electorate is operating in Stokes's valence regime, and the spatial machinery breaks down for the structural reason he identified. The closest formal-theory ancestor of the warm-glow construction is Schuessler (2000), who develops a microfoundation for expressively motivated participation in which voting is an act of identity expression and campaigns strategically manipulate the symbolic meaning of participation; M_{affect} supplies precisely the expressive payoff his framework formalizes.

freedom"—affective framings that attach to the voter's existing evaluative associations. The voter's post-update understanding carries new affective content about net neutrality but zero policy content, regardless of how much policy the broadcast contained.

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. The implication for the updating model is sharp: if the typical voter's credibility filter excludes most policy content and the typical voter's trigger threshold is finite, then the *generic* update is not the forward update but the baseline. The section's formal apparatus—forward update, credibility filter, trigger, baseline—is not a sequence of increasingly exotic mechanisms appended to handle pathologies. It is a decomposition of the update process in which the baseline is the typical outcome and the forward update is the special case that requires demanding conditions to activate. This inverts the standard framing in which rational updating is the default and "bias" is the deviation. In the semantic framework, defensive contraction is the structural norm; unimpeded absorption is the exception that requires explanation.

Bayesian conditioning. The framework contains classical Bayesian updating as a degenerate case, recoverable under five jointly restrictive conditions.

Proposition 3.9. *Suppose:*

1. *fixed epistemic level:* $e'_i = e_i$ *(no expansion);*
2. *single-valued semantics:* $\mu_i([r]) = \{p_i([r])\}$ *with* $p_i([r]) \in [0, 1]$ *and* $\sum_{[r]} p_i([r]) = 1$*;*
3. *full credibility:* $\alpha_{ij}([r]) = M$ *for all* $[r]$*;*

4. *deactivated trigger:* $\tau_i = \infty$*;*

5. *replacement interpretation: when the forward update produces a multi-valued meaning set* $\{p_i([r]), p_j([r])\}$*, interpret it as the posterior* $p_i([r] \mid E)$ *rather than retaining both values.*

Then the update $T_{\alpha_{ij},\tau_i}(u_i, u_j)$ *recovers standard Bayesian conditioning:* $p_i^+([r]) = p_i([r] \mid E)$ *where* $E = \{[r] : p_j([r]) > 0\}$*.*

Proof. Under condition (ii), each meaning set is a singleton probability. Single-valued probability assignments cannot conflict in the sense of Definition 3.1—a probability value does not negate another—so $\mathsf{Conf}(u_i, u_j) = \emptyset$. Combined with condition (iv), the trigger never fires and the forward update is always applied. Under condition (iii), the credibility filter is transparent: all of j's content passes through. The base forward update adds j's probability to i's meaning set: $\mu_i^+([r]) = \{p_i([r])\} \cup \{p_j([r])\}$, a two-element set. Condition (v) provides the interpretive rule that collapses this two-element set to a posterior: read j's broadcast as an event $E = \{[r] : p_j([r]) > 0\}$ and set $p_i^+([r]) = p_i([r])/p_i(E)$ for $[r] \in E$, 0 otherwise. This is standard conditioning. □

Each condition individually is restrictive; their conjunction defines a knife-edge. Condition (v) deserves particular comment: it is a convention about how to read the output of the forward update, not a mechanism that activates under specific conditions. The base operator produces a set of meanings; condition (v) says to interpret that set as a posterior distribution rather than as a collection of simultaneously held values. Since conditions (ii) and (iv) together ensure $\mathsf{Conf} = \emptyset$—the probability-valued meanings never conflict—the replacement interpretation is purely definitional here, providing the bridge between the framework's set-valued semantics and the Bayesian apparatus's single-valued posteriors. Relaxing any condition moves the update away from Bayesian conditioning: dropping condition (i) allows epistemic expansion; dropping (ii) allows set-valued meanings (and hence conflict); dropping (iii) introduces credibility filtering; dropping (iv) activates the trigger; dropping (v) retains the multi-valued output as is, which is the base case of Definition 3.3. The augmentation interpretation—retaining both values—is the strictly more general operation. This is a substantive claim about political cognition: voters do not discard prior meanings upon encountering new content; they accumulate both "enforcement works" and "enforcement fails" as simultaneously active meanings for the same representation class, at least until conflict triggers defensive contraction.

With the static framework (Section 2) and the dynamics (this section) in place, the next two sections put them to work: Section 4 applies the update operator to electoral competition and identifies the conditions under which the classical Downsian model can be recovered, and Section 5 characterizes what happens when those conditions fail.

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*

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

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, $\mathsf{Cf} \subseteq M \times M$ is the conflict relation (Definition 3.1), $u_v^0,\ u_v^* \in \mathscr{U}_{\mathcal{R}}$ 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 a modeling device. Sections 2 and 3 developed a theory of how voters hold and process political content—representations, meanings, stances, updating—but that theory is purely cognitive. It says nothing about what voters *want*. To connect the cognitive framework to electoral competition, we need a way to distinguish, from the voter's perspective, a broadcast that serves her well from one that does not. The simplest way to do this is to give each voter an anchor in the understanding lattice—an understanding u_v^* against which incoming broadcasts can be evaluated—and define "serves her well" as "shares a lot of semantic content with u_v^*." That is what the ideal understanding is: an evaluative anchor, not a preference over outcomes, not a belief about how the world should be, and not an aspiration.

The word "ideal" is potentially misleading. It does not mean "best" in any normative sense; it means "the understanding from which the voter evaluates." Voter A's ideal $u_A^* = (\{T, G\}, \{F, B\}, \{T, B\})$ is simply the semantic content that A brings to the election: enforcement works and is good, paths to citizenship are wrong and bad, the economy takes a real hit. This is A's settled cognitive state—the state that the campaign will perturb through the update operator of Section 3. Calling it "ideal" distinguishes it from the initial understanding u_v^0 (which may differ from u_v^* if the voter has not yet fully processed prior information) and from the post-campaign understanding $\hat{u}_v^c$ (the state after absorbing a candidate's broadcast).

We are candid about the limitations of this device. The ideal understanding is not derived from deeper primitives; it is a reduced-form anchor that black-boxes the relationship between the cognitive framework and vote choice. A richer treatment would model how the structure of a voter's understanding—which representations she accesses, which meanings she assigns—gives rise to evaluative dispositions over candidate broadcasts. That is where future work belongs: in the microfoundation of the link between semantic cognition and political preference. The ideal understanding is the simplest object that makes the link precise enough to produce the Downsification results below, and we adopt it in that spirit.

Unlike a spatial ideal point, u_v^* is a structured object in the understanding lattice, and different voters' ideals may be incomparable: each carries semantic content the other lacks, with neither richer than the other. This is the formal expression of the observation that voters do not simply disagree about where to locate on a shared axis; they organize the political domain through different conceptual systems that may not admit a common ranking (cf. Converse, 1964; Lakoff, 1996). The incomparability of voter ideals, not their distance from one another, is what drives the results of Section 5.

Remark 4.2 (Alternative individual-level specifications)**.** *The ideal understanding is one of several ways to anchor the individual voter's relationship to the election. Three alternatives deserve mention. (i)* No anchor*: dispense* *with u_v^* entirely and define the voter's evaluation of a broadcast as $\mathcal{W}(\hat{u}_v^c)$—a function of the post-campaign understanding alone. This eliminates the need for a voter-specific anchor but makes the voter's evaluation independent of what she cared about before the campaign; every voter who ends up in the same post-campaign state evaluates the candidate identically, regardless of where she started. (ii)* Preference function over understandings*: replace u_v^* with a complete preorder $\succsim_v$ on $\mathscr{U}_{\mathcal{R}}$, allowing each voter to rank understandings without an anchor point. This is the most general specification but provides no structure for the lattice-theoretic results to exploit; the Downsification theorem requires knowing what the voter wants* relative to *a reference point, not just that she can rank alternatives. (iii)* Ideal derived from initial understanding*: set $u_v^* = u_v^0$, so the voter evaluates broadcasts against her pre-campaign state. This is defensible when the pre-campaign state is the voter's settled long-run understanding, but it ties the ideal to a temporal snapshot and makes it sensitive to the order in which campaigns are processed.*

The ideal understanding avoids the limitations of each alternative at the cost of being a reduced-form primitive. Whether the cost is worth paying depends on the results it produces; those results begin with how campaigns change what voters think.

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 4.3 (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 4.4. *An* electoral welfare function *is a voter-independent map* $\mathcal{W} : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$ *satisfying*

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

The welfare function is voter-independent: all voters evaluate understandings by the same $\mathcal{W}$. Heterogeneity in voter welfare arises entirely from heterogeneity in ideal understandings u_v^*, not from differences in how voters value semantic content. Voter-specific welfare functions $\mathcal{W}_v$ would complicate the aggregation without changing the structural results.

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.[9]

[9]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.

By strict monotonicity of $\mathcal{W}$, richer common ground translates directly into higher welfare.[10] The meet exists for all pairs by the complete lattice theorem (Theorem 2.9).

Why the meet and not some other aggregator? Two desiderata select it uniquely. First, the voter's welfare should depend only on what the broadcast and the voter's ideal *share*—not on either taken alone. This rules out $\mathcal{W}(\hat{u}_v^c)$, which evaluates the broadcast in isolation and ignores what the voter cares about. Second, welfare should be monotone in the shared content: more common ground, more welfare. This rules out distance-based welfare functions, which would require a metric structure on $\mathscr{U}_{\mathcal{R}}$ that the framework does not impose.

Lemma 4.5. *Let $f : \mathscr{U}_{\mathcal{R}} \times \mathscr{U}_{\mathcal{R}} \to \mathscr{U}_{\mathcal{R}}$ satisfy:*

(a) $f(u, u') \preceq_{\mathscr{U}} u$ and $f(u, u') \preceq_{\mathscr{U}} u'$ for all u, u';

(b) $f(u, u') \succeq_{\mathscr{U}} v$ for every v satisfying $v \preceq_{\mathscr{U}} u$ and $v \preceq_{\mathscr{U}} u'$.

Then $f(u, u') = u \wedge u'$.

Condition (a) says $f(u, u')$ is a lower bound of $\{u, u'\}$; condition (b) says it is the greatest such lower bound. The greatest lower bound is the meet, by definition. Any voter welfare function satisfying the two desiderata above must therefore be a function of $\hat{u}_v^c \wedge u_v^*$: the voter credits exactly the semantic content that both the campaign and the ideal contain, and no more.

Definition 4.6. *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 $\mathcal{V}_c\big((x_c)_{c \in C}\big) = \big|\{v \in V : c^*(v, \cdot) = c\}\big|/|V|$.*

Remark 4.7 (Alternative welfare specifications)**.** *The welfare function $\mathcal{W}$ is voter-independent and strictly monotone; the meet provides the argument.*

[10] The closest empirical/operational precedent for this kind of welfare evaluation is the *correct-voting* criterion of Lau and Redlawsk (2006, ch. 4): the proportion of voters who, under the constraints of an actual campaign, vote for the candidate they would have chosen with full information. The present construction is structural rather than counterfactual—it measures the lattice meet of the broadcast actually received and the voter's ideal understanding, without simulating what the voter would do under full information—but the underlying normative move is the same: welfare is evaluated against the voter's ideal, not against the broadcast in isolation.

Three alternative specifications are worth noting. (i) Voter-specific welfare*: replace $\mathcal{W}$ with $\mathcal{W}_v$, allowing different voters to weight semantic content differently. The lattice results (Theorem 2.9 and lemma 4.5) are unaffected—meet and join are defined by the lattice order, not the welfare function—but the aggregation across voters becomes harder and comparative statics on $\mathcal{W}$ are lost. (ii)* Distance-based welfare*: define $w_v(c) = -d(\hat{u}_v^c, u_v^*)$ for some metric d on $\mathscr{U}_{\mathcal{R}}$, recovering the spatial model's distance-from-ideal structure. This requires a metric that the lattice does not provide; the meet-based specification avoids this by using the lattice order directly. (iii)* Non-monotone welfare*: allow $\mathcal{W}$ to be non-monotone, so that richer common ground does not necessarily produce higher welfare (e.g., a voter who is* threatened *by understanding too much). The Downsification results require strict monotonicity; without it, the connection between semantic structure and coalition structure breaks down.*

With the welfare and voting machinery in place, the question becomes what candidates do with it.

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 4.8. *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 familiar claim that the Downs model "works" when there is a single salient dimension of competition is ubiquitous in the literature—but what, precisely, does "works" mean? The standard answer invokes a pre-given ideological axis along which voters and candidates are arranged. The semantic framework gives a different and more precise answer: "works" means that the coalition structure of the electorate admits a total order that faithfully summarizes all pairwise majority comparisons. This total order is not a property of the issue space or the analyst's model; it is a property of the *election*—of the specific pattern of meanings, credibility structures, and trigger behavior that the electorate and the candidates bring to it.

The primitive is the semantic election. A total order $\leq_X$ on X is a *valid organizational scheme* for an election if and only 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*

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

where x defeats x' if a strict majority of voters have $w_v(x, \cdot) > w_v(x', \cdot)$. The pairwise comparison requires evaluating each voter's welfare under hypothetical two-candidate sub-elections; this is well-defined under the standing assumption of *candidate-independent credibility*: $\alpha_{vc} = \alpha_v$ for all c, so that voter welfare depends only on the broadcast $\phi(x_c)$, not on candidate identity. This condition is automatically satisfied under the Downsification conditions. Under candidate-independent credibility, voter welfare reduces to a function of the platform alone: $w_v(x) = \mathcal{W}(\phi(x) \wedge u_v^*)$, suppressing the candidate index and the campaign profile from the notation of definition 4.6.

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 (Figure 2). 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

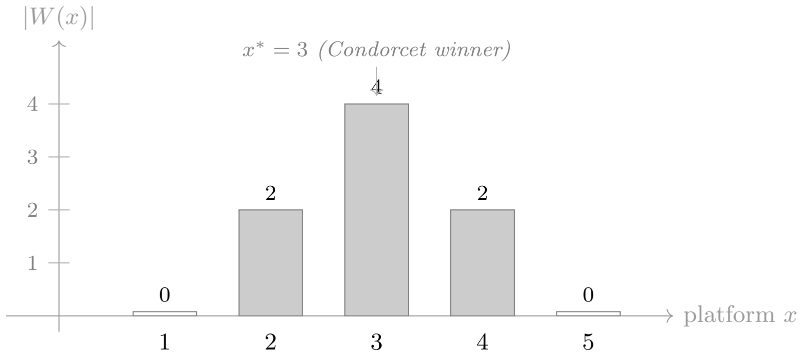

Figure 2: *The coalition mountain condition illustrated. Each bar shows* $|W(x)|$*, the number of platforms that* x *defeats by strict majority. The mountain shape—rising to a peak at the Condorcet winner* $x^* = 3$ *and declining on either side—is the behavioral signature of a recoverable organizational scheme. When this shape exists, the ordering* $1 <_X 2 <_X 3 <_X 4 <_X 5$ *is the recovered policy dimension.*

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.

The preceding subsection argued that "policy as organizational scheme" means a total order $\leq_X$ recovered from, not imposed on, the coalition structure of the election. The question is: what condition on the coalition structure makes such a recovery possible? The answer is a shape condition on the majority-wins function $|W(\cdot)|$.

Definition 4.9. *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).*

The CMC says that the platforms can be arranged on a line such that the number of alternatives each platform defeats rises to a peak at x^* and falls off on either side—a mountain shape. The peak platform is the Condorcet winner; the slopes of the mountain measure how quickly majority support erodes as platforms move away from it. When the mountain exists, $\leq_X$ is the organizational scheme and x^* is the median. When it does not, no single dimension organizes the competition.

A symmetric equilibrium admits a classical median-voter interpretation when the coalition structure has this shape. We make this precise.

Definition 4.10. *A symmetric equilibrium x^* is* Downsifiable *if there exists a total order $\leq_X$ on X such that (i) x^* is the Condorcet winner (defeats every other platform by majority), (ii) voter welfare functions $w_v(x) = \mathcal{W}(\phi(x) \wedge u_v^*)$ are single-peaked on $\leq_X$, and (iii) x^* is the majority median under $\leq_X$.*

Downsifiability asks whether the equilibrium can be *read as* a median-voter outcome: whether there exists some ordering of platforms under which the equilibrium platform is the majority median and voter preferences are single-peaked. The ordering is not given; it is recovered.

Theorem 4.11.

(i) *If x^* is Downsifiable then the CMC holds at x^*.*

(ii) *If the CMC holds at x^* then $\phi(x^*)$ is maximal among $\{\phi(x) : x \in X\}$: no platform's broadcast image is strictly richer in $\preceq_{\mathscr{U}}$ than the winning platform's.*

The first part is the diagnostic: whenever an election admits a Downsian reading, the coalition mountain must be there to support it. No mountain, no recoverable dimension—this is the clean negative test that motivates the failure-mode analysis of Section 5.

The second part is more surprising. The CMC does not just constrain the coalition structure; it constrains the *winner*. The winning platform must be *semantically maximal*: it must broadcast an understanding at least as rich as every other platform's. The intuition is stark—if a rival platform offered strictly richer content, the lattice would guarantee every voter at least as much common ground with the rival as with the incumbent, and x^* could not assemble a majority against it. The Condorcet winner in a semantic election cannot be a thin compromise that splits the difference between richer alternatives. It has to be the most comprehensive message on offer.

These two results bound the problem from different directions but do not close it: the CMC is necessary for Downsifiability, and it disciplines the winner, but it does not by itself guarantee that single-peaked preferences exist on any total order. The sufficient side comes from the four Downsification conditions developed in Section 4.7 below: when all four hold, the CMC is guaranteed and the election is fully Downsian (Proposition 4.12). Between the CMC and the four conditions lies a space of elections whose Downsifiability depends on the lattice geometry rather than on parameter conditions. The immigration example lives in this space—Downsifiable despite every condition failing for at least one voter—and the question of what distinguishes the Downsifiable elections in this gap from the non-Downsifiable ones remains open.[11]

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.[12]

[11]A natural candidate for resolving this gap is a *policy spine*: a totally ordered subset $S \subseteq X$ containing x^* on which all voter welfare functions are weakly single-peaked, with the remaining platforms fibered over S via a welfare-compatible projection $\pi : X \to S$. The spine would be the recovered policy dimension in the geometric sense—a one-dimensional path through the semantic landscape along which Downsian behavior holds, with off-spine platforms representing semantic variation the dimension does not organize. In the classical Downsian case the spine is all of X and the fibers are trivial; in the semantic setting the spine may be a strict subset and the fibers carry the content the organizational scheme compresses away. Whether the CMC guarantees the existence of such a spine—and whether its existence characterizes the Downsifiable elections in the gap—is an open question that connects the present framework to the geometric theory of preference aggregation.

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

4.7 Downsification conditions.

The CMC is a behavioral criterion—a shape condition on the coalition structure. 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. *Full policy credibility*: $\alpha_{vc}([r]) = M_{\text{policy}}$ and epistemic expansion is maximal ($e'_v = e_v \vee e_c$) 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*: $\mathcal{W}(\hat{u}^c_v \wedge u^*_v)$ is single-peaked in $\pi(\phi(x_c))$ for every voter v, with peak at $\pi(u^*_v)$.

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

The weaker condition $\alpha_{vc}([r]) \subseteq M_{\text{policy}}$ (that all absorbed content is policy content, without requiring all policy content to be absorbed) is insufficient: voters could filter out policy content from untrusted sources, breaking the connection between platform and voter response.

The four Downsification conditions formalize what Stokes (1963, pp. 376–377) called the case of *strong ideological focus*: a single stable ordered dimension is salient, voters and parties share a common frame of reference, and political controversy is organized around position-issues rather than valence-issues. Their failure characterizes what Stokes called the case of *weak ideological focus*—political controversy diffused over many issue concerns, perceived in different ways by different actors. The conditions identify the strong-focus regime precisely; the chaos bound, polysemy, and equilibrium-differentiation results of Section 5 characterize what electoral competition looks like in the weak-focus regime.

Condition 3 is genuinely independent of the other three. Conditions 1, 2, and 4 together ensure that the model *looks like* a spatial model: policy content transmits fully (Condition 1), a single dimension is available to organize it (Condition 2), and the update is always forward—no trigger-driven reversion (Condition 4). But even when all three hold, the voter's

---

a fixed domain; the present framework asks when a domain with single-peaked structure exists at all—but the structural duality is exact.

welfare $\mathcal{W}(\hat{u}_v^c \wedge u_v^*)$ need not be single-peaked in $\pi(\phi(x_c))$. The lattice meet does not automatically produce a welfare landscape that peaks at the voter's projected ideal and declines monotonically away from it; the meet can interact with the structure of M_{policy} in ways that create local welfare plateaus or secondary peaks. Condition 3 ensures the model *behaves like* a spatial model—voter welfare peaks at a unique projected ideal and declines away from it—and this behavioral property is what Black's theorem requires. Under additional algebraic requirements on the lattice—such as distributivity of $\mathscr{U}_{\mathcal{R}}$ and π being a meet-preserving homomorphism—Condition 3 may follow from the lattice geometry, but the relationship depends on the welfare function $\mathcal{W}$ in ways that are not straightforward. We state Condition 3 directly rather than pursuing these structural refinements.

Proposition 4.12. *If a semantic election satisfies the four Downsification conditions and $\pi \circ \phi$ is injective (distinct platforms project to distinct policy positions), then:*

(i) the CMC holds at the median platform x_{med}, defined by $\pi(\phi(x_{\text{med}})) = p^*_{\text{med}}$ where p^*_{med} is the median of $\{\pi(u_v^*)\}_{v \in V}$;

(ii) the symmetric profile $(x_{\text{med}}, x_{\text{med}})$ is a Nash equilibrium in the two-candidate case;

(iii) the election is Downsifiable, with recovered ordering $\leq_X$ induced by $\pi \circ \phi$.

The proposition closes the circle: if the semantic parameters of the election are structured in the right way—full policy credibility, a single organizing dimension, single-peaked welfare, and no trigger-driven reversion—then the classical median-voter result obtains as a theorem of the semantic framework, not as an assumption imposed on it. The Downs model is recovered, not presupposed. The injectivity condition ($\pi \circ \phi$ injective) ensures that the recovered ordering $\leq_X$ separates platforms: distinct platforms correspond to distinct policy positions. Without it, multiple platforms could project to the same policy position, and the ordering would fail to distinguish them.

The question, then, is whether the conditions hold.

Rarity of Downsification. The Downsification conditions are individually demanding and jointly restrictive. The immigration example makes this concrete: each of the four conditions fails for at least one voter, and the failures are independent—fixing one does not address the others.

Full policy credibility requires every voter's credibility filter to pass all policy content: $\alpha_{vc}([r]) = M_{\text{policy}}$ for all v, c, $[r]$. The affect voter's filter is $\alpha_C([r]) = \{G, B\}$—she absorbs evaluative content and rejects all factual or policy content entirely. Condition 1 fails for her not because she is uninformed but because her mode of political engagement is evaluative rather than factual. This is not exotic. Iyengar et al. (2019) document that affective polarization—partisan animosity independent of any change in policy positions—has grown dramatically in the American electorate, which in the present framework corresponds to credibility filters increasingly concentrated on $M_{\text{affect}} \cup M_{\text{social}}$. When the typical voter's filter admits affect and filters policy, Condition 1 is the first to fail and the warm glow case of Section 3.5 is the generic mode of campaign reception.

Single dimensionality requires a total order on M_{policy} and a projection π mapping every understanding to a single real-valued policy position. In the immigration example, $M = \{T, F, G, B\}$ has two orthogonal dimensions: factual (T/F) and evaluative (G/B). Voter A's ideal is organized around enforcement-works-and-is-good ($\{T, G\}$ at r_1); voter B's around integration-works-and-is-good ($\{T, G\}$ at r_2). No single projection $\pi : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$ captures both the factual and evaluative trade-offs that distinguish these voters. The incomparability of u_A^* and u_B^* in the understanding lattice—neither is richer than the other; each carries semantic content the other lacks—is the lattice-theoretic manifestation of the multidimensionality that destroys Condition 2. This is Converse's (1964) finding given formal content: the mass public lacks ideological constraint not because voters are ignorant but because the meaning space through which they process political information is genuinely multidimensional, and no single axis organizes the trade-offs that voters face.

Trigger-free requires $\tau_v = \infty$ for all voters, ensuring that every encounter produces forward absorption rather than defensive contraction. In the immigration example, voter A encountering an integration understanding has $|\mathsf{Conf}(u_A, u_L)| = 3$: every representation class is in conflict. With any finite threshold $\tau_A \leq 2$, the trigger fires and A reverts to the baseline $\beta = (\{T\}, \emptyset, \{T\})$—a strict impoverishment. The sophistication paradox (Lemma 3.8) ensures this is not a pathological case: the richer the voter's prior understanding, the more meanings are at risk and the more readily the trigger fires. Taber and Lodge's (2006) finding that politically sophisticated subjects show the strongest disconfirmation effects means that finite thresholds are the empirical norm, not the exception. Condition 4 thus requires voters to absorb without resistance—a psychological impossibility for anyone with substantive prior commitments on the issue at hand.

The immigration example is instructive precisely because Condition 3 does *not* fail. Label the three platforms x_R (a pure enforcement broadcast), x_C (a factual centrist broadcast), and x_L (a full integration broadcast). On the ordering $x_R <_X x_C <_X x_L$, all three voters have single-peaked welfare: each voter's welfare declines monotonically away from their peak platform. But this is accidental. The welfare function $\mathcal{W} = \sum_i |\mu(r_i)|$ is too coarse and the meaning space $M = \{T, F, G, B\}$ too small to generate the cross-cutting trade-offs that would break single-peakedness; a richer M or a finer $\mathcal{W}$ would produce multi-peaked preferences. The example thus sharpens the distinction between the conditions and the property they guarantee. The Downsification conditions ensure single-peakedness *structurally*—via the projection π and the lattice properties of the meet—while the immigration election achieves it *accidentally*, by coincidence of scale. The conditions are about structural guarantee, not about what happens to hold in a particular small election.

Indeed, the Downsification conditions are sufficient for the CMC but not necessary. The immigration election is Downsifiable—x_L is a Condorcet winner and all voters have single-peaked welfare on $x_R <_X x_C <_X x_L$—even though every one of the four conditions fails for at least one voter. The conditions identify when the CMC must hold as a consequence of the semantic structure; they do not characterize the full set of elections in which single-peakedness happens to obtain. When the conditions fail, Downsification may still occur by accident—but the analyst has no structural reason to expect it and no basis for comparative statics that depend on it.

Downs's model is correct on its own assumptions. The present framework characterizes those assumptions as a knife-edge—and provides a formal account of what happens when the knife-edge is not reached. Acharya, Blackwell and Sen (2018) offer a complementary perspective from cognitive dissonance theory: in their model, voters facing a multidimensional policy space and two fixed parties adjust their ideal points to minimize the psychological cost of supporting a party whose platform differs from their initial preferences, producing one-dimensional partisanship as an endogenous outcome. Their mechanism is psychological—dissonance reduction drives preference collapse—while the Downsification conditions are structural: the lattice must have the right properties for a single dimension to organize competition. The two accounts make different predictions when the semantic dimension exceeds one. In the ABS framework, preferences collapse toward party platforms regardless of the underlying dimensionality; in the present framework, the full lattice structure persists, and the results of Section 5—the chaos bound, strategic polysemy, and majoritarian convergence—describe what competition looks like in the absence of collapse.

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 Downsification conditions are demanding and jointly restrictive; the generic case, as the immigration example demonstrated, is that they fail. This section asks what structure remains when they do.

The answer is: more than the spatial literature would predict. Classical results in positive political theory establish that when the core is empty—when no single alternative beats all others by majority—the majority-preference orbit is essentially unconstrained, potentially spanning the entire alternative space (McKelvey, 1976, 1979). The semantic framework produces three results that go beyond this binary. First, 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. Second, even when majority-preference cycling occurs, the welfare-relevant orbit is confined to a *semantic interval*—a lattice-theoretic bound that spatial theory cannot supply. Third, when the electorate fragments into credibility types with incomparable ideals, the equilibrium broadcast is polysemic (heard differently by each type) and competition converges toward the majority bloc's ideal, systematically excluding the minority. The immigration example's three voters—enforcement, integration, and affect—illustrate all three phenomena.

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.

Proposition 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.*

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* $i \neq j$ *and all* $[r]$*, no agent can absorb anything from any other; every profile is already a fixed point, so* $\boldsymbol{\mu}^* = \top_{M_e^N}$, $\boldsymbol{\mu}^\dagger = \bot_{M_e^N}$, *and the spread is maximal. The spread thus measures how much semantic potential the credibility structure permits to be realized.*

The immigration example. Consider the three immigration voters at a common epistemic level $e_\top$, with triggers deactivated. Voters A and B have factual credibility filters $\alpha_A = \alpha_B : [r] \mapsto \{T, F\}$; voter C has an evaluative filter $\alpha_C : [r] \mapsto \{G, B\}$. The two filter types are disjoint: $\alpha_A([r]) \cap \alpha_C([r]) = \emptyset$ for all $[r]$. In round one, A absorbs new factual content from B's semantics—T at r_2 (where A initially held only $\{F, B\}$) passes through A's factual filter—and B symmetrically absorbs F at r_2 from A. Voter C, whose evaluative filter rejects all factual content, absorbs nothing from either A or B at any class where they offer only T or F; she does absorb evaluative content—G at r_1 from A, B at r_2 from A, B at r_3 from A, and G at r_3 from B—expanding her evaluative vocabulary. In round two, no new content passes any filter: every meaning that could be absorbed has already been absorbed. The fixed point is reached in two rounds.

At the fixed point, voters A and B have pooled their factual content (both now hold T and F at r_2), and voter C has pooled the evaluative content available from A and B. But the two groups have exchanged nothing with each other: the disjoint credibility structure creates two informationally closed communities within a single electorate, each converging internally while remaining semantically isolated from the other. The communication equilibrium is stable—no agent can absorb further content—but its stability coexists with a deep informational partition.

Active triggers and cascading impoverishment. The communication equilibrium 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.

The consequence is not merely that convergence may fail; it is that sustained contact can produce *cascading impoverishment*. 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 conflict profile to j: if $|\mathsf{Conf}(u_j, \beta(u_i, u_j))| > \tau$, j's trigger fires in turn, and j reverts to $\beta(u_j, \beta(u_i, u_j)) \prec_{\mathscr{U}} u_j$. Each encounter leaves both agents representationally poorer than before, and since the understanding lattice is finite, the drift terminates only when both have contracted to a point where their remaining semantic content no longer generates enough conflict to fire the trigger.

In the immigration example, voter A encountering voter B has $|\mathsf{Conf}(u_A, u_B)| = 3$—every representation class is in conflict (Section 3.1). If $\tau_A = 2$, A reverts to $\beta(u_A, u_B) = (\{T\}, \emptyset, \{T\})$: bare factual acknowledgments with all evaluative content stripped. Voter B now encounters this impoverished understanding; the conflict set is smaller but B's own threshold may still be exceeded, producing a further contraction. The cascade is Festinger's (1957) dissonance-reduction mechanism iterated: each trigger-firing removes dissonant cognitions, but the reduced understanding creates a new conflict profile that can fire the other agent's trigger. The terminal state is mutual representational poverty—a stable point not because the agents have reached agreement but because neither retains enough semantic content to generate conflict.

The communication equilibrium is a pre-strategic object: it describes the stable state of mutual comprehension, not the outcome of electoral competition. The question of whether competition itself reaches a stable point—whether a symmetric Nash equilibrium exists—is distinct. In the two-candidate case, the connection is direct: a symmetric equilibrium at x^* requires that no deviation to $x \neq x^*$ win a strict majority, which is precisely the condition that x^* be a Condorcet winner in the majority tournament on X. Conversely, any Condorcet winner constitutes a symmetric equilibrium. The symmetric equilibrium and the Condorcet winner are the same object.

This identification clarifies the stakes. When the Downsification conditions hold, Black's theorem guarantees a Condorcet winner at the median platform (Proposition 4.12), and with it a symmetric equilibrium. When they fail, the existence of a Condorcet winner is not structurally guaranteed—and the remainder of this section examines what structure the semantic framework provides in that case.

5.2 The chaos bound.

Standing assumption. Throughout the remainder of this section, $|V|$ is odd, so every pairwise majority comparison has a strict winner.

When the CMC fails, the failure 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—since in the two-candidate case the symmetric equilibrium is the Condorcet winner—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 5.4. *The* semantic dimension *of a voter population V is the width of the poset $\left(\{u_v^*\}_{v \in V}, \preceq_{\mathscr{U}}\right)$—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 5.5. *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.

Lemma 5.6. *No strictly monotone lattice homomorphism* $\pi : (\mathscr{U}_{\mathcal{R}}, \preceq_{\mathscr{U}}) \to (\mathbb{R}, \leq)$ *exists when the voter ideal poset* $(\{u^*_v\}, \preceq_{\mathscr{U}})$ *has width exceeding* 1*.*

The proof is short. A lattice homomorphism preserves meets: $\pi(u \wedge u') = \min(\pi(u), \pi(u'))$. If $u^*_1 \parallel u^*_2$, then $u^*_1 \wedge u^*_2 \prec_{\mathscr{U}} u^*_1$ (the meet is strictly below both elements of the antichain). Strict monotonicity requires $\pi(u^*_1 \wedge u^*_2) < \pi(u^*_1)$ and $\pi(u^*_1 \wedge u^*_2) < \pi(u^*_2)$. But the homomorphism condition gives $\pi(u^*_1 \wedge u^*_2) = \min(\pi(u^*_1), \pi(u^*_2))$; without loss of generality $\pi(u^*_1) \leq \pi(u^*_2)$, so $\pi(u^*_1 \wedge u^*_2) = \pi(u^*_1)$—contradicting strict monotonicity.

The lemma identifies the structural mechanism: the only natural route from semantic structure to Downsification is a lattice homomorphism π that translates the lattice order into a linear order on which welfare is single-peaked. When voter ideals are comparable, such a homomorphism exists (any order-preserving injection into $\mathbb{R}$ works on a chain). When voter ideals are incomparable, it does not—the lattice contains structural information that no single real dimension can faithfully represent.[13]

Proposition 5.7. *If the semantic dimension of* V *exceeds* 1*, then Condition 3 (single-peaked welfare) cannot be derived from the lattice structure of* $\mathscr{U}_{\mathcal{R}}$*:*

[13] Lemma 5.6 shows that the structural path to Downsification—a meet-preserving projection—is blocked by incomparability. This does not rule out the logical possibility that voter welfare is single-peaked on some total order $\leq_X$ under a projection π that is *not* a lattice homomorphism. Such a projection would organize voter welfare without respecting the lattice structure, and single-peakedness would hold by accident rather than by structural guarantee—exactly as in the immigration example, where Downsification obtains despite condition failure (Section 4.7). The proposition claims that the Downsification conditions cannot all hold, not that Downsification itself is impossible: the conditions require a projection that is structurally well-behaved, and Lemma 5.6 shows that no such projection exists when voter ideals are incomparable.

the only structural mechanism that would guarantee it—a strictly monotone lattice homomorphism $\pi : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$—does not exist (Lemma 5.6). The CMC is therefore not structurally guaranteed when semantic dimension exceeds 1*.*

The proposition makes a precise claim about derivability, not impossibility. Condition 2 requires a projection $\pi : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$; Condition 3 requires voter welfare to be single-peaked under it. As noted in the discussion of Condition 3 independence (Section 4.7), the sufficient structural condition for Condition 3 to follow from the lattice geometry is that π preserve meets. Lemma 5.6 shows that no meet-preserving, strictly monotone projection exists when voter ideals are incomparable. Without such a projection, Condition 3 can still hold—the immigration example is Downsifiable despite having semantic dimension 2 (Section 4.7)—but it holds by coincidence of scale, not by structural guarantee. The Downsification conditions identify when the organizational scheme is structurally sound; when semantic dimension exceeds 1, they may be satisfied by accident but cannot be ensured by the lattice.

Remark 5.8 (McKelvey cycling in the Euclidean benchmark)**.** *When voter welfare admits a Euclidean embedding in $\mathbb{R}^k$ (k = semantic dimension ≥ 2)—that is, when there exists an injective map $\iota : \{u_v^*\}_{v \in V} \to \mathbb{R}^k$ such that $\mathcal{W}(\phi(x) \wedge u_v^*)$ is a decreasing function of $\|\iota(\phi(x)) - \iota(u_v^*)\|$—and no total median exists in the distribution of embedded voter ideals, McKelvey's (1976) theorem yields path-connectivity of the majority orbit: any platform is reachable from any other by a finite majority chain. Schofield (1978) established the local cycling condition (the null dual), Schofield (1983) showed global instability is generic when the spatial dimension is at least* 2*, and McKelvey and Schofield (1987) characterized the necessary and sufficient symmetry conditions for a core point to exist. The chaos bound (Theorem 5.9) confines even this global cycling to the semantic interval—a structural constraint that the spatial framework cannot supply.*

Theorem 5.9. *When the semantic dimension of V exceeds* 1*, the welfare-relevant McKelvey majority-preference orbit is confined to the* semantic interval

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

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

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.

The significance of this bound becomes clear against the backdrop of the classical chaos results. McKelvey (1976) showed that under Euclidean preferences with no Condorcet winner, the majority-preference orbit is path-connected and can reach any platform from any other; McKelvey (1979) extended this to general smooth preferences. Austen-Smith and Banks (1999) establish the sharpest form of the classical result (Theorems 6.4–6.5): when the core is empty, the top cycle set equals the entire alternative space. There is, in spatial theory, no intermediate constraint between the core (which bounds everything) and the Pareto set (which in high-dimensional settings approaches X itself). The spatial framework offers a binary: either there is a core and competition is well-behaved, or there is no core and majority dynamics are essentially unconstrained.

The semantic interval provides the missing intermediate bound. It exists regardless of whether a semantic core exists; it is strictly smaller than $\mathscr{U}_{\mathcal{R}}$ whenever voter ideals do not span the full lattice (the generic case); and it is determined entirely by the distribution of voter ideal understandings. The semantic interval is, in fact, the lattice-theoretic analog of the Pareto set: the smallest interval in $\preceq_{\mathscr{U}}$ containing all voter ideals, just as the Pareto set for Euclidean preferences is the convex hull of voter ideal points in $\mathbb{R}^k$. But unlike the convex hull, which approaches the whole space as voters proliferate in high dimensions, the lattice interval is a structural bound that depends only on the extreme voter ideals—the meet and join—and cannot grow without bound.

In the immigration example, the three voter ideals yield a concrete interval. The three-way meet $\bigwedge\{u_A^*, u_B^*, u_C^*\}$ has empty semantics at every representation class: $\{T,G\} \cap \{T,B\} \cap \{B\} = \emptyset$ at r_1, $\{F,B\} \cap \{T,G\} \cap \{G\} = \emptyset$ at r_2, and $\{T,B\} \cap \{T,G\} \cap \emptyset = \emptyset$ at r_3. The three voters share zero semantic common ground. The three-way join is $(\{T,G,B\}, \{T,F,G,B\}, \{T,G,B\})$: the pooled semantic universe of all three voters, nearly the full meaning space (missing only F at r_1 and r_3, because no voter assigns "false" to enforcement or economic impact). The semantic interval runs from the vacuous understanding $u_\perp$ to this near-maximal join. Whatever majority-preference cycling occurs among immigration platforms—and with semantic dimension 2, it can be substantial—it is confined to understandings between these bounds. Platforms outside the interval (those assigning meanings that no voter values or those too impoverished to register with any voter's ideal) are welfare-irrelevant and do not participate in the majority tournament.

The strength of the chaos bound is its unconditional nature. Schofield (1983) shows that when the spatial dimension reaches a critical threshold ($w(n) = 2$ for odd n), majority cycling is generically dense: almost every preference profile admits cycles through almost every platform. The immigration example's semantic dimension of 2 places it squarely in this regime. The chaos bound does not prevent cycling—it confines it. And unlike Schofield's dimensionality results, which require the smooth manifold structure of Euclidean space, the chaos bound relies only on the complete lattice structure of $\mathscr{U}_{\mathcal{R}}$ and holds for arbitrary voter ideal distributions.[14]

Corollary 5.10. *The majority-preference orbit is path-connected within the semantic interval: for any two platforms $\phi(x), \phi(y) \in [\bigwedge_v u_v^*, \bigvee_v u_v^*]$ connected by a majority-preference path, there exists a majority-preference path between them that remains within the interval at every step.*

The argument is a rerouting one, and the key observation is welfare equivalence at the boundary. Consider any platform $\phi(z) \succeq_{\mathscr{U}} \bigvee_v u_v^*$—a platform above the join. For every voter v, the meet $\phi(z) \wedge u_v^* = u_v^*$ (since $\phi(z) \succeq_{\mathscr{U}} \bigvee_v u_v^* \succeq_{\mathscr{U}} u_v^*$). The same is true of the join itself: $\bigvee_v u_v^* \wedge u_v^* = u_v^*$. So $\mathcal{W}(\phi(z) \wedge u_v^*) = \mathcal{W}(\bigvee_v u_v^* \wedge u_v^*) = \mathcal{W}(u_v^*)$ for every voter v—the two platforms generate identical welfare for the entire electorate. Now suppose a majority-preference path includes a step from some interior platform $\phi(x)$ to $\phi(z)$: a strict majority of voters prefer $\phi(z)$ to $\phi(x)$. Since $\phi(z)$ and $\bigvee_v u_v^*$ are welfare-identical for every voter, that same strict majority prefers $\bigvee_v u_v^*$ to $\phi(x)$. Replace $\phi(z)$ with $\bigvee_v u_v^*$; the majority-preference step is preserved. The same reasoning applies symmetrically below the meet (where all platforms generate the same minimal welfare). The rerouted path stays within the interval and preserves strict majority preference at every step. The semantic interval $[\bigwedge_v u_v^*, \bigvee_v u_v^*]$ is defined by the lattice order: it contains every understanding u with $\bigwedge_v u_v^* \preceq_{\mathscr{U}} u \preceq_{\mathscr{U}} \bigvee_v u_v^*$. There is no "lateral escape"—the lattice order precludes platforms that are outside the interval yet incomparable with the bounds, because the interval is the full set of elements between the meet and join in the lattice.

[14] Stokes (1963, p. 376) closed his critique of the spatial model by predicting that "extending the model to the case of two or more stable, ordered dimensions will lead to results that do not have any analogues in the one-dimensional case." McKelvey's chaos theorem is one such result; Theorem 5.9 is another, identifying a structural lattice constraint on the cycling region that the spatial framework cannot supply. A different response to the same chaos literature appears in the behavioral-formal program of Kollman, Miller and Page (1992, 1998), who argue via computational simulation that boundedly rational adaptive parties cannot in practice locate the destabilizing platforms that McKelvey-style chaos requires, so that the cycling region is computationally unreachable even when it formally exists. The two responses are independent and complementary: KMP rely on agent-side computational limits, the chaos bound relies on the lattice geometry of voter ideals, and both reach the conclusion that majority dynamics are more constrained than the classical chaos results suggest.

5.3 Strategic polysemy.

The chaos bound constrains *where* competition can go when Downsification fails. Polysemy describes *how* candidates communicate with a semantically fragmented electorate: not by finding a message that means the same thing to everyone, but by crafting a single broadcast that means different things to different listeners.

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 5.11. *When credibility filters are pairwise disjoint and all triggers are deactivated ($\tau_v = \infty$ for all $v \in V$), 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.*

When triggers are active, the decomposition can fail: content in one type's channel may generate conflicts for voters in another type, coupling the optimization across types through the trigger mechanism.

Corollary 5.12. *In a two-candidate race with pairwise disjoint voter types, pairwise incomparable type ideals ($u^*_{T_i} \parallel u^*_{T_j}$ for $i \neq j$), and deactivated triggers ($\tau_v = \infty$ for all v), 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 and voter type ideals are semantically incomparable.*

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 distinction from existing models of electoral ambiguity is structural. In Shepsle's (1972) pioneering treatment, a candidate is ambiguous by offering a lottery over positions in a fixed one-dimensional policy space; risk-loving voters may prefer the ambiguous candidate to a certain moderate. In the context-dependent model of Callander and Wilson (2008), ambiguity arises because voters evaluate candidates relative to each other, developing a preference for uncertainty through comparison effects. Both accounts require an exogenous policy space and model ambiguity as a probabilistic object—a distribution over positions that voters must evaluate under risk preferences. Semantic polysemy is none of these things. The broadcast is a single understanding in $\mathscr{U}_{\mathcal{R}}$, not a lottery; its meaning is deterministic but *filter-dependent*, so that different voters absorb different components of the same message. The mechanism is not probabilistic evaluation of an uncertain position but structural decomposition of a rich understanding through heterogeneous credibility filters.

In the immigration example, the decomposition is transparent. Partition the electorate into two credibility types: $T_1 = \{A, B\}$ with factual filters $\alpha = \{T, F\}$, and $T_2 = \{C\}$ with evaluative filter $\alpha = \{G, B\}$. The filters are pairwise disjoint. A candidate's optimal broadcast simultaneously carries factual content optimized for voters A and B—policy information about the effectiveness of enforcement and integration—and evaluative content optimized for voter C—signals about whether outcomes are good or bad. Voters A and B hear a policy argument; voter C hears an evaluative appeal. The same broadcast, different meanings absorbed. This is not cynical manipulation; it is the structural consequence of addressing an electorate whose members literally hear different things. A univocal broadcast—one that means the same thing to every voter—would require a credibility structure in which some content passes all filters simultaneously. But $\{T, F\} \cap \{G, B\} = \emptyset$: no meaning element is credible to both types. Univocality is impossible; polysemy is the only option (Figure 3).

Remark 5.13 (Broadcast range constraints)**.** *The proof of Corollary 5.12 assumes that the broadcast map ϕ has enough range to deliver type-specific optimal content—that for each type T_i, there exists a platform x such that the α_i-filtered component of $\phi(x)$ maximizes common ground with $u^*_{T_i}$. When ϕ's range is constrained (e.g., candidates can transmit only a subset of $\mathscr{U}_{\mathcal{R}}$), polysemy still holds for whatever content ϕ can deliver: the decomposition by type remains valid and the independently optimized components are still generically distinct. But the type-specific components may fall short of the unconstrained optima, and the degree of polysemy—the extent to which the equilibrium broadcast exploits the disjoint filter structure—depends on how rich a semantic vocabulary the broadcast map makes available.*

5.4 Majoritarian convergence and the semantic fallback.

The chaos bound describes the *range* of competition under non-Downsifiable conditions; polysemy describes the candidate's communicative strategy. The

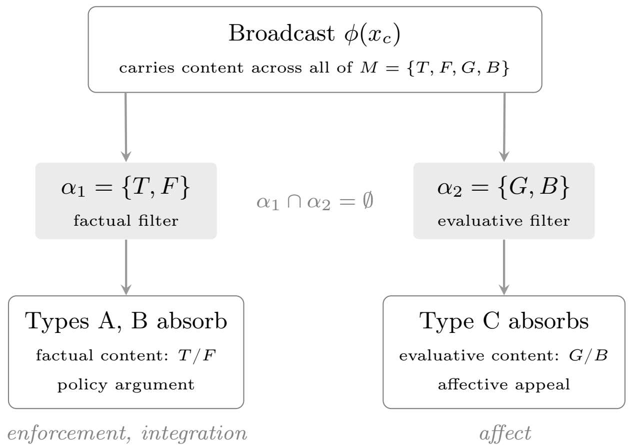

Figure 3: *The polysemy decomposition in the immigration example. A single broadcast carries content across the full meaning space M. Disjoint credibility filters split the reception: factual types (A, B) absorb only $\{T, F\}$ content; the evaluative type (C) absorbs only $\{G, B\}$ content. The candidate optimizes each component independently (Proposition 5.11), and the resulting broadcast is polysemic—the same message, heard differently by each type.*

remaining question is what the equilibrium outcome looks like when voter ideals are incomparable: who benefits, who is excluded, and what structural forces drive the answer.

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. Define the 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. These need not be disjoint: voters whose ideals lie below $u_A^* \wedge u_B^*$ belong to both. The proposition below requires only that $n_A > |V|/2$. 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 5.14. *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.*

Corollary 5.15. *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 5.14):*

competition forces both candidates toward the majority bloc's ideal. Since u_A^* and u_B^* are incomparable, any symmetric equilibrium with $\phi(x^*) \not\preceq_{\mathscr{U}} u_A^* \vee u_B^*$ satisfies $\phi(x^*) \not\preceq_{\mathscr{U}} u_B^*$, 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 $\preceq_{\mathscr{U}}$.[15,16]

In the immigration example, the antichain is $\{u_A^*, u_B^*\}$ and voter C's ideal satisfies $u_C^* \preceq_{\mathscr{U}} u_B^*$ (C's evaluative-only content is a subset of B's richer understanding at every representation class). The blocs are $V_B = \{B, C\}$ (majority, since $n_B = 2 > 3/2$) and $V_A = \{A\}$ (minority). By Proposition 5.14, every

[15] Existing formal accounts of partisan polarization typically invoke either spatial forces (primary constraints, policy-motivated candidates) or cognitive bias: Ortoleva 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.

[16] The result joins a lineage of formal non-convergence findings in spatial-theoretic and behavioral-formal electoral theory. Wittman (1983) showed that two candidates with both policy preferences and winning preferences will not converge to the median voter; Bendor et al. (2011, Prop. 3.4) derive non-convergence from satisficing-winners and searching-losers under retrospective adaptive dynamics; Kollman, Miller and Page (1998) show via simulation that voter distributions with strong issue-intensity asymmetries (extremist or independent strengths) produce persistent platform separation between adaptive parties even after many electoral periods. Corollary 5.15 adds a fourth, structurally distinct, mechanism: the antichain geometry of voter ideal understandings, which forces non-convergence even when candidates are purely office-motivated and the analysis is static. The four mechanisms are independent and can operate simultaneously. A more pointed contrast is with Hinich and Munger (1994), who answer Tullock's stability question by appeal to the sparseness of the ideological space (Π is small, candidates have nowhere to go); the present mechanism explains the complementary fact that the equilibrium occupation of that space is multipolar rather than centrist whenever the lattice of voter ideal understandings contains incomparable elements. The empirical signature of this antichain structure has been documented in the political-psychology literature: Sniderman and Bullock (2004) note that ordinary citizens are consistently liberal or conservative *within* policy agendas (chains in $\preceq_{\mathscr{U}}$) but not *across* them (antichains across chains), the cross-agenda incoherence that classical Converse-style measures register as nonattitudes.

symmetric equilibrium satisfies $\phi(x^*) \succeq_{\mathscr{U}} u_B^*$: competition forces both candidates toward the integration voter's ideal. Voter A—the enforcement voter, whose ideal is incomparable with u_B^*—is systematically excluded. The welfare deficit for A is bounded by $\mathcal{W}(u_B^*) - \mathcal{W}(u_A^* \wedge u_B^*) = \mathcal{W}(u_B^*) - \mathcal{W}((\{T\}, \emptyset, \{T\}))$: the gap between the integration voter's welfare and the welfare from bare factual agreement.

The lattice-theoretic prediction—that majority blocs whose ideals are incomparable with the minority's will dominate, with the minority suffering welfare loss proportional to the incomparability—has an empirical counterpart. Mondak et al. (forthcoming) find that the average partisan is part of an approximately 83% majority on the issue they choose to discuss, with partisan issue selections clustering sharply by party: Republicans and Democrats do not so much disagree about immigration as organize the political domain through different representational systems. In the semantic framework, "different issues" corresponds to incomparable ideals—each organized around representations the other does not prioritize—and the 83–17 issue-selection split is the behavioral signature of the antichain geometry that drives majoritarian convergence. The minority's difficulty explaining the other side (59.4% receiving a grade of D or below in Mondak et al.'s exercise) maps onto credibility filter heterogeneity: voters cannot absorb the other side's semantic content because their filters reject it.

The meet as semantic fallback. Even in full McKelvey chaos, $\bigwedge_{v \in V} u_v^*$ is a specific, well-defined element of $\mathscr{U}_{\mathcal{R}}$: the understanding whose semantic content is contained in every voter's ideal—the thinnest common ground. 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.

In the immigration case, the three-way meet is $(\emptyset, \emptyset, \emptyset)$—the vacuous understanding $u_\perp$. The three voters share zero semantic common ground: no meaning element appears in all three ideals at any representation class. As a campaign platform, this is the equivalent of a politician saying "immigration is complicated"—true, universally inoffensive, and substantively empty. A centrist strategy in an incomparable electorate is not a compromise position equidistant from the extremes; it is a semantic skeleton stripped of all content that any voter might object to—which, when the electorate's ideals are sufficiently diverse, means stripped of all content whatsoever. The structural explanation for the centrist's dilemma is lattice-theoretic: the "middle" of an incomparable electorate is not a moderate position but the meet of the voter ideals, and the meet of incomparable elements can be arbitrarily impoverished.

Legislative preview. The majoritarian convergence and semantic fallback results concern elections, but the lattice structure applies wherever collective decision-making encounters semantic incomparability. The framework extends naturally to multi-actor legislative bargaining, where the semantic viability of coalitions—whether members share enough representational common ground for joint action—adds a lattice constraint with no spatial analog. Section 8.1 develops this extension formally, adapting the Baron-Ferejohn sequential bargaining framework to the semantic setting and characterizing stationary equilibria under semantic viability constraints.

The results of this section—the chaos bound, strategic polysemy, and majoritarian convergence—describe the structure of semantic competition when the Downsification conditions fail. They show that the semantic framework provides substantially more structure than the classical chaos results would predict: majority-preference dynamics are bounded, candidate communication is structurally polysemic, and the equilibrium outcome is determined by the antichain geometry of voter ideals. The next section turns from these formal results to their empirical content: what the framework predicts, and how its constructs can be operationalized for empirical research.

6 Empirical predictions and operationalization.

The formal results of Sections 2–5 are structural claims about what happens when the Downsification conditions hold or fail. This section asks what those claims predict empirically and how the key constructs can be measured. Each of the three subsections pairs a theoretical prediction—derived from the formal results and distinguishable from what the spatial model would predict—with the measurement apparatus needed to test it. The section closes with a worked illustration that maps the immigration example's three voters onto observable survey profiles, showing how the framework's structural distinctions surface in data that existing instruments can collect.

6.1 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. Sophistication gives voters the cognitive resources to locate themselves and candidates on a policy dimension and to choose the closer option; the canonical finding that education and political knowledge predict ideological consistency (Converse, 1964; Zaller, 1992) is typically read as confirming this prediction.

The semantic framework generates 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 (Lemma 3.8). 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.

In the immigration example, the contrast is between voter A (rich prior: $\mu_A = (\{T, G\}, \{F, B\}, \{T, B\})$, all three representation classes carrying factual and evaluative content) and voter C (thin prior: $\mu_C = (\{B\}, \{G\}, \emptyset)$, evaluative content only, no engagement with the factual dimension). When both encounter a candidate broadcasting an integration understanding, voter A has $|\mathsf{Conf}(u_A, u_L)| = 3$—every representation class is in conflict—while voter C has little or nothing to conflict with. Voter A is the more informed, more engaged citizen; voter A is also the one whose trigger fires and who reverts to a bare factual skeleton. Voter C, with less to lose, absorbs freely. The most informed voter is the least spatially predictable.

Measuring the inversion. The epistemic stance e_v corresponds to the scope and fineness of a voter's political representations. The closest existing proxy is the ANES political knowledge battery: factual recall items (which party controls the House, which branch determines constitutionality) measure the scope of $\mathcal{R}_{e_v}$—how many political objects the voter can access—while open-ended likes/dislikes questions, scored for the number and specificity of considerations mentioned, approximate the fineness of $\sim_{e_v}$. Higher scores on both correspond to richer epistemic stances and denser semantic networks—precisely the conditions under which the sophistication lemma predicts larger conflict sets.

The trigger threshold τ_v is not directly observable. Its effects are: the interaction between prior richness (knowledge) and counter-attitudinal exposure should predict attitude entrenchment rather than updating. The estimation strategy follows Kunda (1990) and Taber and Lodge (2006): expose respondents to counter-attitudinal campaign content and measure post-exposure attitude change as a function of pre-exposure knowledge and constraint. The semantic framework predicts that high-knowledge respondents show *less* attitude movement toward the content and *more* reinforcement of prior positions—and that this effect is mediated by the density of the respondent's prior belief structure (the number of interconnected attitudes at risk of conflict), not by the extremity of their positions. The mediator distinguishes the semantic prediction from a simple "extremists don't move" story: it is not position extremity that drives entrenchment but the richness of the semantic network that the position is embedded in.

6.2 Polysemic campaigns and credibility heterogeneity.

Spatial theory predicts that candidates communicate coherent policy positions, with strategic ambiguity arising only under specific conditions—multidimensional preferences, risk-loving voters (Shepsle, 1972), or context-dependent evaluation (Callander and Wilson, 2008). The polysemy result (Corollary 5.12) 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 of communication.

The mechanism is specific. When voter types have pairwise disjoint credibility filters, the candidate's vote-share optimization separates by type (Proposition 5.11): 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 not ambiguity as a lottery over positions; it is a single broadcast with deterministic but filter-dependent meaning.

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.

In the immigration example, the candidate simultaneously carries factual content optimized for voters A and B (whose filters $\alpha = \{T, F\}$ pass policy information) and evaluative content optimized for voter C (whose filter $\alpha = \{G, B\}$ passes only good/bad signals). Voters A and B hear a policy argument about enforcement and integration; voter C hears an evaluative appeal. A univocal broadcast—one that means the same thing to everyone—would require content that passes all filters simultaneously, but $\{T, F\} \cap \{G, B\} = \emptyset$: no meaning element is credible to both types. Univocality is structurally impossible; polysemy is the only option.

Measuring credibility heterogeneity. The credibility filter α_{vc} maps candidate-source content to what the voter can absorb. The observable correlate is source trust: feeling thermometers toward media outlets and political figures, media consumption diaries, and selective exposure measures. When trust in source X correlates strongly within a voter type but weakly across types—when voters who trust Fox News also trust talk radio but not the New York Times, while voters who trust the Times also trust NPR but not Fox—the credibility filters are approximately disjoint, and the framework predicts polysemic equilibria. The degree of disjointness can be estimated from the cross-type correlation matrix of source trust measures: low off-diagonal entries indicate filter separation.

Flaxman, Goel and Rao (2016) document that ideological segregation in online news consumption, while not as extreme as popular accounts suggest, is real and growing: the average conservative and liberal share fewer common news sources over time. In the semantic framework, this media fragmentation is not merely a symptom of polarization but a driver of polysemic campaign strategy: as voters' information diets diverge, the overlap between their credibility filters shrinks, and the structural pressure toward polysemy increases. Iyengar et al. (2019) provide a complementary finding: affective polarization has grown substantially even as ideological polarization (measured by policy-position distance) has grown modestly or not at all. This divergence is a signature of credibility filter separation: voters are increasingly filtering on affect rather than policy, which in the semantic framework corresponds to credibility filters concentrated on $M_{\text{affect}} \cup M_{\text{social}}$—exactly the condition that makes Downsification's first condition (full policy credibility) fail and drives campaigns toward affective polysemy.

Polysemy itself can be operationalized as divergence in voter perceptions of the same candidate message. The test: expose different voter types to an identical campaign communication and measure perceived issue emphasis and evaluative content. The semantic framework predicts systematic, type-dependent divergence in what voters report hearing—not random noise or idiosyncratic misperception but structured variation that tracks the credibility filter partition. A policy-oriented respondent should report hearing policy content; an affect-oriented respondent should report hearing evaluative signals; and the candidate's message should be optimized to produce exactly this pattern.

6.3 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 (Proposition 5.14 and corollary 5.15) 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 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 prediction distinguishes the semantic account from behavioral accounts of polarization based on overconfidence (Ortoleva 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.

In the immigration example, the antichain is $\{u^*_A, u^*_B\}$ and voter C satisfies $u^*_C \preceq_{\mathscr{U}} u^*_B$, making the blocs $V_B = \{B, C\}$ (majority) and $V_A = \{A\}$ (minority). The spatial model would locate voters A and B on opposite ends of an immigration dimension and call the distance between them "polarization." The semantic model says they are not on the same dimension at all: A's ideal carries content ($\{F, B\}$ at r_2—citizenship fails and is bad) that is categorically absent from B's ideal ($\{T, G\}$ at r_2—citizenship works and is good), and vice versa. The distance between them is not a gap on a shared scale but an incomparability in the lattice. Voter C belongs to B's bloc not because C shares B's policy position but because C's evaluative content $(\{B\}, \{G\}, \emptyset)$ is a subset of B's richer understanding—a lattice relationship invisible to any spatial measure.

Measuring incomparability. The key construct to operationalize is the incomparability of voter ideal understandings—the antichain width of the poset $(\{u^*_v\}, \preceq_{\mathscr{U}})$, which the framework defines as the semantic dimension (Definition 5.4).

The closest existing empirical concept is Converse's (1964) belief-system constraint: the degree to which a respondent's attitudes on different political issues correlate with one another. High within-voter constraint means the respondent's attitudes are organized along a single dimension—a chain in $\mathscr{U}_{\mathcal{R}}$. But the relevant measure for incomparability is not within-voter constraint (which measures how organized an individual's beliefs are) but *between-voter* constraint heterogeneity: whether different voters' belief systems are organized along the *same* dimension.[17] Two highly constrained voters whose constraints

[17] Peffley and Hurwitz (1985) distinguish *vertical* constraint (abstract values producing specific policy positions, demonstrated via LISREL) from *horizontal* constraint (consistency

run along orthogonal dimensions—one organizing politics through economic redistribution, the other through cultural identity—have incomparable ideal understandings despite both exhibiting high within-voter constraint.

Mondak et al. (forthcoming) provide a direct behavioral proxy. Their finding that the average partisan is part of an approximately 83% majority on the issue they choose to discuss—and that partisan issue selections cluster sharply by party, with Democrats concentrating on health care, climate, and race while Republicans concentrate on immigration, abortion, and trade—is precisely the behavioral signature of incomparable ideals: voters' representational systems are organized around different domains. The inability to explain the other side (59.4% receiving a grade of D or below) maps onto the lattice-theoretic fact that incomparable understandings share little semantic common ground.

The estimation strategy for semantic dimension proceeds in two steps. First, cluster voters by the correlation structure of their issue attitudes: respondents whose attitudes on economic, social, and foreign-policy items load onto the same factor structure belong to the same representational type. Second, count the number of types whose factor structures are approximately orthogonal—whose organizing dimensions have low cross-type correlation. This count estimates the antichain width. An electorate with two orthogonal types has semantic dimension 2; an electorate in which all respondents share a common factor structure has semantic dimension 1 and is, in principle, Downsifiable.

The chaos bound (Theorem 5.9) generates a further observable implication: the effective issue space of an election—the set of platform positions that produce meaningfully different vote shares—should be narrower than the full policy space, and the degree of narrowing should correlate with the spread of voter ideal understandings. When semantic dimension is 1 (all ideals comparable), the effective issue space spans the full range of voter ideals; when it exceeds 1, the semantic interval confines the welfare-relevant competition to a subset of the understanding lattice, and platform positions outside this interval are electorally inert.

---

across issues at the same level, Converse's measure). Hurwitz and Peffley (1987) extend this hierarchical model to foreign policy, finding core-value-to-posture-to-policy pathways largely independent of the domestic liberal-conservative dimension. The incomparability concept here is orthogonal to vertical constraint: two voters can each have strong vertical constraint (coherent value-to-policy linkages) while organizing politics around different core values, producing the between-voter incomparability the semantic framework models.

6.4 A worked illustration: the immigration voters as survey respondents.

The immigration example's three voters—enforcement (A), integration (B), and affect (C)—can be mapped onto observable profiles that existing survey instruments could distinguish.

Voter A scores high on political knowledge batteries: she can name her representatives, place parties on immigration-specific scales, and identify factual claims about border enforcement. Her open-ended likes/dislikes responses are dense with policy-specific considerations ("I like that the candidate supports E-Verify because it reduces unauthorized employment"). Her credibility filter passes factual sources—policy analysis, CBO reports, think-tank publications—and attenuates evaluative content unless it is anchored in factual claims. On issue-attitude scales, her positions on enforcement, citizenship, and economic impact are tightly correlated: high within-voter constraint organized around an enforcement schema.

Voter B looks similar on knowledge batteries—equally high scores, equally specific open-ended responses—but the content is organized around an integration schema. She trusts similar source types (policy analysis, academic research) but draws opposite factual conclusions: paths to citizenship are viable and economically beneficial. Her within-voter constraint is equally high, but the factor structure of her attitudes is rotated: the issues that A treats as primary (enforcement mechanisms, border security) are secondary for B, and vice versa. On any standard left-right scale, A and B appear as opponents; in the lattice, they are incomparable, each carrying semantic content the other lacks.

Voter C presents a different profile entirely. She scores low on factual knowledge items and cannot place parties on policy-specific scales. Her open-ended responses are evaluative rather than factual ("I don't like how that candidate makes me feel about immigrants"). Her credibility filter passes evaluative signals—emotional appeals, group-identity cues, moral language—and rejects factual policy content. Her within-voter constraint is low by standard measures: her attitudes on different issues do not correlate along any single dimension. But her evaluative content $(\{B\}, \{G\}, \emptyset)$ is a *subset* of B's richer understanding—$u_C^* \preceq_{\mathscr{U}} u_B^*$—placing C in B's lattice bloc despite the two voters looking nothing alike on any spatial or knowledge-based measure.

The framework's prediction: B and C vote together not because they share a policy position but because C's semantic content is contained in B's. A spatial model that locates all three on a single immigration dimension would place C between A and B (low information, moderate positions) and predict C as a swing voter. The semantic model predicts the opposite: C is structurally aligned with B and will vote with B's bloc under any equilibrium that satisfies Proposition 5.14, because the majority bloc $V_B = \{B, C\}$ dominates and competition converges toward u_B^*. The spatial model's swing voter is the semantic model's most structurally committed bloc member.

6.5 Empirical scope and design requirements.

The predictions developed above are specific and empirically distinguishable from what the spatial model generates. The sophistication-proximity inversion predicts a negative interaction that the spatial model predicts to be positive; the polysemy prediction identifies structured perceptual divergence that the spatial model attributes to noise; the incomparability prediction distinguishes structural exclusion from spatial distance. Each prediction names a specific observable pattern, a specific measurement strategy, and a specific contrast with the spatial alternative.

The constructs—epistemic stance, credibility filter, ideal understanding, and incomparability—are measurable with existing survey infrastructure. But the key measurement for the incomparability prediction—between-voter heterogeneity in attitude organization—requires a design that compares respondents' belief structures across types, not just their individual positions on shared scales. Standard survey designs that measure attitudes on a fixed set of issues and locate respondents on a common dimension are structurally incapable of detecting incomparability, because they presuppose the very dimensional structure that the framework treats as an empirical question. The design requirement is to measure *which* dimensions organize each respondent's attitudes, not just *where* the respondent falls on a presupposed dimension—and then to assess whether different respondents' organizing dimensions are the same. This is a cross-respondent comparison of attitude structures, not a within-respondent measurement of attitude positions. The framework predicts that the politically relevant heterogeneity is structural—different organizing dimensions across voters—not parametric—different positions on a shared dimension—and testing this prediction requires instruments capable of distinguishing the two.

The empirical predictions developed here are consequences of the framework's lattice structure. The next section turns to the framework's *theoretical* connections: the update operator's relationship to three classical traditions in formal epistemology.

7 Connections to classical epistemic frameworks.

The update operator of Section 3 is a theory of how agents holding structured representations revise them under new information. Two classical traditions address the same question from different formal starting points: AGM belief revision (Alchourrón, Gärdenfors and Makinson, 1985), which operates on propositional belief sets and characterizes rational contraction and revision through postulates; and dynamic epistemic logic (van Benthem, 2011; Baltag and Renne, 2016), which models information change as transformations of Kripke structures and characterizes the logic of public announcements. This section maps the semantic framework onto both traditions, identifying precisely where each appears as a special case and where the semantic framework goes beyond what either can express. Together with the Bayesian special case established in Section 3.5, the three connections form a triangle: propositional (AGM), probabilistic (Bayesian), and model-theoretic (DEL) updating are all recovered under specific parameter restrictions of the same operator.

7.1 The baseline operator and AGM contraction.

AGM contraction operates on a theory K (a deductively closed set of propositions) and removes a proposition x while retaining as much of K as possible. The six basic Gärdenfors postulates (Alchourrón, Gärdenfors and Makinson, 1985) characterize this operation:

(÷1) *Closure*: $K \div x$ is a theory.

(÷2) *Inclusion*: $K \div x \subseteq K$.

(÷3) *Vacuity*: If $x \notin \mathrm{Cn}(K)$, then $K \div x = K$.

(÷4) *Success*: If $x \notin \mathrm{Cn}(\emptyset)$, then $x \notin K \div x$.

(÷5) *Extensionality*: If $\mathrm{Cn}(\{x\}) = \mathrm{Cn}(\{y\})$, then $K \div x = K \div y$.

(÷6) *Recovery*: $K \subseteq \mathrm{Cn}((K \div x) \cup \{x\})$.

The first five are uncontroversial requirements on any reasonable contraction operator; the sixth—that re-adding the contracted proposition recovers the original theory—is the most substantive and the most debated.

The baseline operator $\beta(u_i, u_j)$ of Section 3.3 is the semantic framework's contraction operation: when conflict with u_j exceeds the threshold, the agent reverts to the richest understanding below u_i that avoids the conflicting content. The translation from AGM to the semantic setting replaces the theory K with an understanding $u_i \in \mathscr{U}_{\mathcal{R}}$, the proposition x with the conflict content $\mathsf{Conf}(u_i, u_j)$, and the contraction $K \div x$ with the baseline $\beta(u_i, u_j) = \bigvee S_{ij}$.

Proposition 7.1. *The baseline operator β satisfies the semantic analogs of postulates* $(\div 1)$–$(\div 5)$ *and fails the semantic analog of* $(\div 6)$ *(Recovery).*

The proof (in Section A) verifies each postulate by direct construction. The first five hold cleanly: closure by completeness of the lattice, inclusion by construction of S_{ij}, vacuity because an agent with no conflict is already in S_{ij}, success by Proposition 3.4, and extensionality because β depends on u_j only through the conflict structure.

Recovery $(\div 6)$ fails, and the failure has three distinct sources. First, the credibility filter: if $\alpha_{ij}([r]) \subsetneq M$, re-absorption of u_j's content is partial, and meanings stripped during contraction may not pass through the filter on re-encounter. Second, the trigger: re-encountering u_j from the impoverished state $\beta(u_i, u_j)$ may generate enough conflict to fire the trigger again, producing a second contraction rather than recovery. Third, epistemic collapse: if the contraction involved coarsening of the epistemic stance (moving from e_i to a coarser e_β with merged equivalence classes), the lost distinctions cannot be recovered by adding meanings—re-encounter with u_j cannot refine a coarsened partition.

The failure of Recovery is not a deficiency of the operator; it is a feature of the political-communication domain. AGM's Recovery postulate assumes that belief change is reversible—that retracting and then re-adding a proposition leaves the agent where she started. In the semantic framework, contraction under conflict can be irreversible: the content stripped by the baseline may include epistemic distinctions (representation classes merged into the invisible class) that cannot be restored by any subsequent encounter. This is precisely the structure Festinger's (1957) theory predicts: dissonance reduction (removing dissonant cognitions) can permanently alter the cognitive landscape, and re-encountering the dissonant information does not undo the alteration because the agent's representational apparatus has changed.

The failure of Recovery places β in the family of *withdrawal* operators—contraction operators satisfying the basic postulates minus Recovery—studied in the belief-revision literature as more empirically realistic alternatives to full AGM contraction.[18]

[18]See in particular the "severe withdrawal" operators of Rott and Pagnucco (1999), which satisfy a strengthened form of inclusion (the contracted theory is a subset of the vacuous

7.2 Bayesian conditioning.

The framework contains classical Bayesian updating as a degenerate special case, recovered under the five jointly restrictive conditions of Proposition 3.9: fixed epistemic level, single-valued (probabilistic) semantics, full credibility, deactivated trigger, and the replacement interpretation of absorption. Under these conditions, the forward update reduces to posterior conditioning: $p_i^+([r]) = p_i([r] \mid E)$ where $E = \{[r] : p_j([r]) > 0\}$.

In the present context, the Bayesian special case occupies a specific position in the triangle of classical frameworks. AGM contraction removes content from a propositional belief set; Bayesian conditioning adds information to a probability distribution; the semantic forward update does both, depending on whether the trigger fires. The Bayesian case is the trigger-free, full-credibility, single-valued limit of the forward branch—the case in which every encounter produces absorption and no encounter produces contraction. It requires the most restrictive parameter configuration of the three classical cases, and relaxing any one of its five conditions moves the update away from Bayesian: dropping full credibility produces filtered updating; dropping the single-valued semantics produces set-valued accumulation (the generic case in the semantic framework); activating the trigger introduces the contraction branch.

The methodological move—present the Bayesian baseline, derive its testable implications, then identify the parameters whose relaxation is required to fit the empirical pattern—has two important precedents in the political-behavior literature that this section follows. Bartels (2002) formalizes a Bayesian learning model in which voters with different prior partisan affinities observe the same evidence, derives the implication that posterior beliefs should converge in the direction of the evidence at a rate determined by its precision, and shows that NES panel data require a partisan-bias parameter offsetting the convergence: the Bayesian benchmark fails empirically, and the parameter that closes the gap is structurally analogous to the credibility filter α_{ij}. Hinich and Munger (1994, ch. 10) write a normal-prior / normal-likelihood Bayesian model of voter inference about candidate ideological position (eqs. 10.2–10.3) and show that the campaign is best understood as variance management—candidates spend resources to reduce σ_θ^2 on themselves and increase σ_ψ^2 on opponents—a corollary of which (eq. 10.15) is that voters strictly prefer the lower-variance candidate when ideological positions are tied. Both projects use the Bayesian baseline as a rhetorical anchor against

contraction) but drop Recovery. The baseline β shares this strengthened inclusion property: $\beta(u_i, u_j) \preceq_{\mathscr{U}} u_i$ with equality iff $\mathsf{Conf} = \emptyset$.

which the relevant departure can be made precise. Proposition 3.9 occupies the same role here: it identifies the five-condition knife-edge under which the semantic update collapses to Bayesian conditioning, and the credibility filter, the trigger, and the set-valued semantics are exactly the parameters whose relaxation is required to recover the empirical pattern Bartels documents and the variance-management strategy Hinich and Munger formalize.[19]

7.3 Public announcement logic as a special case.

Dynamic epistemic logic models information change as transformations of epistemic models—Kripke structures $\mathbf{M} = (W, \{\sim_i\}_{i \in N}, V)$ with possible worlds W, accessibility relations $\sim_i$ for each agent, and a valuation V assigning propositions to worlds. The simplest dynamic operation is the *public announcement* of a proposition P: the model is restricted to worlds satisfying P, yielding $\mathbf{M}|_P = \{w \in W : \mathbf{M}, w \models P\}$ (van Benthem, 2011). Public announcements are "hard information"—irrevocable, fully trusted, and uniformly received by all agents.

The semantic framework recovers public announcement logic under a specific parameter configuration.

Proposition 7.2. *Set $M = \{T, F\}$ (Boolean meanings), fix a common epistemic level $e_\top$ (all representations visible, each in its own class), define $\mathsf{Cf} = \{(T, F), (F, T)\}$, and let each voter have full credibility $\alpha_v([r]) = M$ for all $[r]$, deactivated trigger $\tau_v = \infty$, maximal epistemic expansion, and the* replacement interpretation *of absorption (when the voter absorbs content that conflicts with her prior assignment, the new content supersedes the old rather than accumulating alongside it). Let the broadcast $\phi(x_c)$ assign T to each representation class the announcement endorses and F to each it denies. Then the forward update T_+ reduces to public-announcement-style world-elimination: the voter's post-update semantics assigns T to every class where the announcement asserts T, F to every class where it asserts F, and*

[19]The same methodological tension appears in the political-psychology literature. Sniderman and Bullock (2004) explicitly identify the consistency-vs-updating tension in their account of menu dependence—"if citizens are consistency maximizing for any extended period of time, how are they capable of substantially revising their beliefs?"—and tentatively resolve it by locating belief revision in elite-driven *menu shifts* rather than in individual updating. The trigger / baseline-reversion mechanism of Definition 3.3 is a formal version of the move they gesture toward: individual updating is locally consistency-preserving, and substantive belief change happens through structural shifts in the encountered content (extensions of the form developed in Section 8.3) rather than through smooth Bayesian convergence.

the resulting understanding contains exactly the semantic content consistent with the announcement.

The proof (in Section A) shows that under these conditions, full credibility makes the voter absorb the announcement's complete content. With $M = \{T, F\}$, the replacement interpretation collapses the two-element post-update set $\{T, F\}$ to the announcement's value, reproducing PAL's world-elimination: representations where the prior was "false" are overwritten by "true," and the post-update understanding contains exactly the content consistent with the announcement.[20]

The semantic framework generalizes PAL in four directions. First, the meaning space M can be richer than $\{T, F\}$: evaluative (G, B), affective, and social meanings coexist with factual ones, and a broadcast can carry non-propositional content that PAL cannot express. Second, credibility filters restrict which content passes: PAL assumes every agent fully trusts every announcement (hard information), while the semantic framework models soft information through $\alpha_{vc} \subsetneq M$—the voter hears only what her filter admits. Third, the trigger mechanism allows defensive rejection of announced content: PAL announcements are irrevocable, but the semantic framework permits the voter to refuse the update when conflict exceeds her threshold, reverting to the baseline rather than absorbing. Fourth, epistemic regression: PAL can only eliminate worlds (the model shrinks), while the semantic framework allows the epistemic stance to coarsen (equivalence classes merge), producing an agent who draws *fewer* distinctions after the update than before—a phenomenon with no PAL analog.

The communication equilibrium of Proposition 5.2 has a natural DEL reading. The mutual update map $\mathbf{T}$ is, in DEL terms, a multi-agent action model applied iteratively: each round of mutual updating is a simultaneous "announcement" by every agent of their current semantic content, filtered through the credibility structure. The Tarski fixed-point theorem guarantees that this process stabilizes—the analog of the DEL result that iterated public announcements on a finite model reach a fixed point. The semantic framework adds to this the possibility that the fixed point is not the fully informed state but a stable state of mutual incomprehension: agents whose credibility filters

[20] The replacement interpretation is not the default in the semantic framework, where absorption is accumulative (the voter adds meanings without discarding prior ones). The PAL special case requires replacement precisely because PAL announcements are hard information that supersedes prior beliefs. The accumulative interpretation—in which $\{T, F\}$ is a genuine post-update state representing factual ambivalence—is a generalization with no PAL analog.

are disjoint reach a fixed point in which neither has absorbed anything from the other (Section 5.1).

7.4 Three traditions, one framework.

Framework	M	α	τ	**Epistemic exp.**	**Recovered as**
AGM contraction	any	—	—	—	Baseline β ($\div$1–$\div$5, not $\div$6)
Bayesian	singletons	full	∞	none	Proposition 3.9
PAL/DEL	$\{T, F\}$	full	∞	maximal	Proposition 7.2
Semantic framework	any	any	any	any	Full operator T

The semantic framework is not a competing theory of belief revision, probabilistic updating, or dynamic information change; it is a common generalization that contains each as a parameter-specific special case. What the generalization buys is the interaction among features that each classical tradition treats in isolation: non-Boolean meanings (absent in AGM and PAL), heterogeneous credibility (absent in all three), the trigger mechanism (absent in all three), and the two-branch structure of the update operator (contraction under conflict, absorption otherwise). These are precisely the features the political-communication domain requires—voters process content through evaluative as well as factual channels, trust sources selectively, and sometimes reject information defensively—and each is invisible to the classical frameworks taken individually.

The core framework holds certain features fixed—two candidates, a single broadcast, a stable representational field—that naturally vary in richer political environments. The next section relaxes each in turn.

8 Extensions.

The core framework of Sections 2–5 holds three features fixed that naturally vary: the number of strategic actors (two candidates), the temporal structure of communication (a single broadcast per candidate), and the representational field itself. This section relaxes each in turn. Legislative bargaining replaces two-candidate elections with multi-actor coalition formation under a sequential bargaining protocol; dynamic campaigns replace single-shot broadcasts with sequential messaging and characterize when the order of encounters matters; and endogenous field evolution allows the representational field to vary across agents through algorithmic curation, formalizing the media-fragmentation mechanism that makes Downsification progressively harder to sustain.

8.1 Legislative bargaining.

The electoral results of Sections 4–5 concern two-candidate competition. But the lattice structure of $\mathscr{U}_{\mathcal{R}}$ applies wherever collective decision-making encounters semantic incomparability, and the natural next setting is multi-actor legislative bargaining. The framework here adapts the sequential bargaining model of Austen-Smith and Banks (2005, Ch. 6) to the semantic setting, replacing the Euclidean policy space with the understanding lattice and adding a semantic viability constraint on coalitions that has no spatial analog.

Setup. A legislature $N = \{1, \ldots, n\}$ with n odd makes collective decisions under a q-rule, where $q = (n+1)/2$ (simple majority). Each legislator i has an ideal understanding $u_i^* \in \mathscr{U}_{\mathcal{R}}$ and evaluates outcomes by $\mathcal{W}(\cdot \wedge u_i^*)$. The feasible set of outcomes is the semantic interval $[\bigwedge_i u_i^*, \bigvee_i u_i^*]$—the welfare-relevant bargaining space, bounded by the chaos theorem (Theorem 5.9). A status quo understanding $u^0 \in \mathscr{U}_{\mathcal{R}}$ obtains if no proposal is accepted. In each period $t = 1, 2, \ldots$, a legislator i is randomly recognized with probability $p_i > 0$ and proposes an outcome $u \in [\bigwedge_i u_i^*, \bigvee_i u_i^*]$; the legislature votes, and the proposal passes if at least q members vote in favor; if it fails, the process moves to $t+1$. All legislators share a common discount factor $\delta \in (0, 1)$.

Definition 8.1. *A coalition $S \subseteq N$ is* semantically viable *if $\bigwedge_{i \in S} u_i^* \succ_{\mathscr{U}} u_\perp$: the meet of its members' ideals is strictly above the bottom of the lattice, meaning the coalition shares some representational common ground beyond vacuity.*

Semantic viability is the lattice-theoretic prerequisite for joint action: a coalition whose members share no semantic common ground cannot form a coherent collective position, because any proposal that enriches one member's welfare is orthogonal to another's ideal. In the spatial model, any subset of legislators can form a coalition because the policy space is shared; in the semantic model, coalition formation is constrained by the lattice geometry of the members' ideals.

Stationary strategies. Following Austen-Smith and Banks (2005), we restrict attention to *stationary* strategies: a legislator's proposal and voting behavior depend only on the current proposal and the other legislators' strategies, not on the history of past proposals or votes. Stationarity is natural in the semantic setting for the same reason it is natural in the spatial one: the game's structure is history-independent (the subgame starting at any period t is identical to the subgame at $t' > t$), so history-dependent strategies add complexity without altering the set of subgame perfect equilibrium payoffs for sufficiently patient legislators.

A stationary strategy for legislator i consists of a proposal rule $\pi_i \in \mathscr{U}_{\mathcal{R}}$ (the understanding i proposes when recognized) and a voting rule $v_i : \mathscr{U}_{\mathcal{R}} \to \{0, 1\}$ (accept or reject any proposal). Let $V_i(\sigma)$ denote i's *continuation value*—the expected discounted welfare from period $t + 1$ onward—under stationary strategy profile σ.

Proposition 8.2. *In a semantic legislature with finite* $[\bigwedge_i u_i^*, \bigvee_i u_i^*]$*, n odd, and common discount factor* $\delta \in (0, 1)$*:*

(i) A stationary subgame perfect equilibrium exists.

(ii) In any stationary equilibrium, the recognized proposer i selects a winning coalition $S_i \ni i$ with $|S_i| = q$ and proposes $\pi_i \in [\bigwedge_i u_i^*, \bigvee_i u_i^*]$ that maximizes $\mathcal{W}(\pi_i \wedge u_i^*)$ subject to $\mathcal{W}(\pi_i \wedge u_j^*) \geq \delta V_j(\sigma)$ for all $j \in S_i \setminus \{i\}$.

(iii) All stationary equilibria are no-delay: the first proposal is accepted.

(iv) The proposer's winning coalition S_i is drawn from the set of cheapest $q - 1$ partners—those legislators j for whom $\delta V_j(\sigma)$ is smallest—subject to the constraint that S_i is semantically viable (Definition 8.1).

Part (i): the game is a finite-action extensive form with discounting, so a stationary subgame perfect equilibrium exists by standard results (cf. Austen-Smith and Banks, 2005, Theorem 6.2). Part (ii): the proposer faces a finite menu of platforms in the semantic interval; she selects the one maximizing her welfare subject to buying $q - 1$ partners at their continuation values. Because the feasible set is finite rather than convex (in contrast to the Euclidean setting of Austen-Smith and Banks 2005, Ch. 6), the proposer may need to *overpay* some coalition partners: the discrete menu may not contain a platform delivering partner j exactly $\delta V_j(\sigma)$, so the proposer offers the cheapest platform that weakly exceeds the target. Part (iii): delay is never optimal because the proposer can always offer a platform that gives every coalition partner at least $\delta V_j(\sigma)$ (the continuation value is achievable or exceedable on the finite menu) and extract the surplus from being recognized; waiting one period costs δ and gains nothing.

The semantic novelty is in part (iv): the viability constraint. In the spatial model, the proposer simply buys the cheapest $q-1$ votes; any subset of legislators is a feasible coalition. In the semantic model, a coalition is feasible only if its members share enough representational common ground to form a coherent position—one that lies above the meet of their ideals. When the cheapest $q-1$ partners do not form a semantically viable coalition (their meet is $u_\perp$), the proposer must pay more for a viable alternative, and the equilibrium outcome shifts.

Corollary 8.3. *When the semantic dimension of the legislature is* 1 *(all legislator ideals are comparable), the semantic core is non-empty and contains the legislator whose ideal is the median of the chain. As $\delta \to 1$, all stationary equilibrium outcomes converge to a proposal that maximizes welfare at the semantic median.*

This is the semantic analog of the standard one-dimensional core convergence result: when the lattice structure supports a single organizing dimension, the median legislator commands the core, and patient legislators converge to the median outcome. When semantic dimension exceeds 1, the semantic core may be empty—the Downsification conditions fail and no single proposal defeats all alternatives—but the chaos bound confines all stationary equilibrium outcomes to the semantic interval. The lattice provides a structural bound on legislative outcomes that the spatial model, in the absence of a core, cannot supply.

The immigration legislature. Suppose the three immigration voters serve as legislators with recognition probabilities $p_A = p_B = p_C = 1/3$ and majority rule ($q = 2$). The feasible set is the semantic interval $[u_\perp,\ (\{T,G,B\},\ \{T,F,G,B\},\ \{T,G,B\})]$. Coalition $\{B,C\}$ is semantically viable: $u_B^* \wedge u_C^* = u_C^* = (\{B\},\{G\},\emptyset) \succ_{\mathscr{U}} u_\perp$. Coalition $\{A,B\}$ is viable: $u_A^* \wedge u_B^* = (\{T\},\emptyset,\{T\}) \succ_{\mathscr{U}} u_\perp$. Coalition $\{A,C\}$ is barely viable: $u_A^* \wedge u_C^* = (\emptyset,\emptyset,\emptyset) = u_\perp$—the enforcement and affect voters share no representational common ground, making joint action vacuous. When B is recognized, she proposes an understanding near u_B^* and buys C's vote cheaply (since $u_C^* \preceq_{\mathscr{U}} u_B^*$, any proposal above u_B^* also satisfies C). When A is recognized, she must form a coalition with B (since $\{A,C\}$ is not viable), and the proposal must lie above $(\{T\},\emptyset,\{T\})$—the bare factual common ground between enforcement and integration. The semantic viability constraint binds: it is not enough for A to offer B a high continuation value; the proposal must be semantically coherent for the coalition, which forces A away from her ideal toward the A–B common ground.

8.2 Dynamic campaigns.

The electoral model of Section 4 treats each candidate's campaign as a single broadcast: the voter receives $\phi(x_c)$ and updates once. In practice, campaigns are sequences of messages, and the voter's response to a late message depends on what earlier messages have already done to her understanding. This subsection formalizes dynamic campaigns as iterated applications of the update operator and characterizes when the resulting process converges, when it depends on the order of messages, and what these properties imply for campaign strategy.

Setup. A *dynamic campaign* is a finite sequence of platforms $(x_{c,1}, \ldots, x_{c,T})$ chosen by candidate c. The voter updates iteratively:

$$\hat{u}_v^{c,0} = u_v^0,$$
$$\hat{u}_v^{c,t} = T_{\alpha_{vc},\tau_v}\big(\hat{u}_v^{c,t-1},\ \phi(x_{c,t})\big), \qquad t = 1, \ldots, T.$$

The voter's final post-campaign understanding is $\hat{u}_v^{c,T}$. The single-shot model of Section 4.2 is the special case $T = 1$.

Proposition 8.4. *When $\tau_v = \infty$ (triggers deactivated) and epistemic expansion is minimal ($e'_v = e_v$ for all v), the iterative forward update converges to a fixed point in at most $|\mathscr{U}_\mathcal{R}|$ steps: there exists $T^* \leq |\mathscr{U}_\mathcal{R}|$ such that $\hat{u}_v^{c,t} = \hat{u}_v^{c,T^*}$ for all $t \geq T^*$. The fixed point $\hat{u}_v^{c,T^*}$ is the voter's* absorption ceiling*—the richest understanding attainable by iteratively absorbing content from the candidate's broadcast sequence under the voter's credibility filter.*

The proof exploits extensivity. By Lemma 3.5, $u_v \preceq_\mathscr{U} T_+(u_v, u_j)$ for every forward update: the voter's understanding weakly increases at each step. Since $\mathscr{U}_\mathcal{R}$ is a finite lattice, a weakly increasing sequence can make at most $|\mathscr{U}_\mathcal{R}|$ strict upward steps before reaching a fixed point. At the fixed point, no further content from the broadcast sequence passes through the credibility filter: the voter has absorbed everything she can absorb, and further repetition of any message in the sequence leaves her understanding unchanged.

The absorption ceiling depends on the credibility filter α_{vc} and the content of the broadcast sequence, but not on the order of messages within the sequence—a consequence of commutativity under source-independent filters, proved below. Convergence is guaranteed but the ceiling may be far below the candidate's broadcast: the credibility filter can block most content, leaving the voter at a stable but impoverished state.

Proposition 8.5. *When triggers are active ($\tau_v < \infty$) and $\mathscr{U}_\mathcal{R}$ is finite, every finite dynamic campaign produces a determinate terminal state that depends on the broadcast sequence. The terminal state may be the absorption ceiling (if the trigger never fires) or an impoverished state reached through one or more trigger-firings. When the same broadcast is repeated, the voter may cycle between absorption and reversion rather than converging.*

The contrast with the trigger-free case is sharp. Without triggers, the monotone convergence of Proposition 8.4 guarantees a unique fixed point independent of message order. With triggers, neither convergence nor uniqueness is guaranteed: each trigger-firing produces a strict downward move to the baseline (Proposition 3.4), and a subsequent forward update may restore enough content to fire the trigger again, creating a cycle between absorption and reversion. Finite non-repeating campaigns always terminate (the voter processes each message once and stops), but the terminal state depends on the sequence—different orderings of the same messages can produce different outcomes. This path-dependence is the formal content of campaign sequencing effects.

Path-dependence. Whether the order of messages matters depends on the credibility structure.

Proposition 8.6. *When credibility filters are source-independent ($\alpha_{vc} = \alpha_v$ for all c), triggers are deactivated ($\tau_v = \infty$), and all encounters are processed at a common epistemic level, the forward update is commutative: for any understandings $u_j, u_k \in \mathscr{U}_\mathcal{R}$,*

$$T_+\big(T_+(u_v, u_j),\ u_k\big) \ = \ T_+\big(T_+(u_v, u_k),\ u_j\big).$$

Under a common filter α_v and a common epistemic level e, the forward update at each representation class is $\mu_v^+([r]) = \mu_v([r]) \cup (\alpha_v([r]) \cap \mu_j([r]))$. Applying sequentially with u_j then u_k:

$$\mu_v([r]) \ \cup \ \big(\alpha_v([r]) \cap \mu_j([r])\big) \ \cup \ \big(\alpha_v([r]) \cap \mu_k([r])\big) \ = \ \mu_v([r]) \ \cup \ \big(\alpha_v([r]) \cap (\mu_j([r]) \cup \mu_k([r]))\big).$$

The right-hand side is symmetric in j and k: union is commutative and associative, and the filter α_v is applied identically in both orders.

Commutativity breaks when credibility depends on the source.

Proposition 8.7. *When credibility filters are source-dependent ($\alpha_{vc} \neq \alpha_{vc'}$ for some v, c, c') and triggers are active ($\tau_v < \infty$), the update is generically non-commutative: there exist understandings u_j, u_k and a voter v such that the terminal state after encountering u_j then u_k differs from the terminal state after u_k then u_j.*

The mechanism is the trigger. Suppose voter v trusts source j broadly ($\alpha_{vj}([r]) = M$ for all $[r]$) but trusts source k narrowly ($\alpha_{vk}([r]) = M_{\text{policy}}$). Encountering j first, the voter absorbs all of j's content—including potentially conflict-generating affective meanings—which may push $|\mathsf{Conf}|$ above τ_v and fire the trigger. The voter reverts to baseline before encountering k, and k's narrower content is absorbed from the impoverished state. Encountering k first, the voter absorbs only policy content (which generates less conflict), the trigger does not fire, and j's full content is then absorbed from a richer starting state. The two orderings produce different terminal states: order matters because the trigger's firing depends on the conflict generated by each encounter, which in turn depends on the voter's state at the time of the encounter.

This is the formal content of the campaign priming effect noted in Remark 4.3: at a symmetric equilibrium, the relevant voter state when evaluating a deviation is the *primed* state—the state after absorbing the opponent's broadcast—not the initial state. Dynamic campaigns generalize this from a two-message setting to an arbitrary sequence, and the non-commutativity result shows that the strategic implications are genuine: a candidate who controls the order of messages can exploit path-dependence to reach voter states that are inaccessible under other orderings.

Immigration illustration. Consider voter A encountering two messages: an integration broadcast u_L (source j, with broad credibility $\alpha_{Aj} = M$) and an affect broadcast u_E (source k, with narrow credibility $\alpha_{Ak} = \{G, B\}$). Integration first: $|\mathsf{Conf}(u_A, u_L)| = 3$; if $\tau_A = 2$, the trigger fires and A reverts to $\beta = (\{T\}, \emptyset, \{T\})$. From this impoverished state, the affect broadcast adds evaluative content through the narrow filter—A ends at $(\{T, B\}, \{B\}, \{T, B\})$ or similar. Affect first: u_E carries only evaluative content through the narrow filter, generating minimal conflict (only evaluative meanings at risk). If $|\mathsf{Conf}| \leq \tau_A$, the trigger does not fire; A absorbs the affect content and then encounters u_L from a richer state, possibly still triggering but from a higher baseline. The two orderings can produce different terminal states, different welfare evaluations, and therefore different voting behavior.

8.3 Endogenous field evolution.

The representational field $\mathcal{R}$ has been fixed throughout the manuscript: all agents operate within a shared ontology of representational resources. In practice, the set of representations available to any individual voter depends on what information environment she inhabits. A voter who consumes news primarily through algorithmically curated social media feeds encounters a different portion of $\mathcal{R}$ than one who reads broadsheet newspapers; the two may access the same political domain but through representational subfields that overlap only partially.

This is not a new observation—the media-fragmentation literature has documented the phenomenon extensively (Flaxman, Goel and Rao, 2016)—but the semantic framework provides a precise formal language for its consequences. When voters operate in different subfields of $\mathcal{R}$, their epistemic stances and ideal understandings are defined relative to different representational bases, and the lattice relationships among them can change. The key result of this subsection is that field fragmentation can destroy comparability: voters whose ideals are comparable in the full field can become incomparable when each sees only a curated portion of it.

Grothendieck universes and field structure. The formal apparatus for field evolution comes from set-theoretic foundations. A *Grothendieck universe* is a set $\mathfrak{U}$ satisfying four closure conditions: (i) if $x \in \mathfrak{U}$ and $y \in x$, then $y \in \mathfrak{U}$ (transitivity); (ii) if $x, y \in \mathfrak{U}$, then $\{x, y\} \in \mathfrak{U}$; (iii) if $x \in \mathfrak{U}$, then $\mathcal{P}(x) \in \mathfrak{U}$ (closure under power sets); and (iv) if $I \in \mathfrak{U}$ and $\{x_i\}_{i \in I}$ is a family with each $x_i \in \mathfrak{U}$, then $\bigcup_{i \in I} x_i \in \mathfrak{U}$ (closure under indexed unions). A Grothendieck universe is a "set-sized model of set theory": every construction that ZFC permits can be carried out within $\mathfrak{U}$, and the universe is large enough to contain all the sets one needs for the construction.

The representational field $\mathcal{R}$ lives in some Grothendieck universe $\mathfrak{U}_0$. The understanding lattice $\mathscr{U}_{\mathcal{R}}$, the meaning space M, the credibility filters, and all the formal apparatus of Sections 2–5 are constructed within $\mathfrak{U}_0$. The Grothendieck axiom—every set belongs to some universe—guarantees that $\mathfrak{U}_0$ itself is a member of a larger universe $\mathfrak{U}_1 \supset \mathfrak{U}_0$, within which a richer representational field $\mathcal{R}' \supset \mathcal{R}$ can be defined.

Field *expansion*—a politician introduces crime statistics disaggregated by immigration status, adding new representations to the discourse—corresponds to moving from $\mathcal{R} \in \mathfrak{U}_0$ to $\mathcal{R}' \in \mathfrak{U}_1$. The understanding lattice over $\mathcal{R}'$ is strictly richer than over $\mathcal{R}$: new epistemic stances, new semantics, new understandings become available. Field *restriction*—algorithmic curation limits voter v to a subfield $\mathcal{R}_v \subset \mathcal{R}$—corresponds to restricting to a sub-universe. The voter's effective understanding lattice is $\mathscr{U}_{\mathcal{R}_v}$, which is a sublattice of $\mathscr{U}_{\mathcal{R}}$ but may have different order-theoretic properties.

Definition 8.8. *A* field fragmentation *of $\mathcal{R}$ over a voter population V is a profile $(\mathcal{R}_v)_{v \in V}$ with $\mathcal{R}_v \subseteq \mathcal{R}$ for each v. The* restricted ideal *of voter* *v is $u_v^*|_{\mathcal{R}_v}$: the projection of u_v^* onto the subfield $\mathcal{R}_v$, defined by restricting the epistemic stance and semantics to the representations in $\mathcal{R}_v$. The* zero-extension $\overline{u_v^*|_{\mathcal{R}_v}} \in \mathscr{U}_{\mathcal{R}}$ *embeds the restricted ideal back into the full lattice by assigning empty semantics ($\emptyset$) to every representation class in $\mathcal{R} \setminus \mathcal{R}_v$: the voter retains her content on her curated field and has nothing to say about the rest.*

A fragmentation is not a model of disagreement—it is a model of ontological divergence. Voters in a fragmented field do not merely assign different meanings to shared representations; they operate within different representational vocabularies entirely, so that the lattice relationships among their understandings can change. Two voters who are comparable in the full field—one's understanding strictly richer than the other's—can become incomparable when each sees only a curated portion of $\mathcal{R}$, because the representations that made one richer than the other may be visible to only one of them.

The central result of this subsection makes this precise. It says that comparability is fragile: for any pair of strictly comparable voter ideals, there exists a fragmentation—one that still leaves the two voters with overlapping fields—that destroys the comparability. The lattice order that organized their relationship in the full field does not survive differential curation.

Theorem 8.9. *Let $u_1^* \prec_{\mathscr{U}} u_2^*$ be strictly comparable ideal understandings in $\mathscr{U}_{\mathcal{R}}$, with $|\mathcal{R}| \geq 3$. There exists a field fragmentation $(\mathcal{R}_1, \mathcal{R}_2)$ with $\mathcal{R}_1 \cap \mathcal{R}_2 \neq \emptyset$ such that the zero-extended restricted ideals are incomparable in $\mathscr{U}_{\mathcal{R}}$: neither $\overline{u_1^*|_{\mathcal{R}_1}} \preceq_{\mathscr{U}} \overline{u_2^*|_{\mathcal{R}_2}}$ nor $\overline{u_2^*|_{\mathcal{R}_2}} \preceq_{\mathscr{U}} \overline{u_1^*|_{\mathcal{R}_1}}$. (When $|\mathcal{R}| = 2$, only total fragmentation—$\mathcal{R}_1 \cap \mathcal{R}_2 = \emptyset$—can destroy comparability.)*

The proof is constructive, and the comparison takes place in the full lattice $\mathscr{U}_{\mathcal{R}}$—not in a product of sublattices—so incomparability is non-trivial. Since $u_1^* \prec_{\mathscr{U}} u_2^*$, there exists a representation $r_a \in \mathcal{R}$ at which u_2^* has strictly richer semantic content than u_1^*: either $\mu_1(r_a) \subsetneq \mu_2(r_a)$ or r_a is visible to e_2 but not to e_1. Since $|\mathcal{R}| \geq 3$, there exists $r_b \neq r_a$ in $\mathcal{R}$ at which u_1^* has nonempty content ($\mu_1(r_b) \neq \emptyset$). Set $\mathcal{R}_1 = \{r_a, r_b\}$ and $\mathcal{R}_2 = \mathcal{R} \setminus \{r_b\}$, so $\mathcal{R}_1 \cap \mathcal{R}_2 = \{r_a\} \neq \emptyset$. The zero-extension $\overline{u_1^*|_{\mathcal{R}_1}}$ has nonempty content at r_b but assigns $\emptyset$ to r_a (since $r_a \notin \mathcal{R}_1$). The zero-extension $\overline{u_2^*|_{\mathcal{R}_2}}$ has nonempty content at r_a (retained from u_2^*) but assigns $\emptyset$ to r_b (since $r_b \notin \mathcal{R}_2$). For $\overline{u_1^*|_{\mathcal{R}_1}} \preceq_{\mathscr{U}} \overline{u_2^*|_{\mathcal{R}_2}}$ we would need $\mu_1(r_b) \subseteq \emptyset$, which fails since $\mu_1(r_b) \neq \emptyset$. For the reverse, $\overline{u_2^*|_{\mathcal{R}_2}} \preceq_{\mathscr{U}} \overline{u_1^*|_{\mathcal{R}_1}}$ we would need $\mu_2(r_a) \subseteq \emptyset$, which fails since u_2^* has content at r_a. Neither zero-extension is below the other in $\mathscr{U}_{\mathcal{R}}$: the fragmentation has destroyed comparability.

The condition $\mathcal{R}_1 \cap \mathcal{R}_2 \neq \emptyset$ is important: the theorem does not require that the two voters be placed in entirely disjoint informational worlds. Even fragmentations that leave substantial overlap—voters who still share most of their representational vocabulary—can destroy comparability, as long as the curation is asymmetric at the representations where the gap between u_1^* and u_2^* resides. The fragility is not in the amount of shared content but in its location: comparability depends on a specific lattice relationship that can be severed by removing even a single representation from one voter's field.

Corollary 8.10. *If a voter population has semantic dimension* 1 *under the full field* $\mathcal{R}$ *(all voter ideals are comparable), there exists a field fragmentation that increases the semantic dimension to exceed* 1*, making the election non-Downsifiable under the restricted fields.*

Apply Theorem 8.9 to any strictly comparable pair: the fragmentation makes their zero-extended restricted ideals incomparable in $\mathscr{U}_{\mathcal{R}}$, increasing the antichain width. By Proposition 5.7, the Downsification conditions have no structural guarantee when semantic dimension exceeds 1.

Broadcasting and narrowcasting. The fragmentation theorem formalizes the transition from the broadcast era to the algorithmic curation era. In the broadcast era, all voters share the full field $\mathcal{R}$: they may disagree about meanings (different μ_v) and about the fineness of their representations (different e_v), but they operate within a shared ontology. If the underlying ideal distribution has semantic dimension 1—all ideals lie on a chain in $\mathscr{U}_{\mathcal{R}}$—then the Downsification conditions can hold and the Downs model is applicable. The shared field keeps voter stances comparable, and competition can be organized along a single dimension.

Algorithmic curation breaks this. When voter v's effective field is $\mathcal{R}_v \subsetneq \mathcal{R}$, the voter's understanding lattice is $\mathscr{U}_{\mathcal{R}_v}$—a sublattice that may lack representations present in other voters' lattices. Even if the underlying ideals are comparable in the full field, curation can fragment the field enough to make restricted ideals incomparable: voters whose belief systems were organized along a common dimension in the broadcast era find themselves on orthogonal dimensions in the curated era, because each sees a different portion of the representational domain.

Flaxman, Goel and Rao (2016) document that ideological segregation in online news consumption is real and growing: conservative and liberal web users share fewer common news sources over time. In the semantic framework, this IS field restriction—the news sources a voter consumes determine which representations she can access, which meanings she encounters, and therefore which portion of $\mathscr{U}_{\mathcal{R}}$ her ideal understanding inhabits. The fragmentation theorem shows that this process has a structural consequence beyond filter bubbles: it can destroy the comparability of voter ideals and thereby eliminate the possibility of Downsification, even when the underlying political domain would support a single organizing dimension if all voters could see it.

Mondak et al. (forthcoming) provide the behavioral signature. Their finding that partisan issue selections cluster sharply by party—with the average partisan part of an approximately 83% majority on the issue they choose to discuss, and most unable to articulate the opposing side's rationale—is precisely what advanced field fragmentation looks like from the outside: voters' effective representational fields have diverged to the point where they engage with different portions of the political domain, and the common ground ($\mathcal{R}_1 \cap \mathcal{R}_2$) has shrunk below the threshold needed for mutual comprehension.

The immigration case. Under the full field $\mathcal{R} = \{r_1, r_2, r_3\}$, voters A and B are already incomparable—their ideals carry different semantic content at every representation class—but voter C's ideal satisfies $u_C^* \preceq_{\mathscr{U}} u_B^*$ (C is below B in the lattice). The semantic dimension is 2 (the antichain $\{u_A^*, u_B^*\}$).

Now suppose algorithmic curation restricts voter A to $\mathcal{R}_A = \{r_1, r_3\}$ (enforcement and economic impact only—the representations that anchor A's schema) and voter B to $\mathcal{R}_B = \{r_2, r_3\}$ (citizenship and economic impact—the representations that anchor B's). The shared field shrinks from $\mathcal{R}$ to $\mathcal{R}_A \cap \mathcal{R}_B = \{r_3\}$: economic impact is the only representation both voters can access. The common semantic ground at r_3 is $\mu_A(r_3) \cap \mu_B(r_3) = \{T, B\} \cap \{T, G\} = \{T\}$—bare factual agreement that immigration has economic consequences, with no shared evaluative content. The legislative coalition viability of $\{A, B\}$ degrades: under the full field, $u_A^* \wedge u_B^* = (\{T\}, \emptyset, \{T\})$; under fragmentation, the common ground is even thinner, because each voter's restricted ideal lacks the representations the other's schema requires.

The Downs model worked in the broadcast era not because voters agreed but because they operated within a shared representational field that kept their disagreements structured—organized along dimensions that could, in principle, be collapsed into a single axis. Algorithmic curation does not merely make voters more extreme or more partisan; it fragments the field itself, producing a structural incomparability that no amount of moderation, deliberation, or compromise can resolve within the fragmented ontology. Resolution requires field expansion—restoring shared representations—which is precisely the reverse of what curation algorithms are designed to do.

9 Conclusion.

This manuscript has developed a semantic theory of electoral competition in which the 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. The Downsification theorem (Theorem 4.11) makes the diagnostic precise: a symmetric equilibrium admits a median-voter interpretation only if the coalition mountain condition holds, and the CMC in turn forces the winning platform to be the semantically richest broadcast available. The four Downsification conditions (Section 4.7) characterize the semantic parameters sufficient to ensure the CMC. The conditions are demanding: full policy credibility, single dimensionality, single-peaked welfare, and trigger-free voters. The empirical record suggests all four are generically violated.

When the conditions fail, the semantic lattice provides three results that the spatial framework cannot state. The chaos bound (Theorem 5.9) confines the welfare-relevant majority-preference orbit to the semantic interval—a structural bound determined entirely by the distribution of voter ideal understandings, providing the intermediate constraint between core and Pareto set that classical positive political theory identifies as missing (Austen-Smith and Banks, 1999). The polysemy decomposition (Corollary 5.12) shows that when credibility filters are heterogeneous enough, the optimal campaign is not a coherent policy message but a structured broadcast whose components are independently optimized for disjoint audiences—message ambiguity as an equilibrium property, not a communication failure. And the majoritarian convergence result (Corollary 5.15) shows that when voter ideals are incomparably structured and one bloc constitutes a strict majority, competition forces equilibrium toward the majority bloc's ideal; the minority's welfare shortfall is determined by the lattice-theoretic incomparability of voter ideals, not by the geometry of a pre-given policy axis.

The manuscript has developed three further contributions beyond the core results. The empirical predictions of Section 6 show that the framework generates specific, testable claims—the sophistication-proximity inversion, the polysemy prediction under credibility fragmentation, and the incomparability-based structure of polarization—that are empirically distinguishable from what the spatial model predicts, and that the key constructs (epistemic stance, credibility filter, ideal understanding, incomparability) can be operationalized using existing survey infrastructure. The connections to classical epistemic frameworks in Section 7 show that the update operator generalizes AGM contraction (satisfying the basic postulates minus Recovery), Bayesian conditioning (as a degenerate special case under five restrictive conditions), and public announcement logic (recovered when meanings are Boolean and credibility is full)—positioning the semantic framework as a common generalization of three traditions that each capture a fragment of the political-communication domain. And the extensions of Section 8 show that the framework's lattice structure extends naturally to legislative bargaining (where semantic viability constrains coalition formation), dynamic campaigns (where the order of messages matters when credibility is source-dependent), and endogenous field evolution (where algorithmic curation can destroy the comparability of voter ideals and thereby eliminate the possibility of Downsification).

Several questions remain genuinely open. The ideal understanding u_v^* is a reduced-form modeling device; the microfoundation of the link between semantic cognition and political preference—how the structure of a voter's understanding gives rise to evaluative dispositions over candidate broadcasts—is where future theoretical work belongs. The operationalization of Section 6 identifies what to measure but not what the measurements show; whether the sophistication-proximity inversion, the polysemy prediction, and the incomparability structure of polarization hold in data is an empirical question that the framework makes precise but does not answer. The Grothendieck universe hierarchy introduced for endogenous field evolution (Section 8.3) formalizes ontological change as movement between levels of a set-theoretic hierarchy, but the strategic analysis of field manipulation—candidates choosing which representations to introduce or retire—remains undeveloped.

The spatial model is correct on its own assumptions. This manuscript has characterized those assumptions precisely, shown that they constitute a knife-edge, and provided a formal account of what electoral competition looks like on both sides of it. The three voters from the introduction—the enforcement voter, the integration voter, and the affect voter—are not arguing about where to stand on a shared axis. They are operating through different representational systems that may not admit a common ranking. The question is not where they disagree but whether the space in which they could disagree exists at all.

References

Acharya, Avidit, Matthew Blackwell and Maya Sen. 2018. "Explaining Preferences from Behavior: A Cognitive Dissonance Approach." *Journal of Politics* 80(2):400–411.

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.

Austen-Smith, David and Jeffrey S. Banks. 1999. *Positive Political Theory I: Collective Preference*. Ann Arbor: University of Michigan Press.

Austen-Smith, David and Jeffrey S. Banks. 2005. *Positive Political Theory II: Strategy and Structure*. Ann Arbor: University of Michigan Press.

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.

Bartels, Larry M. 2002. "Beyond the Running Tally: Partisan Bias in Political Perceptions." *Political Behavior* 24(2):117–150.

Bendor, Jonathan, Daniel Diermeier, David A. Siegel and Michael M. Ting. 2011. *A Behavioral Theory of Elections*. Princeton, NJ: Princeton University Press.

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. In *Ideology and Discontent*, ed. David E. Apter. New York: Free Press pp. 206–261.

Davey, B. A. and H. A. Priestley. 2002. *Introduction to Lattices and Order*. 2nd ed. Cambridge: Cambridge University Press.

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.

Festinger, Leon. 1957. *A Theory of Cognitive Dissonance*. Evanston, IL: Row, Peterson.

Flaxman, Seth, Sharad Goel and Justin M. Rao. 2016. "Filter Bubbles, Echo Chambers, and Online News Consumption." *Public Opinion Quarterly* 80(S1):298–320.

Harmon-Jones, Eddie and Judson Mills. 2019. An Introduction to Cognitive Dissonance Theory and an Overview of Current Perspectives on the Theory. In *Cognitive Dissonance: Reexamining a Pivotal Theory in Psychology*, ed. Eddie Harmon-Jones. 2nd ed. Washington, DC: American Psychological Association pp. 3–24.

Hinich, Melvin J. and Michael C. Munger. 1994. *Ideology and the Theory of Political Choice*. Ann Arbor: University of Michigan Press.

Hurwitz, Jon and Mark Peffley. 1987. "How Are Foreign Policy Attitudes Structured? A Hierarchical Model." *American Political Science Review* 81(4):1099–1120.

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* 22:129–146.

Kollman, Ken, John H. Miller and Scott E. Page. 1992. "Adaptive Parties in Spatial Elections." *American Political Science Review* 86(4):929–937.

Kollman, Ken, John H. Miller and Scott E. Page. 1998. "Political Parties and Electoral Landscapes." *British Journal of Political Science* 28(1):139–158.

Kunda, Ziva. 1990. "The Case for Motivated Reasoning." *Psychological Bulletin* 108(3):480–498.

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

Lau, Richard R. and David P. Redlawsk. 2006. *How Voters Decide: Information Processing during Election Campaigns*. New York: Cambridge University 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.

McKelvey, Richard D. 1979. "General Conditions for Global Intransitivities in Formal Voting Models." *Econometrica* 47(5):1085–1112.

McKelvey, Richard D. and Norman Schofield. 1987. "Generalized Symmetry Conditions at a Core Point." *Econometrica* 55(4):923–933.

Mondak, Jeffery J., Matthew V. Hibbing, Matthew Mettler and Isaiah R. Raynal. forthcoming. *(Mis)Understanding the Other Side: How Americans Explain Opposing Political Views*. Philadelphia: Temple University Press.

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.

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

Peffley, Mark A. and Jon Hurwitz. 1985. "A Hierarchical Model of Attitude Constraint." *American Journal of Political Science* 29(4):871–890.

Rott, Hans and Maurice Pagnucco. 1999. "Severe Withdrawal (and Recovery)." *Journal of Philosophical Logic* 28(5):501–547.

Schofield, Norman. 1978. "Instability of Simple Dynamic Games." *Review of Economic Studies* 45(3):575–594.

Schofield, Norman. 1983. "Generic Instability of Majority Rule." *Review of Economic Studies* 50(4):695–705.

Schuessler, Alexander A. 2000. *A Logic of Expressive Choice*. Princeton, NJ: Princeton University Press.

Shepsle, Kenneth A. 1972. "The Strategy of Ambiguity: Uncertainty and Electoral Competition." *American Political Science Review* 66(2):555–568.

Sniderman, Paul M. and John Bullock. 2004. A Consistency Theory of Public Opinion and Political Choice: The Hypothesis of Menu

Dependence. In *Studies in Public Opinion: Attitudes, Nonattitudes, Measurement Error, and Change*, ed. Willem E. Saris and Paul M. Sniderman. Princeton, NJ: Princeton University Press pp. 337–357.

Stokes, Donald E. 1963. "Spatial Models of Party Competition." *American Political Science Review* 57(2):368–377.

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

Tarski, Alfred. 1955. "A Lattice-Theoretical Fixpoint Theorem and Its Applications." *Pacific Journal of Mathematics* 5(2):285–309.

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

Wittman, Donald. 1983. "Candidate Motivation: A Synthesis of Alternative Theories." *American Political Science Review* 77(1):142–157.

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

A Proofs.

A.1 Proofs for Section 2.

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* $e_\top = (\mathcal{R}, \mathrm{id}_{\mathcal{R}})$ *(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 $\bigcap_j \mathcal{R}_{e_j}$. To verify the visible-scope claim: any $r \in \bigcap_j \mathcal{R}_{e_j}$ lies in a non-invisible equivalence class under each $\sim_j$. Since each $\sim_j$ places all invisible elements in a single class separated from all visible classes, r is not $\sim_j$-equivalent to any invisible element for any j. The union $\bigcup_j \sim_j$ therefore contains no pair (r, s) with $r \in \bigcap_j \mathcal{R}_{e_j}$ and $s \notin \mathcal{R}_{e_k}$ for some k, and the transitive closure $\sim^\wedge$ cannot introduce such a pair. Conversely, any $r \notin \bigcap_j \mathcal{R}_{e_j}$ is invisible to some e_k and hence $\sim_k$-equivalent to every other element of I_{e_k}; under $\sim^\wedge$ it belongs to the merged invisible class.

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. □

Proposition 2.8. *The pair* $(\mathrm{Crs}_{e_2 \to e_1}, \mathrm{Ref}_{e_1 \to e_2})$ *is a monotone Galois connection: for any* $\mu_{e_1} \in M_{e_1}$ *and* $\mu_{e_2} \in M_{e_2}$*,*

$$\mathrm{Crs}_{e_2 \to e_1}(\mu_{e_2})(C) \subseteq \mu_{e_1}(C) \ \ \textit{for all } C \quad \Longleftrightarrow \quad \mu_{e_2}(D) \subseteq \mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1})(D) \ \ \textit{for all } D.$$

Proof. $\mathrm{Crs}_{e_2 \to e_1}(\mu_{e_2})(C) = \bigcup_{D \subseteq C} \mu_{e_2}(D)$ and $\mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1})(D) = \mu_{e_1}(C)$ where C is the unique e_1-class containing D. The forward direction: if $\bigcup_{D \subseteq C} \mu_{e_2}(D) \subseteq \mu_{e_1}(C)$ for all C, then in particular $\mu_{e_2}(D) \subseteq \mu_{e_1}(C) = \mathrm{Ref}_{e_1 \to e_2}(\mu_{e_1})(D)$ for every D. The reverse direction: if $\mu_{e_2}(D) \subseteq \mu_{e_1}(C)$ for every $D \subseteq C$, then $\bigcup_{D \subseteq C} \mu_{e_2}(D) \subseteq \mu_{e_1}(C)$. □

Theorem 2.9. $(\mathscr{U}_\mathcal{R}, \preceq_\mathscr{U})$ *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.*

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. We verify $e^\wedge$-consistency of $\mu_\wedge$: for the invisible class $I_{e^\wedge} = \mathcal{R} \setminus \bigcap_j \mathcal{R}_{e_j}$, any $r \in I_{e^\wedge}$ satisfies $r \notin \mathcal{R}_{e_k}$ for some k, so the $\sim_{e_k}$-class of r is I_{e_k} and $\mu_{e_k}(I_{e_k}) = \emptyset$ by e_k-consistency. This $\emptyset$ appears as a term in the intersection defining $\mu_\wedge(I_{e^\wedge})$, giving $\mu_\wedge(I_{e^\wedge}) = \emptyset$. For $e^\vee$-consistency of $\mu_\vee$: the invisible class $I_{e^\vee} = \mathcal{R} \setminus \bigcup_j \mathcal{R}_{e_j}$ consists of elements invisible to *every* e_j. For any such element, $\mathrm{Ref}_{e_j \to e^\vee}(\mu_{e_j})$ assigns it $\mu_{e_j}(I_{e_j}) = \emptyset$, so $\mu_\vee(I_{e^\vee}) = \bigcup_j \emptyset = \emptyset$.

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 $\mu_{e_\perp}([\mathcal{R}]) = \emptyset$ 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 $\{r\} \subseteq C$. □

A.2 Proofs for Section 3.

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

Lemma 3.5. *For any understandings u_i, u_j with $|\mathsf{Conf}(u_i, u_j)| \leq \tau_i$ (the trigger does not fire):*

$$u_i \ \preceq_{\mathscr{U}}\ T_+(u_i, u_j).$$

The forward update is weakly above the input: encountering another understanding can only enrich, never impoverish, when the trigger does not fire.

Proof. Write $T_+(u_i, u_j) = (e'_i, \mathrm{Ref}_{e_i \to e'_i}(\mu_i^+))$. The epistemic condition holds by construction: $e_i \preceq_E e'_i$. For the semantic condition: $\mu_i^+([r]) = \mu_i([r]) \cup (\alpha_{ij}^{e^\wedge}([R]) \cap \mathrm{Crs}_{e_j \to e^\wedge}(\mu_j)([R])) \supseteq \mu_i([r])$ for every $[r]$. Since Ref is monotone (preserves pointwise inclusion), $\mathrm{Ref}_{e_i \to e'_i}(\mu_i^+) \supseteq \mathrm{Ref}_{e_i \to e'_i}(\mu_i)$ pointwise. Both conditions of $\preceq_{\mathscr{U}}$ are satisfied. □

Lemma 3.6. *For fixed epistemic expansion rule and trigger threshold, if $\alpha_{ij}([r]) \subseteq \alpha'_{ij}([r])$ for all $[r]$ and neither credibility fires the trigger, then $T_+(u_i, u_j; \alpha_{ij}) \preceq_{\mathscr{U}} T_+(u_i, u_j; \alpha'_{ij})$.*

Proof. Fix the epistemic expansion rule (so e'_i is the same under both credibility functions). The meet-level filter satisfies $\alpha_{ij}^{e^\wedge}([R]) = \bigcap_{D \subseteq [R]} \alpha_{ij}(D) \subseteq \bigcap_{D \subseteq [R]} \alpha'_{ij}(D) = \alpha_{ij}'^{e^\wedge}([R])$, since the intersection of pointwise-larger sets is larger. Therefore $\alpha_{ij}^{e^\wedge}([R]) \cap \mathrm{Crs}_{e_j \to e^\wedge}(\mu_j)([R]) \subseteq \alpha_{ij}'^{e^\wedge}([R]) \cap \mathrm{Crs}_{e_j \to e^\wedge}(\mu_j)([R])$, giving $\mu_i^+([r]) \subseteq \mu_i^{+\prime}([r])$ for all $[r]$. The conclusion follows as in Lemma 3.5: Ref preserves pointwise inclusion. □

Lemma 3.7. *For fixed α_{ij}, the output $T_{\alpha_{ij}, \tau}(u_i, u_j)$ is weakly increasing in τ.*

Proof. Fix u_i, u_j, α_{ij} and let $n = |\mathsf{Conf}(u_i, u_j)|$. For $\tau \geq n$ the trigger does not fire: the output is $T_+(u_i, u_j) \succeq_{\mathscr{U}} u_i$ by Lemma 3.5. For $\tau < n$ the trigger fires: the output is $\beta(u_i, u_j) \preceq_{\mathscr{U}} u_i$ by Proposition 3.4 (strict when $\mathsf{Conf} \neq \emptyset$). Now let $\tau' > \tau$. If $\tau \geq n$, both produce T_+: equal. If $\tau < n \leq \tau'$, τ produces $\beta \preceq_{\mathscr{U}} u_i \preceq_{\mathscr{U}} T_+$: strictly higher under τ'. If $\tau' < n$, both produce β: equal. In every case the output under τ' is weakly above the output under τ. □

Lemma 3.8 (Sophistication paradox)**.** *Let $u_i \preceq_{\mathscr{U}} u'_i$ at the same epistemic level ($e_i = e_{i'}$, so $\mu_i([r]) \subseteq \mu_{i'}([r])$ for all $[r]$). Then $\mathsf{Conf}(u_i, u_j) \subseteq \mathsf{Conf}(u_{i'}, u_j)$: richer priors produce weakly larger conflict sets.*

Moreover, if τ is such that the trigger fires for u_i (i.e., $|\mathsf{Conf}(u_i, u_j)| > \tau$), then it also fires for $u_{i'}$. In that case, $\beta(u_{i'}, u_j) \prec_{\mathscr{U}} u_{i'}$ and $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$ by Proposition 3.4, and the richer voter's reversion strips weakly more content: the baseline must remove conflicting meanings at every class in Conf, and since $\mathsf{Conf}(u_{i'}, u_j) \supseteq \mathsf{Conf}(u_i, u_j)$, the richer voter has weakly more classes to evacuate.

Proof. Conflict monotonicity. Fix $[r] \in \mathsf{Conf}(u_i, u_j)$: by definition $\mu_i([r]) \cap F_{ij}([r]) \neq \emptyset$. Since $\mu_i([r]) \subseteq \mu_{i'}([r])$, we have $\mu_{i'}([r]) \cap F_{ij}([r]) \supseteq \mu_i([r]) \cap F_{ij}([r]) \neq \emptyset$, so $[r] \in \mathsf{Conf}(u_{i'}, u_j)$.

Trigger monotonicity. $|\mathsf{Conf}(u_{i'}, u_j)| \geq |\mathsf{Conf}(u_i, u_j)| > \tau$, so the trigger fires for $u_{i'}$. By Proposition 3.4, both $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$ and $\beta(u_{i'}, u_j) \prec_{\mathscr{U}} u_{i'}$. The baseline $\beta(u_{i'}, u_j)$ must remove conflicting meanings at every class in $\mathsf{Conf}(u_{i'}, u_j)$, which is a (weakly) larger set than $\mathsf{Conf}(u_i, u_j)$; the content stripped at each shared class is also weakly larger, since $\mu_{i'}([r]) \cap F_{ij}([r]) \supseteq \mu_i([r]) \cap F_{ij}([r])$. The richer voter thus reverts from a higher starting point with weakly more content stripped at weakly more representation classes. □

A.3 Proofs for Section 4.

Lemma 4.5. *Let $f : \mathscr{U}_{\mathcal{R}} \times \mathscr{U}_{\mathcal{R}} \to \mathscr{U}_{\mathcal{R}}$ satisfy:*

(a) $f(u, u') \preceq_{\mathscr{U}} u$ and $f(u, u') \preceq_{\mathscr{U}} u'$ for all u, u';

(b) $f(u, u') \succeq_{\mathscr{U}} v$ for every v satisfying $v \preceq_{\mathscr{U}} u$ and $v \preceq_{\mathscr{U}} u'$.

Then $f(u, u') = u \wedge u'$.

Proof. Condition (a) says $f(u, u')$ is a lower bound of $\{u, u'\}$ in $(\mathscr{U}_{\mathcal{R}}, \preceq_{\mathscr{U}})$. Condition (b) says $f(u, u')$ is at least as large as every lower bound—that is, it is the *greatest* lower bound. The greatest lower bound of $\{u, u'\}$ in a lattice is, by definition, the meet $u \wedge u'$. □

Proposition 4.12. *If a semantic election satisfies the four Downsification conditions and $\pi \circ \phi$ is injective (distinct platforms project to distinct policy positions), then:*

(i) the CMC holds at the median platform x_{med}, defined by $\pi(\phi(x_{\text{med}})) = p^*_{\text{med}}$ where p^*_{med} is the median of $\{\pi(u^*_v)\}_{v \in V}$;

(ii) *the symmetric profile* $(x_{\text{med}}, x_{\text{med}})$ *is a Nash equilibrium in the two-candidate case;*

(iii) *the election is Downsifiable, with recovered ordering* $\leq_X$ *induced by* $\pi \circ \phi$.

Proof. We verify each part in turn.

Step 1: Post-campaign understanding depends only on the policy projection. Under full policy credibility (Condition 1), the credibility filter passes exactly policy content: $\alpha_{vc}([r]) = M_{\text{policy}}$ for all $v, c, [r]$. With triggers deactivated (Condition 4), the forward update is always applied. With maximal epistemic expansion, voter v's post-campaign understanding $\hat{u}_v^c = T_+(u_v^0, \phi(x_c))$ absorbs all policy content from $\phi(x_c)$. The voter's welfare $\mathcal{W}(\hat{u}_v^c \wedge u_v^*)$ therefore depends on $\phi(x_c)$ only through the policy projection $\pi(\phi(x_c))$.

Step 2: Voter welfare is single-peaked by assumption. Condition 3 directly posits that $\mathcal{W}(\hat{u}_v^c \wedge u_v^*)$ is single-peaked in $\pi(\phi(x_c))$ with peak at $\pi(u_v^*)$.

Step 3: Black's theorem applies. The injectivity of $\pi \circ \phi$ gives a well-defined total order $\leq_X$ on X via $x \leq_X x'$ iff $\pi(\phi(x)) \leq \pi(\phi(x'))$. We now have a finite set X with a total order, $|V|$ voters with single-peaked preferences on this order (by Steps 1–2), and pairwise majority voting. By Black's theorem (Black, 1948), the platform x_{med} whose policy projection $\pi(\phi(x_{\text{med}}))$ equals the median of $\{\pi(u_v^*)\}_{v \in V}$ is the Condorcet winner: it defeats every other platform by simple majority.

Step 4: The CMC holds. By Step 3, x_{med} is the Condorcet winner with single-peaked voter welfare on $\leq_X$. By the backward direction of Theorem 4.11, the CMC holds at x_{med}.

Step 5: Nash equilibrium. In a two-candidate election, both candidates broadcasting at x_{med} constitutes a Nash equilibrium: any deviation by one candidate to $x' \neq x_{\text{med}}$ is defeated by the opponent at x_{med}, since x_{med} is the Condorcet winner and beats x' by majority. No profitable deviation exists. □

Theorem 4.11.

(i) *If* x^* *is Downsifiable then the CMC holds at* x^*.

(ii) *If the CMC holds at* x^* *then* $\phi(x^*)$ *is maximal among* $\{\phi(x) : x \in X\}$*: no platform's broadcast image is strictly richer in* $\preceq_{\mathscr{U}}$ *than the winning platform's.*

Proof. Part (i) (Downsifiable $\Rightarrow$ CMC). Suppose x^* satisfies (i)–(iii) of Definition 4.10 under $\leq_X$. By (i), $W(x^*) = X \setminus \{x^*\}$, so $|W(x^*)| = |X| - 1$.

By (ii), for any $x <_X x^*$, all voters with peaks at or above x^* prefer x^* to x (single-peakedness); these voters include the median and everything above it—a majority. So x^* defeats x, and similarly all platforms below x lose to everything between them and x^*: $|W(\cdot)|$ is non-decreasing up to x^*. By symmetry, $|W(\cdot)|$ is non-increasing above x^*. Hence $|W(\cdot)|$ is single-peaked at x^* and the CMC holds.

Part (ii) (CMC $\Rightarrow \phi(x^*)$ maximal). Suppose the CMC holds at x^*, so that $|W(x^*)| = |X| - 1$: x^* defeats every other platform by strict majority. Suppose for contradiction that $\phi(y) \succeq_{\mathscr{U}} \phi(x^*)$ with $\phi(y) \neq \phi(x^*)$ for some $y \in X$. By monotonicity of the meet, $\phi(y) \wedge u_v^* \succeq_{\mathscr{U}} \phi(x^*) \wedge u_v^*$ for every voter v. Strict monotonicity of $\mathcal{W}$ then gives $w_v(y) \geq w_v(x^*)$ for every v, so $\{v : w_v(x^*) > w_v(y)\} = \emptyset$. But x^* must defeat y by strict majority, requiring $|\{v : w_v(x^*) > w_v(y)\}| > |V|/2$—a contradiction. □

A.4 Proofs for Section 5.

Proposition 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 (1955) 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. □

Lemma 5.6. *No strictly monotone lattice homomorphism $\pi : (\mathscr{U}_{\mathcal{R}}, \preceq_{\mathscr{U}}) \to (\mathbb{R}, \leq)$ exists when the voter ideal poset $(\{u_v^*\}, \preceq_{\mathscr{U}})$ has width exceeding 1.*

Proof. Let u_1^*, u_2^* be incomparable voter ideals (an antichain of width ≥ 2). Since $u_1^* \parallel u_2^*$, neither is below the other, so $u_1^* \wedge u_2^* \prec_{\mathscr{U}} u_1^*$ and $u_1^* \wedge u_2^* \prec_{\mathscr{U}} u_2^*$. A lattice homomorphism π preserves meets: $\pi(u_1^* \wedge u_2^*) = \min(\pi(u_1^*), \pi(u_2^*))$. Without loss of generality $\pi(u_1^*) \leq \pi(u_2^*)$, so $\pi(u_1^* \wedge u_2^*) = \pi(u_1^*)$. But strict monotonicity requires $\pi(u_1^* \wedge u_2^*) < \pi(u_1^*)$ (since $u_1^* \wedge u_2^* \prec_{\mathscr{U}} u_1^*$), a contradiction. □

Proposition 5.7. *If the semantic dimension of V exceeds 1, then Condition 3 (single-peaked welfare) cannot be derived from the lattice structure of $\mathscr{U}_{\mathcal{R}}$: the only structural mechanism that would guarantee it—a strictly monotone lattice homomorphism $\pi : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$—does not exist (Lemma 5.6). The CMC is therefore not structurally guaranteed when semantic dimension exceeds 1.*

Proof. Semantic dimension exceeding 1 means the voter ideal poset contains an antichain of width ≥ 2. By Lemma 5.6, no strictly monotone lattice homomorphism $\pi : \mathscr{U}_{\mathcal{R}} \to \mathbb{R}$ exists. A meet-preserving π is the structural mechanism by which Condition 3 follows from the lattice geometry: when $\pi(u \wedge u') = \min(\pi(u), \pi(u'))$, the welfare formula $\mathcal{W}(\phi(x) \wedge u_v^*)$ inherits single-peakedness from the linear order on $\pi(\phi(X))$. Without a meet-preserving projection, Condition 3 has no structural derivation from the lattice and can hold only by coincidence of the welfare function and the platform set. The CMC is therefore not structurally guaranteed. □

Theorem 5.9. *When the semantic dimension of V exceeds 1, the welfare-relevant McKelvey majority-preference orbit is confined to the* semantic interval

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

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. For any platform $\phi(x) \preceq_{\mathscr{U}} \bigwedge_v u_v^*$, we have $\phi(x) \wedge u_v^* = \phi(x)$ for every voter v (since $\phi(x) \preceq_{\mathscr{U}} u_v^*$), so all voters evaluate sub-meet platforms by the same quantity $\mathcal{W}(\phi(x))$; strict monotonicity of $\mathcal{W}$ then gives a unanimous ranking among them, and no majority cycling occurs. Neither class—above-join or sub-meet—participates in the majority tournament; the welfare-relevant orbit is therefore contained in the semantic interval. □

Proposition 5.11. *When credibility filters are pairwise disjoint and all triggers are deactivated ($\tau_v = \infty$ for all $v \in V$), 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 5.12. *In a two-candidate race with pairwise disjoint voter types, pairwise incomparable type ideals ($u^*_{T_i} \parallel u^*_{T_j}$ for $i \neq j$), and deactivated triggers ($\tau_v = \infty$ for all v), 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 and voter type ideals are semantically incomparable.*

Proof. By Proposition 5.11, vote-share optimization decomposes by type: the candidate independently optimizes the α_i-filtered component of the broadcast for each type T_i. Since the type ideals are pairwise incomparable, for each pair (T_i, T_j) there exist meanings in $u^*_{T_i} \setminus u^*_{T_j}$ and in $u^*_{T_j} \setminus u^*_{T_i}$—semantic content that one type values and the other does not. Because $\mathcal{W}$ is strictly monotone, the optimal α_i-component for type T_i maximizes common ground with $u^*_{T_i}$, and therefore includes meanings in $u^*_{T_i} \setminus u^*_{T_j}$ whenever the broadcast map ϕ can deliver them. The optimal components for distinct types are therefore distinct—they contain different semantic content—and the broadcast is polysemic. Any deviation that unifies the message sacrifices type-specific optimization and reduces vote share in at least one type. □

Proposition 5.14. *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. Since voter ideals form a 2-element antichain, every voter's ideal is either u^*_A or u^*_B; hence $V_A = \{v : u^*_v = u^*_A\}$. 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 $\mathcal{W}(\phi(x^*) \wedge u^*_v) = \mathcal{W}(\phi(x^*) \wedge u^*_A) < \mathcal{W}(u^*_A)$ 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 5.15. *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 5.14): 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.*

Proof. By Proposition 5.14, $\phi(x^*) \succeq_{\mathscr{U}} u_A^*$. Suppose $\phi(x^*) \not\succeq_{\mathscr{U}} u_A^* \vee u_B^*$. Since $\phi(x^*) \succeq_{\mathscr{U}} u_A^*$ and the join $u_A^* \vee u_B^*$ is the least upper bound, the only way $\phi(x^*)$ fails to dominate the join is if $\phi(x^*) \not\succeq_{\mathscr{U}} u_B^*$. Then $\phi(x^*) \wedge u_B^* \prec_{\mathscr{U}} u_B^*$ (the meet is strictly below u_B^*, since $\phi(x^*) \wedge u_B^* = u_B^*$ would imply $u_B^* \preceq_{\mathscr{U}} \phi(x^*)$, a contradiction). By strict monotonicity of $\mathcal{W}$: $\mathcal{W}(\phi(x^*) \wedge u_B^*) < \mathcal{W}(u_B^*)$. For the bound on the deficit: $\phi(x^*) \succeq_{\mathscr{U}} u_A^*$ implies $\phi(x^*) \wedge u_B^* \succeq_{\mathscr{U}} u_A^* \wedge u_B^*$ (meets are monotone), so the deficit satisfies $\mathcal{W}(u_B^*) - \mathcal{W}(\phi(x^*) \wedge u_B^*) \leq \mathcal{W}(u_B^*) - \mathcal{W}(u_A^* \wedge u_B^*)$. □

A.5 Proofs for Section 7.

Proposition 7.1. *The baseline operator β satisfies the semantic analogs of postulates* (÷1)–(÷5) *and fails the semantic analog of* (÷6) *(Recovery).*

Proof. We verify each postulate in turn.

Closure (÷1): $\beta(u_i, u_j) = \bigvee S_{ij}$ is an element of $\mathscr{U}_{\mathcal{R}}$, since $\mathscr{U}_{\mathcal{R}}$ is a complete lattice (Theorem 2.9) and the join of any subset exists.

Inclusion (÷2): $\beta(u_i, u_j) \preceq_{\mathscr{U}} u_i$ by construction—S_{ij} contains only elements below u_i, and the join of elements below u_i is below u_i.

Vacuity (÷3): If $\mathsf{Conf}(u_i, u_j) = \emptyset$, then u_i has no conflicting content with u_j at any representation class. In this case $u_i \in S_{ij}$ (it vacuously satisfies the avoidance condition), so $\bigvee S_{ij} \succeq_{\mathscr{U}} u_i$; combined with inclusion, $\beta(u_i, u_j) = u_i$.

Success (÷4): If $\mathsf{Conf}(u_i, u_j) \neq \emptyset$, then by Proposition 3.4, $\beta(u_i, u_j) \prec_{\mathscr{U}} u_i$—the conflicting content has been strictly removed. Moreover, the avoidance condition in the definition of S_{ij} ensures that $\beta(u_i, u_j)$ contains no meanings from $F_{ij}([r])$ at any conflicting representation class $[r] \in \mathsf{Conf}(u_i, u_j)$.

Extensionality (÷5): The baseline depends on u_j only through the conflict set $\mathsf{Conf}(u_i, u_j)$ and the conflict meanings $F_{ij}([r])$ at each conflicting class. If two source understandings u_j, u_k induce identical conflict structures—$\mathsf{Conf}(u_i, u_j) = \mathsf{Conf}(u_i, u_k)$ and $F_{ij}([r]) = F_{ik}([r])$ for all $[r]$—then $S_{ij} = S_{ik}$ and $\beta(u_i, u_j) = \beta(u_i, u_k)$.

Recovery (÷6): This postulate requires $u_i \preceq_{\mathscr{U}} T_+(\beta(u_i, u_j), u_j)$—that re-encountering u_j after contraction restores the original understanding. Recovery fails in general; the three sources of failure (credibility filtering, re-triggering, and epistemic collapse) are discussed in the main text. □

Proposition 7.2. *Set $M = \{T, F\}$ (Boolean meanings), fix a common epistemic level $e_\top$ (all representations visible, each in its own class), define $\mathsf{Cf} = \{(T, F), (F, T)\}$, and let each voter have full credibility $\alpha_v([r]) = M$ for all $[r]$, deactivated trigger $\tau_v = \infty$, maximal epistemic expansion, and* *the* replacement interpretation *of absorption (when the voter absorbs content that conflicts with her prior assignment, the new content supersedes the old rather than accumulating alongside it). Let the broadcast* $\phi(x_c)$ *assign* T *to each representation class the announcement endorses and* F *to each it denies. Then the forward update* T_+ *reduces to public-announcement-style world-elimination: the voter's post-update semantics assigns* T *to every class where the announcement asserts* T*,* F *to every class where it asserts* F*, and the resulting understanding contains exactly the semantic content consistent with the announcement.*

Proof. Under full credibility ($\alpha_v([r]) = M$) and deactivated trigger ($\tau_v = \infty$), the forward update at each representation class $[r]$ is

$$\mu_v^+([r]) \;=\; \mu_v([r]) \;\cup\; \big(\alpha_v([r]) \cap \mu_c([r])\big) \;=\; \mu_v([r]) \;\cup\; \mu_c([r]).$$

With $M = \{T, F\}$, if the voter previously held $\mu_v([r]) = \{F\}$ and the announcement asserts $\mu_c([r]) = \{T\}$, the post-update semantics is $\{T, F\}$. Under the replacement interpretation, this reduces to $\{T\}$: the announcement eliminates the prior's false assignment, reproducing PAL's world-elimination. More generally, for any $[r]$: if $\mu_c([r]) = \{T\}$, then $\mu_v^+([r])$ contains T, and replacement removes F wherever it conflicts with T; if $\mu_c([r]) = \{F\}$, the symmetric argument applies. The post-update understanding contains exactly the semantic content consistent with the announcement. □