Diverging Paths to Dark-Matter Discovery

Simon Allzén[1,2,3, ‡]

[1]Department of Philosophy, Stockholm University
[2]Institute of Physics, University of Amsterdam
[3]Vossius Center for the History of Humanities and Sciences, University of Amsterdam
[‡]Correspondence: simon.allzen@philosophy.su.se

Abstract

Dark-matter research brings together two distinct epistemic contexts: laboratory searches for new particles and astronomical inferences about a non-luminous gravitating component. The justification for treating these pursuits as complementary is that they can establish different aspects of the same entity. The weakly interacting massive particle (WIMP) made this expectation scientifically credible by connecting a particle candidate to a cosmological abundance, experimental searches, and explanations of astronomical phenomena. What happens to this expectation as the WIMP loses prominence? I analyze 201,627 NASA Astrophysics Data System records to examine changes in research organization as WIMP prominence declines. Astrophysics and high-energy physics remain jointly dominant, while WIMP-related families lose share of tracked candidate mentions in both fields and successor candidates receive different degrees of attention. Concept networks show that WIMP-related research remains central across the literature, even as cross-field connections become less concentrated around it. These indicators measure organized research attention and cannot establish what scientists would accept as a discovery. I interpret the findings through a framework distinguishing constitutive identification, cosmic attribution, explanatory adequacy, and the cross-domain identification needed to connect them. A particle, compact-object population, or gravitational theory can resolve different parts of the inherited problem, and a partial or multicomponent discovery can leave other parts open. 'Diverging paths' describes this branching of possible resolutions.

Keywords: dark matter; WIMP; astroparticle physics; bibliometrics; scientific discovery

"Superiora de inferioribus, inferiora de superioribus"
Tabula Smaragdina —

1 A particle panacea

It's possible for particle physics to identify a new particle without establishing its role on astronomical scales, just as astronomy can infer the existence, distribution and behavior of gravitating matter without identifying its microscopic constitution. Sufficient evidence in one context is not necessarily sufficient evidence in the other. During much of the late twentieth century, however, dark-matter research largely treated the contexts as coupled pursuits, directed toward *the discovery of the same object*. The weakly interacting massive particle (WIMP) provided scientific credibility to this coupling: WIMPs could fulfill the cosmological requirement for cold, non-baryonic matter, and weak-scale extensions of the Standard Model (SM) supplied stable relics whose freeze-out abundance could be of the order required for the astronomically inferred dark component [10–14].

The appeal of the WIMP consisted in its ability to connect questions that neither astronomy nor particle physics could answer independently. Astronomical observations established mass discrepancies and constrained the distribution of unseen gravitating matter [3–5]. Cosmology connected a non-baryonic component to relic abundance and structure formation, while particle physics supplied stable weak-scale candidates and methods for detecting their non-gravitational interactions [8, 9]. A successful WIMP program could establish with high confidence that a *particle identified experimentally* was the same ontological object responsible for the dynamics which had been inferred on galactic and cosmological scales. Contemporary reviews presented relic abundance, halo properties, and direct and indirect detection as connected parts of this program [13, 14]. Alas, direct and indirect detection searches have for the past 30 or so years progressively restricted the parameter space of a canonical WIMP, without conclusive results supporting its existence [15–17]. The relative prominence of the WIMP in dark-matter research has consequently declined. The bibliometric analysis in this paper documents declining shares of WIMP-related candidate mentions alongside rising shares for axions, axion-like particles (ALPs), and primordial black holes (PBHs) (Fig. 3). Reviews describe an expansion and diversification of theory-space that also includes dark sectors and modified-gravity approaches [18, 27]. The question is whether a successor can mediate between microphysical and macrophysical evidence in the way that made the WIMP so attractive across disciplines.

The prospect of discovering a single object which was equally significant in both particle physics and astrophysics was formative for the organization of dark-matter research. Characterizing dark matter as a particle gave particle physics a substantial role in a problem whose evidential basis was, and remains, astronomical and cosmological. In the bibliographic data, particle phenomenology remains a principal location of dark-matter research, while astrophysics remains dominant when its classifications are aggregated. Papers cross-listed between astrophysics and high-energy physics remain a substantial part of the literature. Dark-matter research thus remains a shared venture even as the candidate that helped connect its pursuits has become less prominent. Does that persistence still give us reason to expect that the evidence provided by each participating program will coalesce into a single discovery?

Existing philosophical work examines the justification and relationship of these pursuits. De Baerdemaeker's account of method-driven experiments explains why the limited gravitational description of the target must be supplemented with assumptions about its properties before particle-search methods can be justified [46]. Her later distinction between astrophysical and fundamental dark-matter concepts separates the component characterized in cosmological and astronomical models from candidate descriptions that add assumptions about its constitution and interactions [47]. Martens argues that detection methods are *epistemically complementary* because they establish different things about dark matter, while emphasizing that detecting its existence and identifying its nature are semantically intertwined [48]. The question I explore concerns how a particular candidate made the prospect of connecting these achievements scientifically credible, and what happens to that expectation when the participating disciplines give different degrees of attention to candidates with different constitutions and conditions for discovery.

After reconstructing the WIMP's role as an evidential mediator, I examine five bibliometric regularities in an ADS dataset covering research records published through 2025: persistent cross-disciplinary participation, declining WIMP prominence, field-relative successor prominence, no increase in divergence across the tracked 13-family candidate-mention profiles, and a shift from WIMP-centered to more distributed cross-field connectivity.[1] I then compare what different candidate programs could establish, distinguishing constitutive identification, cosmic attribution, explanatory adequacy, and the cross-domain identification that connects

[1]These indicators measure organized research attention. They do not *directly* measure epistemic complementarity, discovery standards, or communal judgments.

them. I argue that the WIMP provided a historically influential and scientifically motivated expectation that different investigations could converge on one identification, while its declining prominence undermines that expectation without removing successor programs' capacity to establish and connect the relevant achievements. Continued participation and overlapping vocabularies do not in and of themselves justify treating locally decisive results as parts of a single discovery. These distinctions explain why resolving a question within one program can leave the jointly occupied problem open, and why evidential closure can differ from communal acceptance.

2 The WIMP: an evidential mediator

Dark matter was initially invoked to account for several observational anomalies. Historical studies show that a coherent missing-mass hypothesis became viable and significant only after several dynamical problems were brought together in a cosmological context, long before the constitution of that mass was coupled to particle candidates [19, 20]. Astronomical observations supported the existence of a component that supplied additional gravitating mass, had an appropriate abundance and distribution, and could contribute to the formation of cosmic structure. The evidence characterized this component through the dynamical, cosmological, and explanatory roles it was required to fulfill. Support for *the inference* could therefore become increasingly robust while *the nature* of the inferred component remained unresolved and ontologically flexible. The increasing focus on dark matter *qua* particle has directed attention beyond its evidential foundations in a historically heterogeneous set of astronomical observations toward speculative accounts of its microphysical constitution within physics beyond the Standard Model.[2]

This 'particle turn' had a prehistory. In the mid-1960s, Gershtein and Zel'dovich used cosmological density bounds to constrain neutrino masses, showing that particle properties could be directly relevant to cosmic dynamics [1]. Cowsik and McClelland later proposed that massive neutrinos might dominate the cosmic mass density and therefore account for inferred missing mass in clusters [2]. The significance of this step was that the relationship between the fields acquired an explanatory direction as well as a constraining one. Where cosmological considerations had placed *limits on neutrino masses*, a neutrino population with suitable properties could now be proposed as the constituent of the missing mass. Because the proposal gave particle physics a possible role in explaining an astronomical problem while making cosmological abundance a requirement it had to satisfy, the fields could contribute to assessing a common candidate through their different theoretical and empirical resources.

As rotation-curve and mass-to-light evidence connected galactic and cluster discrepancies to the cosmic matter budget [3–5], particle candidates became subject to both astronomical and cosmological constraints. Tremaine and Gunn derived a phase-space bound limiting the capacity of ordinary stable neutrinos to constitute galactic halos [6], while Peebles investigated structure formation in a universe containing massive neutrinos [7]. Although these contributions did not necessarily identify dark matter as neutrinos, they helped set the stage for missing mass turning into a problem whose possible solutions could be characterized and constrained through particle properties. The connection with particle physics is made unusually explicit in the proceedings of the 117th Symposium of the International Astronomical Union, held in 1985.[3] Peebles, chairing a discussion after a talk by Scott Tremaine, asks what epistemic status astronomers should grant these ideas:

[2]The situation has a particular impact in the debate on scientific realism: see [52], [53], and [50, 51].

[3]The symposium, 'Dark Matter in the Universe' [44], included talks by Jim Peebles, Vera Rubin, Ken Freeman, Mordehai Milgrom, Jacob Bekenstein, Jaan Einasto, Robert Sanders, John Bahcall, Joseph Silk, Michael Turner, Sandra Faber, Brent Tully, Jerry Ostriker, Scott Tremaine, and Jim Gunn, all of whom have contributed significantly to shaping dark-matter research.

> New ideas from particle physics have greatly stimulated some old speculations in cosmology. Inflation provides a beautiful explanation for the isotropy of the Universe and lends respectability to prejudices against cosmological models with non-zero space curvature. Phase transitions may produce computable density fluctuations that end up as black holes or superclusters. *And we are presented with a rich catalogue of forms of matter, from gravitational waves to cosmic strings, whose abundances may be computable and whose physics may help to account for the properties of galaxies and superclusters.* But apart from the general enthusiasm, are we astronomers justified in accepting any of this as received knowledge? [45] [My emphasis]

Peebles's remark suggests that particle physics was viewed, at least in part, as a way to revitalize astronomical and cosmological theory. In the same discussion, Tremaine asks:

> [. . .] *what* is the dark matter? My understanding of the consensus is that we certainly need some baryonic dark matter, given that there is a problem of missing mass in the galactic disk. We don't *need* non-baryonic dark matter, although it is very attractive.

Here, Tremaine treats *baryonic* dark matter as necessary for the missing-mass problem in the galactic disk, while *non-baryonic* dark matter remains an attractive, but perhaps a bit exotic, possibility. Characterizing dark matter as a non-baryonic particle was evidently not yet a settled premise of the discussion. Terrestrial detection proposals added additional connections. Following Drukier and Stodolsky's proposal for a neutral-current detector with applications to neutrino physics and astronomy [8], Goodman and Witten applied terrestrial detection methods to candidate particles in the galactic halo [9]. Supersymmetric theories, meanwhile, supplied stable weak-scale relics motivated independently of the astronomical anomalies [10–12]. Astronomy and cosmology thus found possible explanations for missing mass and structure formation in particle physics, while particle physics found astronomical and cosmological implications for its theories. This connection made the constitution of dark matter appear potentially *knowable* by making it potentially *detectable*. The WIMP became the mediator of all these possibilities: cosmology specified an abundance, particle theory supplied viable constituents, and direct detection offered a non-gravitational test for the presence of WIMPs in the Milky Way's halo [13, 14]. The question was no longer restricted to whether unseen matter was required, but had expanded to which microscopic entity filled the inferred role. As constitutive identification became a further ambition of dark-matter research, a discovery could be expected to establish both *the existence* of additional gravitating matter and *its nature*.

The benefits of the expanded ambitions of dark matter research were mutual. For astronomy and cosmology, a particle candidate promised to replace an inferential placeholder with an entity whose stability and interactions could be physically specified. Direct detection offered a way to seek evidence independent of the gravitational phenomena that had motivated the problem, and to investigate the same entity in a controlled terrestrial setting [8, 9]. For particle physics, dark matter supplied theories beyond the Standard Model with a cosmic role already supported on independent empirical grounds. Cosmological abundance and structure formation constrained candidate properties, while the local halo provided an experimental target. Supersymmetric relics thereby acquired a rationale beyond their original theoretical motivations: they could constitute the matter that empirical evidence in astronomy and cosmology had already suggested should exist [10–13].

Dark matter gave new microphysics an empirical purpose, while new microphysics promised to determine the nature of dark matter. These reciprocal expectations were connected through

a division of epistemic labor centered on the WIMP hypothesis, with the participating fields differing in their immediate objects, vocabularies, and evidential responsibilities. While astronomy established the phenomena and constrained the distribution of the gravitating component, cosmology constrained its abundance, production history, and role in structure formation, and particle physics proposed its constitution and developed detection strategies. The contributions could be considered *complementary* because there was a reason to regard them as different modes of investigation of *the same entity*, despite their different methods, scales, and forms of evidence.

None of these fields could, however, independently provide the evidence sufficient to declare a discovery of dark matter. A relic-density calculation did not establish that WIMPs existed, and a laboratory excess would not, by itself, show that the detected particle accounted for the cosmic abundance. Nor could an astronomical fit determine a unique microscopic constitution. The WIMP was a common candidate whose proposed properties connected the evidential contributions of astronomy, cosmology, and particle physics. It thereby served as an *evidential mediator*, making the prospect of field-relative results converging on a single identification scientifically credible.

3 A jointly occupied problem

I use the arXiv classifications attached to the ADS corpus to establish how this division of labor is reflected in the literature. Of 201,627 unique records published no later than 2025, 143,390 have between one and nine arXiv classes.[4] The two tables summarize the leading primary assignments across this full period, with percentages calculated against all classified records. Aggregating primary arXiv classes into their parent categories gives astrophysics a larger share than high-energy physics across the full period (Table 1). This also holds in the annual comparison throughout 2010–2025, with astrophysics accounting in 2025 for 48.7% of classified papers, compared with 26.8% for high-energy physics. At the level of individual primary classes, however, *hep-ph* first exceeds *astro-ph.CO* in 2015 (Fig. 1). This crossover concerns two classes and does not establish that high-energy physics displaced astrophysics as a whole. Across the entire period, *hep-ph* is the single most frequently assigned primary class, accounting for 21.3% of classified papers, followed by *astro-ph.CO* at 19.9% (Table 2). Together, these two classes account for 41.2% of classified papers, and the gap between them is 2,068 papers. Particle phenomenology has therefore anchored itself as a prominent locus of dark-matter research, situated within a literature that remains predominantly astrophysical. The supplementary classification analysis tests whether the broad-field comparison depends on how multi-classified papers are counted. Astrophysics remains larger under both any-class and fractional assignments, although its estimated share

Tbl 1: Papers by primary arXiv categories

arXiv category	Total	Percent
astrophysics	78,673	54.9%
high-energy physics	39,191	27.3%
gr-qc	14,668	10.2%
physics	6,427	4.5%

Tbl 2: Papers by primary arXiv classes

arXiv classification	Total	Percent
hep-ph	30,540	21.3%
astro-ph.CO	28,472	19.9%
astro-ph	21,450	15.0%
astro-ph.GA	19,731	13.8%
gr-qc	14,668	10.2%
astro-ph.HE	5,456	3.8%
hep-th	4,777	3.3%
hep-ex	3,381	2.4%
physics.ins-det	2,677	1.9%

[4]The metadata are used to locate research within a publication and classification structure at a large scale, which should not be taken to imply homogeneity with respect to the aims and commitments of individual researchers.

changes with the counting rule.[5] Between 16.8% and 20.2% of classified papers each year from 2010 through 2025 belong to both fields under any-class membership. Dark matter thus remains a *jointly occupied problem*: despite the differentiation within the literature, several fields retain a place in formulating the research and contributing to its resolution.

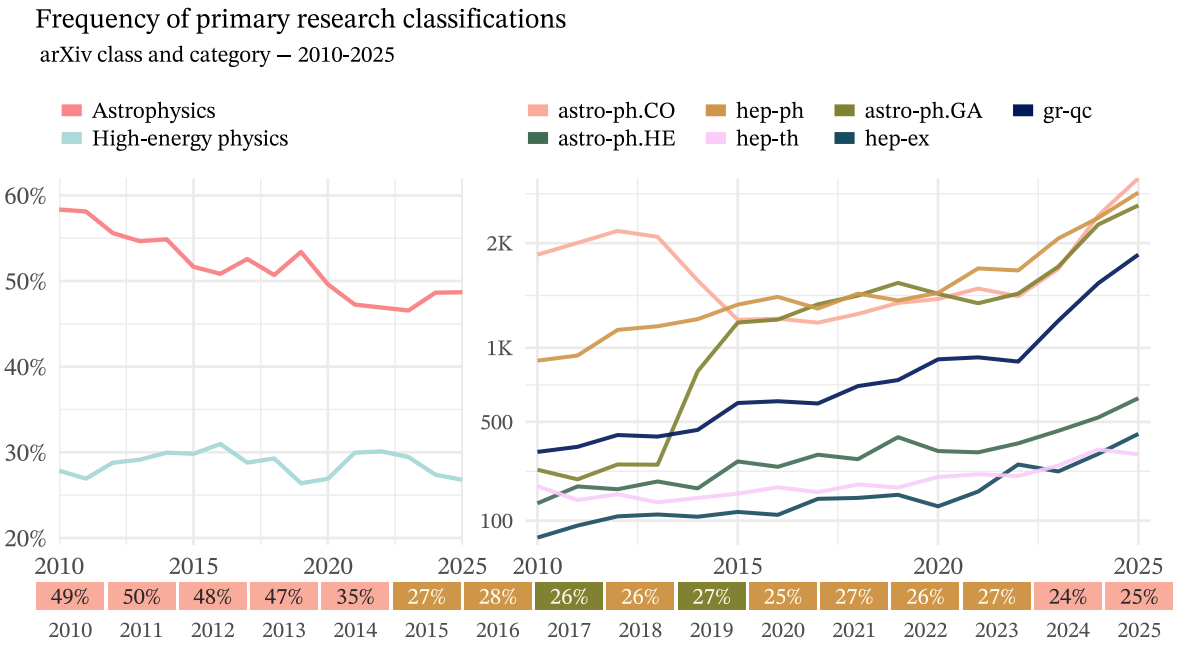


Figure 1: The disciplinary location of dark-matter research. **Left:** Annual shares of classified papers under primary assignments, aggregated into broad fields, with astrophysics remaining larger than high-energy physics through 2025. **Right:** Annual paper counts on a square-root vertical scale for the seven primary arXiv classes with the largest cumulative counts in 2010–2025. Among these, *hep-ph* first exceeds the single class *astro-ph.CO* in 2015, while the class ordering varies thereafter. The strip beneath the right panel identifies the leading displayed class each year and its share of papers across those seven classes.

4 Changing candidate prominence and connectivity

The fact that several disciplines share a research problem does not entail that they share a route to resolving that problem. Candidate families receive unequal attention in the two broad literatures, as Fig. 2 shows by comparing each family's share of tracked mentions in high-energy physics with its share in astrophysics. Axion and several supersymmetry-related families are relatively more prominent in high-energy physics, while PBHs, massive compact halo objects (MACHOs), and ultralight or fuzzy candidates are relatively more prominent in astrophysics. These alignments concern a family's place in each field's candidate mixture: a family can be studied in both fields without being equally prominent in them.

A mention count does not establish whether an author accepts a candidate, since papers can introduce, compare, constrain, or exclude the candidates they name. With this limitation in view, the within-field composition shows a clear temporal pattern: WIMP-related families lose share of tracked mentions in both literatures (Fig. 3). Axion and ALP research expands especially in high-energy physics, where haloscope and helioscope programs provide

[5] A fixed-seed, year-stratified audit of 500 multi-classified papers found 500/500 agreement between the first ADS class and the arXiv primary category (agreement 1.000, Wilson 95% interval 0.992–1.000). While this audit supports the primary-order convention for the sampled records, the alternative assignments show how the estimated shares depend on the counting rule.

experimental avenues [21–24]. PBH research expands especially in astrophysics, where lensing, compact-object dynamics, and gravitational-wave observations constrain candidate populations [25–27]. Research attention has thus shifted toward different successor mixtures in the two fields, although neither mixture is exclusive to one field. These different mixtures do not establish increasing divergence between the 13-family candidate-mention profiles. None of the Jensen–Shannon divergence, total-variation distance, or cosine distance summaries supports an increase over time. In the predefined comparison,
mean Jensen–Shannon divergence is 0.326 in 1995–2004 and 0.293 in 2016–2025, giving a late-minus-early contrast of −0.033. The early annual values also rest on far fewer candidate-mentioning papers. A separate lexical analysis shows rising high-energy-physics terms such as *mediator* and *portal*, but its divergence magnitude depends more strongly on preprocessing and term selection. I therefore retain lexical divergence after the predefined candidate-name exclusions as a supplementary robustness check.

The absence of increasing divergence sharpens the conceptual question because similarity in the distribution of attention does not establish complementarity in the production of evidence. Two fields can study much the same candidates without producing results that support one discovery claim, while a productive division of labor can equally involve different distributions. 'Diverging paths' therefore concerns what the programs could establish and how their achievements could be connected. The findings locate this question within a still-shared problem: the fields continue to participate and share candidates, but have not converged on a replacement with WIMP-comparable prominence.

Declining mention share does not by itself determine a candidate's place in the wider conceptual organization of the literature, so I distinguish centrality across the literature from the position connecting the two fields. In the pooled candidate–concept network, the WIMP-related umbrella remains the highest-ranked candidate by prevalence-weighted role-specific

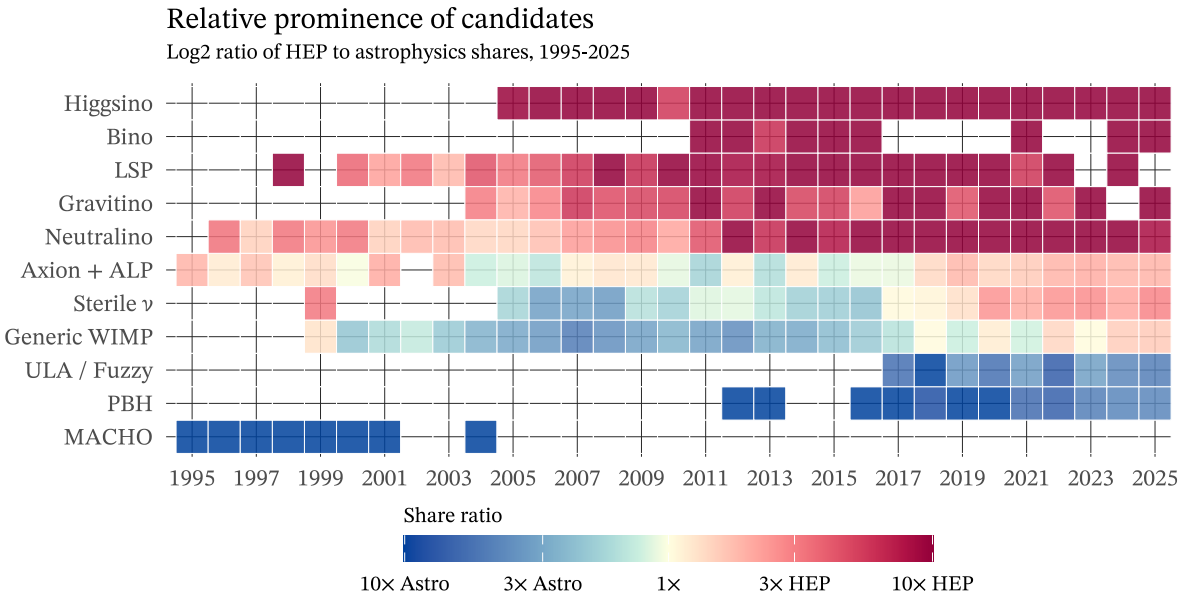


Figure 2: **Candidate prominence differs between astrophysics and high-energy physics.** Each cell shows the base-2 logarithm of a candidate family's share of tracked mentions in high-energy physics divided by its share in astrophysics for that year. Positive values indicate greater relative prominence in high-energy physics, negative values greater relative prominence in astrophysics, and zero parity. Candidate-years with fewer than 20 total mentions across the two fields are left blank. The threshold affects display only, not the underlying estimates or candidate-composition summaries. Ratios compare relative shares rather than raw publication volumes, and mentions include comparisons and exclusions rather than endorsements.

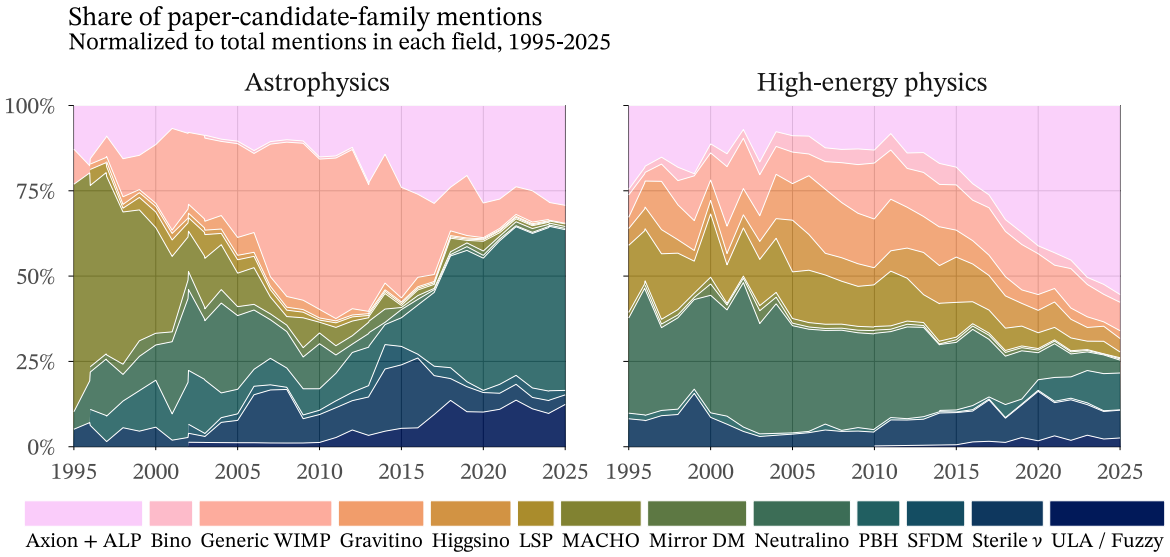


Figure 3: Candidate composition follows different field-specific trajectories. Annual paper–candidate-family mentions are normalized within astrophysics and high-energy physics. A paper can mention more than one family, and mentions do not imply endorsement, only prevalence. Raw counts, equal-paper weighting, and formal divergence summaries are reported in the supplementary information.

spectral centrality, calculated with the HITS algorithm, in both 2006–2015 and 2016–2025. Its declining share of mentions has thus not made it marginal to the discourse as a whole.

The WIMP-related umbrella's position between the fields has nevertheless become less singular: in the field-oriented concept network, it has the highest normalized cross-field betweenness in 2006–2015. Under the primary published-umbrella definition, it ranks second among candidates in 2016–2025, behind Axion + ALP, and its normalized betweenness falls from 0.073 to 0.031 (Fig. 4). Although the exact late-period rank varies across three WIMP definitions, all three yield lower betweenness and removal effects in the later period. In the node-removal analysis, removing the WIMP-related node reduces cross-field efficiency by 41.5% in the earlier network and 6.2% in the later one, where Axion + ALP and PBH have larger effects. Results for the predefined rolling five-year windows show the same broad decline in the concentration of connections around the WIMP.

These measures establish that cross-field discourse depends less on the WIMP-related node, but they do not establish evidential integration or a historical effect of the WIMP's absence. To assess the prospect of integration, we must turn from the organization of attention to the achievements available to each candidate program.

5 Relative paths to discovery

What, then, could the different programs establish? I distinguish four achievements without assuming at the outset that the evidence for each concerns the same entity. Let c_I, c_A, and c_E denote the referents characterized in the constitutive, attribution, and explanatory settings.

Constitutive identification, $I(c_I)$, establishes what kind of particle, field, object, population, or physical structure has been identified. *Cosmic attribution*, $A(c_A)$, establishes that a component or population has the abundance, persistence, and distribution required of the cosmic dark component. *Explanatory adequacy*, $E(c_E)$, establishes that the proposed explanation accounts for the relevant astronomical and cosmological phenomena at the required scales. These achievements can each be warranted while their connection remains open. *Cross-domain*

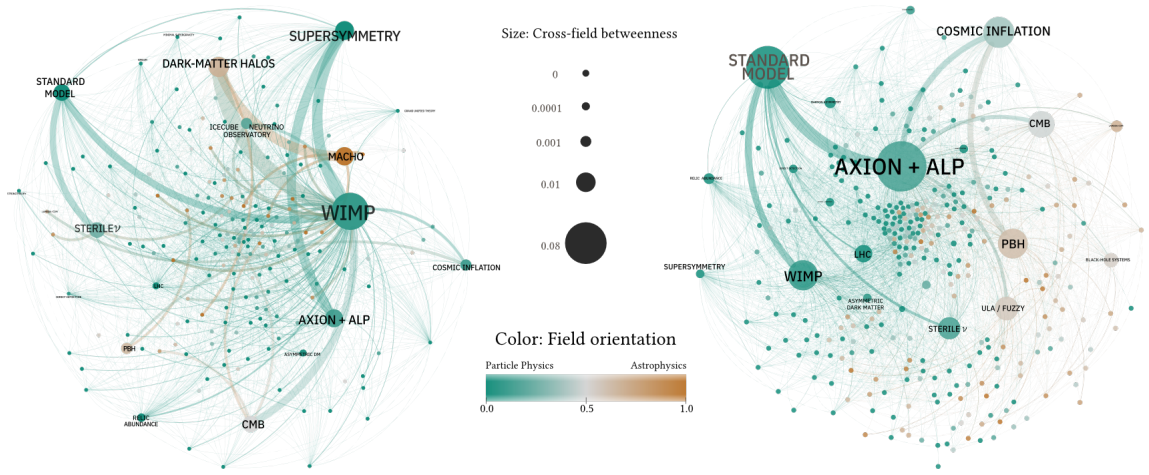


Figure 4: **The WIMP remains cross-field connective, but its position becomes less singular.** Nodes are normalized concepts identified in abstracts, with edges connecting concepts mentioned in the same paper and weighted by their mean co-occurrence prevalence across astrophysics and high-energy physics. Node size encodes normalized weighted betweenness between concepts at least twice as prevalent in one field as in the other. Color encodes relative field prevalence, from high-energy physics (teal) through balanced prevalence (gray) to astrophysics (ochre). The cyan outline identifies the WIMP-related umbrella, which combines generic WIMP, neutralino, LSP, Higgsino, gravitino, and sneutrino mentions. Under this published-umbrella definition, the WIMP-related node ranks first among candidates in 2006–2015 and second, behind Axion + ALP, in 2016–2025, while its normalized cross-field betweenness falls from 0.073 to 0.031. The network measures discursive connection, not evidential complementarity.

identification, $L(c_I, c_A, c_E)$, supplies that connection by warranting that the separately characterized referents are related in the way required for one claim about a common candidate. A strong dark-matter discovery claim therefore has the structure

$$D_{\mathrm{DM}} = I(c_I) \wedge A(c_A) \wedge E(c_E) \wedge L(c_I, c_A, c_E) \tag{1}$$

The expression represents the strong claim that the linked referents identify a common candidate c as the relevant cosmic dark component, rather than defining discovery of c in general. If c supplies only a fraction f of the abundance, $A_f(c)$ records that partial attribution. A multicomponent claim instead replaces c with a set $C = \{c_1, \ldots, c_n\}$ and requires the attributed fractions to supply the relevant abundance and explanatory roles together. Discovering one member of C can be a genuine discovery of a particle or component while leaving the remaining dark-matter problem unresolved.

The achievements represented by the framework need not occur simultaneously, within one discipline, or in a fixed order. Nor should $L(c_I, c_A, c_E)$ be understood as a fourth independent observation, since it concerns the warrant for connecting the laboratory entity, the attributed cosmic population, and the proposed explanation in the way required for one dark-matter claim. Suppose laboratory evidence identifies a particle with specified properties and cosmological evidence independently supports a population compatible with those properties. Reason to accept each result need not supply reason to accept that the laboratory particle constitutes the cosmic population, because the latter claim requires evidence connecting them. This is what L adds to I, A, and E.

Within the WIMP program, these conditions appeared jointly attainable. A robustly interpreted recoil signal could contribute to constitutive identification, although it would not

by itself establish the candidate's cosmic abundance, distribution, or explanatory reach. For those claims, we need halo, cosmological, and astronomical evidence, as well as reason to identify the locally detected entity with the component responsible for the cosmic phenomena. An account of its production history would provide a fuller cosmological explanation, but is not a necessary condition of cosmic attribution. The WIMP's historical coordinating role rested on a scientifically motivated expectation that the participating programs could supply convergent evidence warranting all four conditions.

Bertone and collaborators provide a concrete example of such a prospective connection [49]. In their simplified models, combined simulated collider and direct-detection results constrain particle parameters from which a thermal relic abundance can be inferred and compared with the cosmological abundance. A mismatch can point to additional dark-sector physics or a different thermal history. Such comparisons supply a model-dependent route toward connecting particle identification and cosmic attribution, with the interpretation depending on the particle model and cosmological assumptions.

Axion and ALP searches make the distinction between identifying a particle and attributing it to the cosmic component particularly clear. A haloscope such as ADMX searches for a local dark-matter population, whereas IAXO is a helioscope designed to detect axions or ALPs produced in the Sun [23, 24]. A signal in either setting could establish the existence and properties of a particle or field and thus contribute to constitutive identification, although the implications for cosmic attribution would differ. A haloscope signal would give evidence that the candidate is present in the Milky Way's halo, whereas a helioscope signal would not, by itself, identify the particle as the cosmic dark component. Neither signal alone would establish the required cosmological abundance, distribution, explanatory adequacy, and cross-domain identification. We could therefore discover an axion or ALP before we discover it *as dark matter*. As with a WIMP, a fuller body of convergent evidence could establish the required attribution, explanation, and cross-domain identification. The distinction concerns what a given result establishes and what evidence would connect it to the cosmic component.

The PBH profile reverses the particle-centered expectation. Evidence for an abundant and appropriately distributed PBH population could identify the kind of object constituting the dark component, establish its cosmic attribution, and account for the relevant gravitational phenomena [26, 27]. We would still need reason to identify the characterized population with the population responsible for those phenomena. If that connection were warranted, the dark component could be identified without discovering a new particle beyond the Standard Model. Constitutive identification must therefore remain candidate-neutral. Defining it in particle terms would build one program's ambitions into the conditions for resolving the shared problem.

Modified-gravity approaches differ more fundamentally. A successful theory could explain phenomena attributed to dark matter while removing the need for a material dark component in the relevant domain [28–30]. Constitutive identification, cosmic attribution, and cross-domain identification would then cease to be outstanding stages of the same discovery: the claim that there is a component to identify would have been revised. Such a result could resolve or dissolve part of the inherited problem without being a discovery of dark matter. The candidates therefore offer different forms of resolution, with different implications for the programs that pursue them.

6 Discussion

The transition away from WIMP-centered research is commonly described in terms of diversification: more candidates, broader mass ranges, and new experiments. Although dark-matter research remains jointly occupied and its tracked 13-family candidate profiles do not become

progressively more divergent, no successor has acquired WIMP-comparable prominence across the fields. WIMP-related research remains central across the literature, while its cross-field connective position has become less singular. The WIMP offered a route connecting a particle kind, a production mechanism, a cosmological abundance, a laboratory strategy, and an explanation of astronomical phenomena. A successor may answer some of these questions without supplying the same connection between them.

A result achieves *local closure* when it settles the question posed within a program: whether a signal identifies a particle or field, whether a population supplies a specified abundance fraction, or whether a theory explains a specified class of phenomena. Identifying the relevant cosmic dark component requires these achievements to be warranted and connected in the way specified by the framework. The profiles in Fig. 5 distinguish what an initial result could establish from what further convergent evidence could justify. WIMP and halo-axion signals face the same requirement to connect a local identification to the cosmic component, even though the programs have different histories and methods.

Epistemic closure concerns whether the evidence warrants the relevant claims of identification, attribution, explanation, and connection, whereas communal ratification concerns whether the participating communities accept that these claims jointly resolve the problem. The two can come apart, particularly when a result fulfills one program's ambitions while leaving another's open. Although the bibliometric record establishes continued participation, it cannot determine how the communities would assess a future discovery claim or whose standards they would regard as authoritative.

Particle physics commonly uses a 5σ excess as the convention for calling an effect a discovery [31]. Historical and methodological analyses show how this threshold emerged from changing experimental standards, and why its application still requires judgment about the particular case [32, 33]. Astronomical and cosmological attribution may require consistency across observations, population inferences, and physical scales. PBH assessments, for example, combine abundance with lensing, compact-object, dynamical, and gravitational-wave constraints [25–27]. Practices vary within fields, so neither account provides a uniform rule. Since cross-domain identification requires the relevant warrants to be connected, one field's discovery convention cannot by itself settle every part of the claim.

The disciplinary payoffs are consequently asymmetric. An axion or ALP discovery could vindicate a particle-search program while leaving dark-matter attribution open. A sufficiently comprehensive PBH result could identify the cosmic component without supplying a new particle beyond the Standard Model. Modified-gravity success could remove the need for a component in the domain it explains. In each case, researchers could agree about the local achievement while differing over whether it fulfills the ambitions pursued through the shared problem.

The earlier history of 'dark matter' illustrates the contingency of this unity. Historical studies show that the phenomena later grouped as one dark-matter anomaly arose from previously distinct concerns [19, 20, 42]. Kelvin and Poincaré discussed unseen or extinguished stellar bodies [34, 35]. Kapteyn and Jeans related additional matter to stellar dynamics [36, 37], and Trumpler used 'dark matter' for light-obscuring material [38]. Lundmark, Pannekoek, and Holmberg connected absorption, dynamical mass deficits, and mass-to-light estimates within the same problem [39–41]. As interstellar absorption became a separate topic, the obscuration strand separated from the dynamical one without a single decisive redefinition [42]. The analogy shows how the evidential connections that allow one term to coordinate several problems can change as those problems develop, without predicting another semantic split. The institutional arrangement can likewise persist as the coordinating role of the candidate that helped establish it weakens. This need not lead to fragmentation, since cross-listing remains substantial and the candidate vocabularies continue to overlap. The issue is whether those continuing contributions support a common resolution. A shared problem can persist

Result or body of evidence	Constitutive identification	Cosmic attribution	Explanatory adequacy	Cross-domain identification
WIMP recoil or halo-axion/ALP signal	Could contribute to identification	Could support local halo presence	Cosmic explanatory reach remains open	Link to cosmic component remains open
Solar axion/ALP signal	Could contribute to identification	Halo and cosmic attribution remain open	Cosmic explanatory reach remains open	Link to cosmic component remains open
WIMP or axion/ALP with convergent evidence	Could identify the particle or field	Could establish abundance and distribution	Could explain attributed phenomena	Evidence must connect the achievements
PBH population with convergent evidence	Could identify the population	Could establish abundance and distribution	Could explain gravitational phenomena	Same population must fulfill the roles
Modified-gravity success	No component identification required in the domain	No component attribution required in the domain	Could explain attributed phenomena	Revises the component claim in the domain

Figure 5: **From an initial result to a dark-matter identification.** Columns use $I(c)$, $A(c)$, $E(c)$, and $L(c)$ as compact labels for constitutive identification, cosmic attribution, explanatory adequacy, and relational cross-domain identification. The text distinguishes the potentially different evidential referents related by $L(c_I, c_A, c_E)$. The first two rows describe initial signals. The next two include further convergent evidence and show what a material-candidate program could establish if the relevant claims were warranted and connected. Both WIMP and axion/ALP programs could meet all four conditions, and a PBH identification could do so without a new particle. Modified-gravity success could instead remove the need for a component in the domain explained. These are assessments of stipulated results and bodies of evidence, rather than rankings of candidate viability.

while the routes to resolving it branch.

The empirical interpretation remains subject to limits of coverage and measurement. ADS and arXiv omit some venues, abstract availability varies over time and across classifications, and dictionary matching misses implicit references while counting critical discussion alongside advocacy. Early divergence estimates rest on smaller annual samples. Bootstrap intervals capture sampling variation within the observed corpus, while rarefaction and permutation address sparse support. Alternative classification rules, coverage tables, lexical exclusions, and document-type sensitivities help assess these sources of uncertainty, without exhausting the historical and interpretive uncertainties involved.

7 Conclusion

The WIMP hypothesis supported a scientifically motivated expectation that laboratory particle detection and identification of the cosmic dark component could converge on one discovery. Its declining prominence unsettles that expectation and brings into view the further evidence needed to connect a local result to a cosmic identification. Research remains jointly occupied, and WIMP-related work remains central across the literature, while cross-field connections become less concentrated around it. The leading successor programs could connect constitutive identification, cosmic attribution, explanatory adequacy, and cross-domain identification in different ways. Where the relevant evidence supports their connection, a successor could warrant all four. Continued participation and shared vocabularies alone do not establish that connection. 'Diverging paths' describes this branching of possible resolutions.

The decline in WIMP prominence therefore weakens the presumption that the next locally decisive result will also close the jointly occupied problem. Claims to have discovered dark matter must state what has been identified, whether it constitutes all or part of the cosmic component, which phenomena it explains, and which cross-domain identifications are warranted or remain open.

8 Materials and methods

8.1 Canonical corpus and reproducibility

The retained ADS snapshot was queried with `full:"dark matter"` and indexed on May 21, 2026. Every analysis excludes records dated after 2025 and deduplicates records by ADS bibcode. The manifest `analysis_manifest.json` documents the query, indexing timestamp, cutoff, exclusions, input and output row counts, SHA-256 hashes, software versions, and random seeds. It also identifies and hashes the canonical inputs and provides the headline sample sizes. A clean rebuild from the retained archive regenerates the canonical tables, with earlier duplicate tables excluded from the active inputs.

8.2 Classifications and primary-order audit

The first listed ADS arXiv class defines a paper's primary assignment. Under any-class membership, a paper is counted in every broad field represented by its classes. Fractional allocation gives weight $1/k$ to each of its k listed classes before aggregation. Astro–HEP cross-listing is the annual share assigned to both fields. The primary-order audit compared ADS order with arXiv primary metadata for 500 multi-classified papers drawn with a fixed seed and stratified by year from 2010–2025. Agreement uncertainty was calculated with a Wilson 95% interval. The supplementary information gives the deterministic class-to-field mapping and classification coverage by year and document type.

8.3 Candidate extraction

Dictionary matching uses normalized abstracts to record mentions at the paper level. The candidate ontology maps explicit names and abbreviations to canonical families, with nearby dark-matter context required where appropriate for expressions with substantial uses outside dark-matter research. Each family is counted at most once per paper. The counts include support, comparison, constraint, and exclusion, and do not imply endorsement. Implicit references and expressions absent from the dictionary are not captured. The complete ontology and extraction code are released with the analysis. Regression tests check canonical matches, ambiguity guards, and family collapse without estimating extraction performance.

8.4 Concept-network analysis

Network analyses use dark-matter papers with primary astrophysics or high-energy-physics assignments, divided into 2006–2015 and 2016–2025. Nodes are normalized candidates, theories, phenomena, problems, methods, apparatuses, and mechanisms identified in abstracts. An undirected edge records paper-level abstract co-occurrence, with cross-field edges requiring support from at least five papers in either field. The pooled candidate–concept network measures role-specific spectral centrality using co-occurrence prevalence among candidate-bearing papers. The cross-field network contains one node per concept and weights an edge by the mean of its field-specific co-occurrence prevalences, giving astrophysics and high-energy physics equal influence. Inverse prevalence supplies weighted distance. Non-candidate concepts at least twice as prevalent in one field define the endpoint sets for normalized cross-field

betweenness. Node-removal effects report the relative loss of mean reciprocal distance between those endpoint sets. The primary WIMP-related node combines generic WIMP, neutralino, LSP, Higgsino, gravitino, and sneutrino mentions. Metrics are calculated upstream, while the saved Gephi projects determine only layout and appearance. Full definitions, WIMP-ontology sensitivities, and the removal analysis are reported in the supplementary information.

8.5 Candidate composition and robustness

For each year and field, candidate counts are normalized over the 13 tracked dictionary families. The supplementary information reports annual Jensen–Shannon divergence, total-variation distance, and cosine distance as descriptive robustness checks. Two thousand fixed-seed, paper-level bootstrap replicates are drawn independently within each field-year. This count refers to resamples, not to the number of papers supporting an estimate. The predefined contrast is the mean annual divergence in 2016–2025 minus that in 1995–2004. Lexical divergence uses yearly document frequencies and is recomputed after removing a predefined set of candidate-name unigrams, and the exclusion list, rankings, thresholds, and selection rules are released in full.

8.6 Additional sensitivities

The supplementary information reports abstract and classification coverage, articles/eprints-only results versus the broader eligible-document definition, raw versus normalized candidate counts, primary/any/fractional classifications, and lexical results with predefined candidate-name unigrams excluded. Candidate-composition checks give each paper total weight one, suppress overlapping generic-WIMP labels, collapse tracked WIMP families, end the series in 2024, rarefy annual field samples, and compare observed divergence with a within-year permutation null. Deterministic tests enforce one primary row per classified paper, fractional and equal-paper weights summing to one, candidate-ontology invariants, reproducible samples and bootstraps, and distance values within mathematical bounds.

Funding

Funding for this research was provided by the National Swedish Research Council (Vetenskapsrådet) grant number 2022-06143.

Author contributions

S.A. conceived the study, curated the data, developed the methodology and software, performed the analysis and visualization, and wrote the manuscript.

Competing interests

The author declares no competing interests.

Data availability

The reproduction repository is https://github.com/SALLZEN/dark-matter-plurality. Version 3.0.0 is archived at https://doi.org/10.5281/zenodo.23263996 [43]; the Zenodo concept DOI https://doi.org/10.5281/zenodo.20241725 links the archive lineage. The retained ADS snapshot, manifest and hashes, primary-order audit, yearly analysis tables, lexical rankings, and figure source data are included or rebuildable as documented in the repository.

Code availability

The self-contained six-stage pipeline and deterministic tests are available in the same repository.

Supplementary material

Supplementary information is included at the end of this preprint. The reproduction bundle provides the accompanying data and code.

References

[1] Gershtein, S. S. & Zeldovich, Y. B. Rest mass of muonic neutrino and cosmology. *Soviet Journal of Experimental and Theoretical Physics Letters* **4**, 120–122 (1966).

[2] Cowsik, R. & McClelland, J. Gravity of neutrinos of nonzero mass in astrophysics. *Astrophysical Journal* **180**, 7–10 (1973). https://doi.org/10.1086/151937.

[3] Rubin, V. C. & Ford Jr, W. K. Rotation of the Andromeda nebula from a spectroscopic survey of emission regions. *Astrophysical Journal* **159**, 379–403 (1970). https://doi.org/10.1086/150317.

[4] Ostriker, J. P., Peebles, P. J. E. & Yahil, A. The size and mass of galaxies and the mass of the universe. *Astrophysical Journal Letters* **193**, L1–L4 (1974). https://doi.org/10.1086/181617.

[5] Einasto, J., Kaasik, A. & Saar, E. Dynamic evidence on massive coronas of galaxies. *Nature* **250**, 309–310 (1974). https://doi.org/10.1038/250309a0.

[6] Tremaine, S. & Gunn, J. E. Dynamical role of light neutral leptons in cosmology. *Physical Review Letters* **42**, 407–410 (1979). https://doi.org/10.1103/PhysRevLett.42.407.

[7] Peebles, P. J. E. Primeval adiabatic perturbations: Effect of massive neutrinos. *Astrophysical Journal* **258**, 415–424 (1982). https://doi.org/10.1086/160094.

[8] Drukier, A. & Stodolsky, L. Principles and applications of a neutral-current detector for neutrino physics and astronomy. *Physical Review D* **30**, 2295–2309 (1984). https://doi.org/10.1103/PhysRevD.30.2295.

[9] Goodman, M. W. & Witten, E. Detectability of certain dark-matter candidates. *Physical Review D* **31**, 3059–3063 (1985). https://doi.org/10.1103/PhysRevD.31.3059.

[10] Pagels, H. & Primack, J. R. Supersymmetry, cosmology, and new physics at teraelectronvolt energies. *Physical Review Letters* **48**, 223–226 (1982). https://doi.org/10.1103/PhysRevLett.48.223.

[11] Goldberg, H. Constraint on the photino mass from cosmology. *Physical Review Letters* **50**, 1419–1422 (1983). https://doi.org/10.1103/PhysRevLett.50.1419.

[12] Ellis, J. R., Hagelin, J. S., Nanopoulos, D. V., Olive, K. & Srednicki, M. Supersymmetric relics from the Big Bang. *Nuclear Physics B* **238**, 453–476 (1984). https://doi.org/10.1016/0550-3213(84)90461-9.

[13] Jungman, G., Kamionkowski, M. & Griest, K. Supersymmetric dark matter. *Physics Reports* **267**, 195–373 (1996). https://doi.org/10.1016/0370-1573(95)00058-5.

[14] Bertone, G., Hooper, D. & Silk, J. Particle dark matter: Evidence, candidates and constraints. *Physics Reports* **405**, 279–390 (2005). https://doi.org/10.1016/j.physrep.2004.08.031.

[15] Akerib, D. S., Alsum, S., Araújo, H. M. *et al.* Results from a search for dark matter in the complete LUX exposure. *Physical Review Letters* **118**, 021303 (2017). https://doi.org/10.1103/PhysRevLett.118.021303.

[16] Aprile, E., Aalbers, J., Agostini, F. *et al.* Dark matter search results from a one ton-year exposure of XENON1T. *Physical Review Letters* **121**, 111302 (2018). https://doi.org/10.1103/PhysRevLett.121.111302.

[17] Cui, X., Abdukerim, A., Chen, W., Chen, X. *et al.* Dark matter results from 54-ton-day exposure of PandaX-II experiment. *Physical Review Letters* **119**, 181302 (2017). https://doi.org/10.1103/PhysRevLett.119.181302.

[18] Bertone, G. & Tait, T. M. P. A new era in the search for dark matter. *Nature* **562**, 51–56 (2018). https://doi.org/10.1038/s41586-018-0542-z.

[19] de Swart, J. G., Bertone, G. & van Dongen, J. How dark matter came to matter. *Nature Astronomy* **1**, 0059 (2017). https://doi.org/10.1038/s41550-017-0059.

[20] Bertone, G. & Hooper, D. History of dark matter. *Reviews of Modern Physics* **90**, 045002 (2018). https://doi.org/10.1103/RevModPhys.90.045002.

[21] Sikivie, P. Axion cosmology. In *Axions: Theory, Cosmology, and Experimental Searches* (eds Kuster, M., Raffelt, G. & Beltrán, B.), Vol. 741 of *Lecture Notes in Physics*, 19–50 (Springer, Berlin, Heidelberg, 2008). astro-ph/0610440.

[22] Adams, C. B., Aggarwal, N., Agrawal, A. *et al.* Axion dark matter. Snowmass 2021 white paper, arXiv:2203.14923 (2022).

[23] Du, N., Force, N., Khatiwada, R. *et al.* Search for invisible axion dark matter with the Axion Dark Matter Experiment. *Physical Review Letters* **120**, 151301 (2018). https://doi.org/10.1103/PhysRevLett.120.151301.

[24] Armengaud, E., Attié, D., Basso, S., Brun, P. *et al.* Physics potential of the International Axion Observatory (IAXO). *Journal of Cosmology and Astroparticle Physics* **2019**, 047 (2019). https://doi.org/10.1088/1475-7516/2019/06/047.

[25] Bird, S. *et al.* Did LIGO detect dark matter? *Physical Review Letters* **116**, 201301 (2016). https://doi.org/10.1103/PhysRevLett.116.201301.

[26] Carr, B., Kühnel, F. & Sandstad, M. Primordial black holes as dark matter. *Physical Review D* **94**, 083504 (2016). https://doi.org/10.1103/PhysRevD.94.083504.

[27] Carr, B. & Kühnel, F. Primordial black holes as dark matter: Recent developments. *Annual Review of Nuclear and Particle Science* **70**, 355–394 (2020). https://doi.org/10.1146/annurev-nucl-050520-125911.

[28] Famaey, B. & McGaugh, S. S. Modified Newtonian dynamics (MOND): Observational phenomenology and relativistic extensions. *Living Reviews in Relativity* **15**, 10 (2012). https://doi.org/10.12942/lrr-2012-10.

[29] Bekenstein, J. D. Relativistic gravitation theory for the modified newtonian dynamics paradigm. *Physical Review D* **70**, 083509 (2004). https://doi.org/10.1103/PhysRevD.70.083509.

[30] Verlinde, E. P. Emergent gravity and the dark universe. *SciPost Physics* **2**, 016 (2017). https://doi.org/10.21468/SciPostPhys.2.3.016.

[31] Cowan, G. Statistics. In *Review of Particle Physics* (Workman, R. L. *et al.*, Particle Data Group), *Progress of Theoretical and Experimental Physics* **2022**, 083C01 (2022). https://doi.org/10.1093/ptep/ptac097.

[32] Franklin, A. *Shifting Standards: Experiments in Particle Physics in the Twentieth Century* (University of Pittsburgh Press, Pittsburgh, 2013).

[33] Dorigo, T. Extraordinary claims: The 0.000029% solution. *EPJ Web of Conferences* **95**, 02003 (2015). https://doi.org/10.1051/epjconf/20159502003.

[34] Thomson, W. *Baltimore lectures on molecular dynamics and the wave theory of light* (C. J. Clay and Sons, London, 1904).

[35] Poincaré, H. The Milky Way and the theory of gases. *Popular Astronomy* **14**, 475–488 (1906).

[36] Kapteyn, J. C. First attempt at a theory of the arrangement and motion of the sidereal system. *Astrophysical Journal* **55**, 302–327 (1922). https://doi.org/10.1086/142670.

[37] Jeans, J. H. The motions of stars in a Kapteyn-Universe. *Monthly Notices of the Royal Astronomical Society* **82**, 122–132 (1922). https://doi.org/10.1093/mnras/82.3.122.

[38] Trumpler, R. J. Dark nebulae. *Leaflet of the Astronomical Society of the Pacific* **1**, 179–182 (1932).

[39] Lundmark, K. Über die Bestimmung der Entfernungen, Dimensionen, Massen und Dichtigkeit für die nächstgelegenen anagalacktischen Sternsysteme. *Meddelanden från Lunds Astronomiska Observatorium, Serie I* **125**, 1–13 (1930).

[40] Pannekoek, A. On the possible existence of large attracting masses in the centre of the galactic system. *Bulletin of the Astronomical Institutes of the Netherlands* **4**, 39–40 (1927).

[41] Holmberg, E. A study of double and multiple galaxies together with inquiries into some general metagalactic problems. *Annals of the Observatory of Lund* **6**, 1–173 (1937).

[42] Allzén, S. Dark matter: Explanatory unification and historical continuity. Preprint, arXiv:2412.13404 (2024).

[43] Allzén, S. Reproduction Bundle for "Diverging paths to dark-matter discovery". Version 3.0.0. *Zenodo* (2026). URL https://doi.org/10.5281/zenodo.23263996.

[44] Kormendy, J. & Knapp, G. R. (eds.). *Dark Matter in the Universe*. International Astronomical Union Symposium **117** (D. Reidel, Dordrecht, 1987). ISBN 978-90-277-2356-7.

[45] Tremaine, S. A historical perspective on dark matter. In *Dark Matter in the Universe* (eds. Kormendy, J. & Knapp, G. R.), International Astronomical Union Symposium **117**, 547–549 (D. Reidel, Dordrecht, 1987). https://doi.org/10.1017/S0074180900150818.

[46] De Baerdemaeker, S. Method-driven experiments and the search for dark matter. *Philosophy of Science* **88**, 124–144 (2021). https://doi.org/10.1086/710055.

[47] De Baerdemaeker, S. A tale of two dark matter concepts: The relation between cosmology and high-energy physics in the context of dark matter research. In *Astronomy and Philosophy: Conceptual and Methodological Foundations and Challenges* (ed. Dick, S. J.), 325–336 (Cambridge University Press, 2026). https://doi.org/10.1017/9781009621168.

[48] Martens, N. C. M. Detecting the dark: Indirectness and dark matter epistemology. In *Astronomy and Philosophy: Conceptual and Methodological Foundations and Challenges* (ed. Dick, S. J.), 279–292 (Cambridge University Press, 2026). https://doi.org/10.1017/9781009621168.

[49] Bertone, G., Bozorgnia, N., Kim, J. S. *et al.* Identifying WIMP dark matter from particle and astroparticle data. *Journal of Cosmology and Astroparticle Physics* **2018**, 026 (2018). https://doi.org/10.1088/1475-7516/2018/03/026.

[50] Allzén, S. Dark matter realism: How referential semantics restricts realism in contemporary fundamental physics. *Synthese* **207**(6), 229 (2026). https://doi.org/10.1007/s11229-026-05589-8.

[51] Allzén, S. Scientific realism and empirical confirmation: A puzzle. *Studies in History and Philosophy of Science Part A* **90**, 153–159 (2021). https://doi.org/10.1016/j.shpsa.2021.10.008.

[52] Martens, N. C. M. Dark matter realism. *Foundations of Physics* **52**(1), 16 (2022). https://doi.org/10.1007/s10701-021-00524-y.

[53] Vaynberg, E. Realism and the detection of dark matter. *Synthese* **204**(3), 82 (2024). https://doi.org/10.1007/s11229-024-04728-3.

Supplementary information

S1 Analysis scope

The analysis uses a frozen NASA Astrophysics Data System (ADS) archive queried with `full:"dark matter"`, indexed May 21, 2026, and restricted to publication years through 2025. The corpus contains 201,627 unique records, including 143,390 with arXiv classifications. A machine-readable manifest documents the query, cutoff, exclusions, row counts, software versions, random seed 20250801, and SHA-256 hashes of the retained inputs. Records dated after 2025 are excluded before every aggregation. Candidate extraction additionally requires abstract text and excludes records whose document type is abstract, thesis, proposal, technical report, press release, editorial, notice, dataset, software, or another excluded type specified in the manifest. Articles, eprints, proceedings papers, book chapters, books, and records without a document-type label remain eligible, and the released document-flow table reports every type and count.

S2 Classification robustness

The main analysis assigns each paper to its first listed ADS arXiv class. Two alternatives test the dependence of the broad-field result on that convention: any-class membership counts a paper in every represented field, whereas fractional assignment distributes weight $1/k$ across its k listed classes. Aggregated astrophysics remains larger than aggregated high-energy physics under all three rules, although the estimated shares vary (Fig. S1). Between 16.8% and 20.2% of classified papers each year from 2010 through 2025 belong to both fields under any-class membership.

A fixed-seed, year-stratified audit compared the first ADS class with the arXiv primary category for 500 multi-classified papers from 2010 through 2025. Agreement was 500/500, or 1.000 (Wilson 95% interval 0.992–1.000). This supports the ordering convention for the sampled records without making the alternative assignment rules redundant. Annual classification coverage ranges from 71.7% to 85.3% during 2010–2025 and does not decline monotonically, with coverage at 85.3% in 2025. Missingness is strongly associated with document type: 78.1% of articles, virtually all eprints, and 30.0% of proceedings papers have an arXiv classification. Records without classifications cannot be assigned a missing "field" without assuming the quantity at issue. The released tables therefore report coverage by year and document type instead. Table S1 gives the deterministic broad-field rules, and `arxiv_class_mapping.csv` enumerates all 149 observed classes and their corpus counts.

S3 Candidate measurement

Candidate extraction uses a frozen dictionary applied to normalized abstracts and counts paper-level mentions. The ontology maps explicit names and abbreviations to canonical candidate families, with nearby dark-matter context required where appropriate for expressions with substantial non-dark-matter uses. The unit is the paper–candidate-family pair: repeated occurrences do not increase the count, whereas explicitly named generic and specific families may each contribute once when they co-occur. Synonymous tags mapped to the same family still contribute only once. Mentions can occur in support, comparison, criticism, constraint, or exclusion and therefore must not be interpreted as endorsement. Implicit references and formulations absent from the dictionary are outside the measurement.

Under the primary field assignment, the candidate-composition analysis tracks the union of the nine most-mentioned families in each focal field across 1995–2025, yielding the 13 families

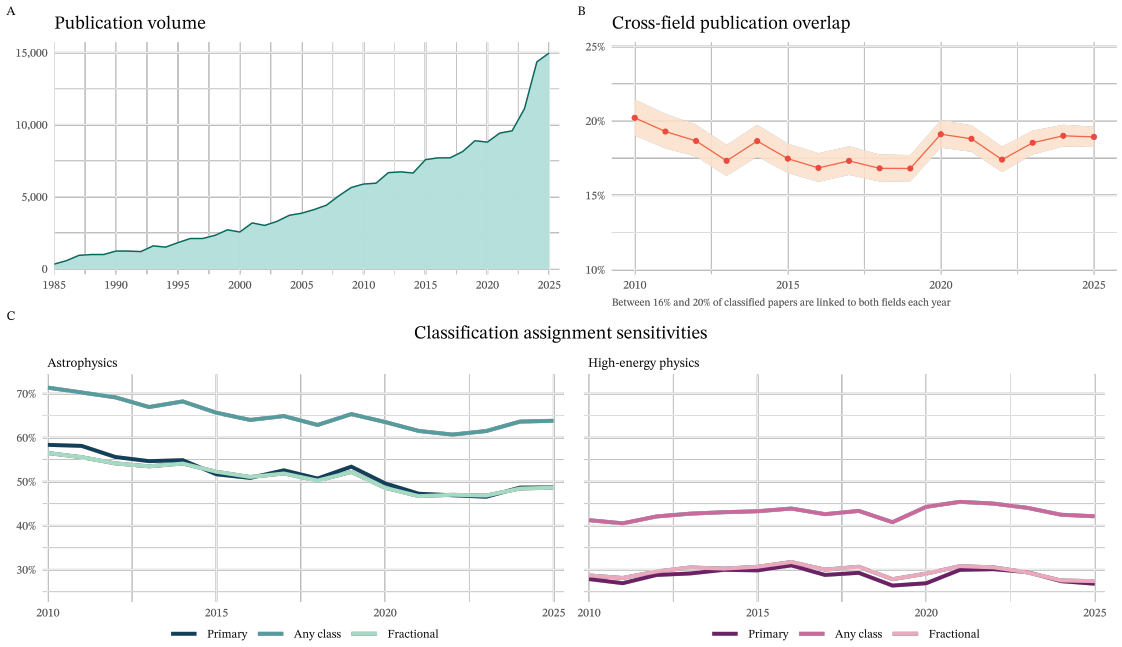


Figure S1: **Corpus scale and classification sensitivity.** (**A**) Annual publication volume in the frozen ADS corpus. (**B**) Annual share of papers linked to both astrophysics and high-energy physics under any-class membership, with a Wilson 95% interval. (**C**) Broad-field shares under primary, any-class, and fractional assignments. Aggregated astrophysics exceeds aggregated high-energy physics under each rule.

Tbl S1: Broad-field mapping of arXiv classes. Prefix matching is applied in the order implemented in the released canonical-data script. The two focal fields are astrophysics and high-energy physics, while all remaining categories are retained for the corpus-scale comparison.

Class prefix	Broad field
`astro*`	Astrophysics
`hep-*`	High-energy physics
`cond-mat*`	Condensed matter
`math*`	Mathematics
`cs*`	Computer science
`quant-ph*`	Quantum physics
`gr-qc*`	General relativity and quantum cosmology
`nucl-*`	Nuclear physics
`physics*`	Physics
`q-bio*`	Quantitative biology
`q-fin*`	Quantitative finance
`stat*`	Statistics
`econ*`	Economics
`eess*`	Electrical engineering/systems
`nlin*`	Nonlinear sciences

in Table S2. The complete regular expressions, the default 60-character context window on each side of context-dependent expressions, a documented 160-character exception for glueball expressions, the wider ontology, and the canonical mappings are released with the analysis.

A manually annotated 180-record candidate-enriched development set was used to refine boundary cases. Because it informed the final rules, it is not used to estimate precision or recall.

Tbl S2: Tracked candidate families and their canonical ontology tags. Multiple tags assigned to one family collapse to one paper–candidate-family mention.

Family	Canonical tag(s)
Axion + ALP	`axion`, `alp`
Generic WIMP	`wimp`
Gravitino	`gravitino_dm`
Bino	`bino_dm`
Higgsino	`higgsino_dm`
LSP	`lsp`
MACHO	`macho`
Mirror DM	`mirror_dm`
Neutralino	`neutralino`
PBH	`pbh`
SFDM	`scalar_field_dm`
Sterile ν	`sterile_neutrino`
ULA / Fuzzy	`ula`, `fuzzy_dm`

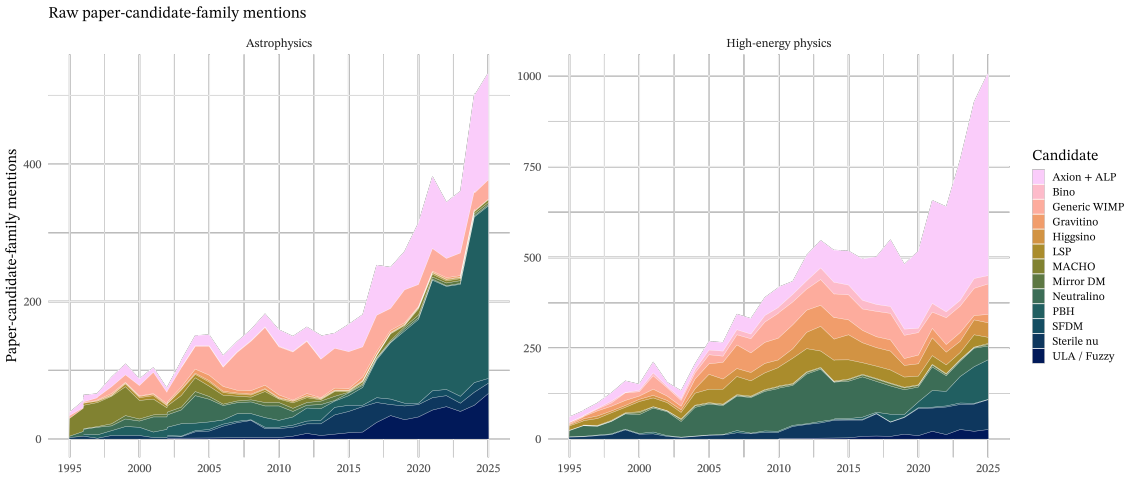


Figure S2: **Raw annual paper–candidate-family mentions.** Counts are shown separately for astrophysics and high-energy physics, with separate vertical scales for the two panels.

Deterministic tests exercise canonical mentions, negated and ambiguous expressions, spelling variants, experiment-name exclusions, and family collapse. These tests check implementation integrity rather than estimate population-level performance.

S4 Candidate-composition robustness

The main text normalizes candidate mentions within each field-year so that publication growth does not determine the comparison. Main-text Fig. 2 reports, for each candidate-year, the base-2 logarithm of the high-energy-physics share divided by the astrophysics share. A pseudocount of 0.5 per family is added before normalizing each field's counts, preventing undefined ratios when one field has no mention. Candidate-years with fewer than 20 paper–candidate-family mentions across the two fields are omitted from the display.

Raw annual paper–candidate-family mention counts are shown in Fig. S2. A paper mentioning more than one family contributes once to each mentioned family.

Tbl S3: Candidate-composition specification checks. Except for the 2024-cutoff row, the early and late periods are 1995–2004 and 2016–2025. The 2024-cutoff row uses 2015–2024 as its late ten-year window.

Jensen-Shannon Divergence			
Specification	Early mean	Late mean	Late - early
Candidate-family scope			
Top 5 per field (7 unique families)	0.149	0.236	+0.087
Top 6 per field (9 unique families)	0.311	0.252	−0.059
Top 9 per field (13 families; primary)	0.326	0.293	−0.033
Top 12 per field (18 unique families)	0.343	0.306	−0.037
All observed ontology families (36)	0.351	0.322	−0.029
Other specification checks			
Equal total weight per paper	0.345	0.281	−0.064
Overlapping generic WIMP suppressed	0.338	0.295	−0.042
Tracked WIMP families collapsed	0.271	0.212	−0.059
Corpus ending in 2024	0.326	0.305	−0.021

For each year, the field-normalized candidate-mention profiles form probability vectors P and Q. Jensen–Shannon divergence (JSD) is

$$\mathrm{JSD}(P,Q) = \tfrac{1}{2}D_{\mathrm{KL}}(P\|M) + \tfrac{1}{2}D_{\mathrm{KL}}(Q\|M), \qquad M = \tfrac{1}{2}(P+Q)$$

using base-2 logarithms.

Total-variation and cosine distances provide alternative summaries. JSD intervals use 2,000 fixed-seed, paper-level bootstrap replicates within each field-year. Annual support is much smaller early in the series: 37 astrophysics and 55 high-energy-physics papers mention a tracked candidate in 1995, compared with 487 and 878, respectively, in 2025.

The predefined comparison is mean annual JSD in 2016–2025 minus mean annual JSD in 1995–2004. Under the primary 13-family specification, mean JSD declines from 0.326 to 0.293, giving a contrast of −0.033 (95% bootstrap interval −0.095 to −0.018). None of the three distance summaries supports increasing distributional divergence (Fig. S3). Although these measures describe candidate-mention profiles under the stated dictionary and aggregation rules, they do not measure endorsement, belief, or the functional relationship among disciplinary results.

Table S3 reports specification checks. Candidate-family scope matters most under aggressive truncation. The union of the five most-mentioned families per field contains seven unique families and yields an increase from 0.149 to 0.236. The contrast becomes negative at a union of the top six per field and remains negative through every tested cutoff up to the top 20. Retaining all 36 observed ontology families gives 0.351 early and 0.322 late. The complete top-k sequence is released with the analysis.

Giving every paper total weight one across all families it mentions strengthens rather than reverses the primary negative contrast. Of 2,560 papers with a generic-WIMP mention, 285 (11.1%) also mention a tracked specific WIMP family. Suppressing the overlapping generic label or collapsing the tracked WIMP-related families likewise preserves the result. Ending the corpus in 2024 leaves the direction unchanged and leaves the first *hep-ph*/*astro-ph.CO* crossover in 2015.

Rarefaction addresses sparse-sample bias directly. Each field-year was repeatedly down-sampled without replacement to the minimum annual support (37 astrophysics and 55 high-energy-physics papers, 1,000 replicates). Mean rarefied JSD is 0.357 in 1995–2004 and 0.337 in 2016–2025, a contrast of −0.020 (simulation interval −0.075 to 0.029). Equalizing support

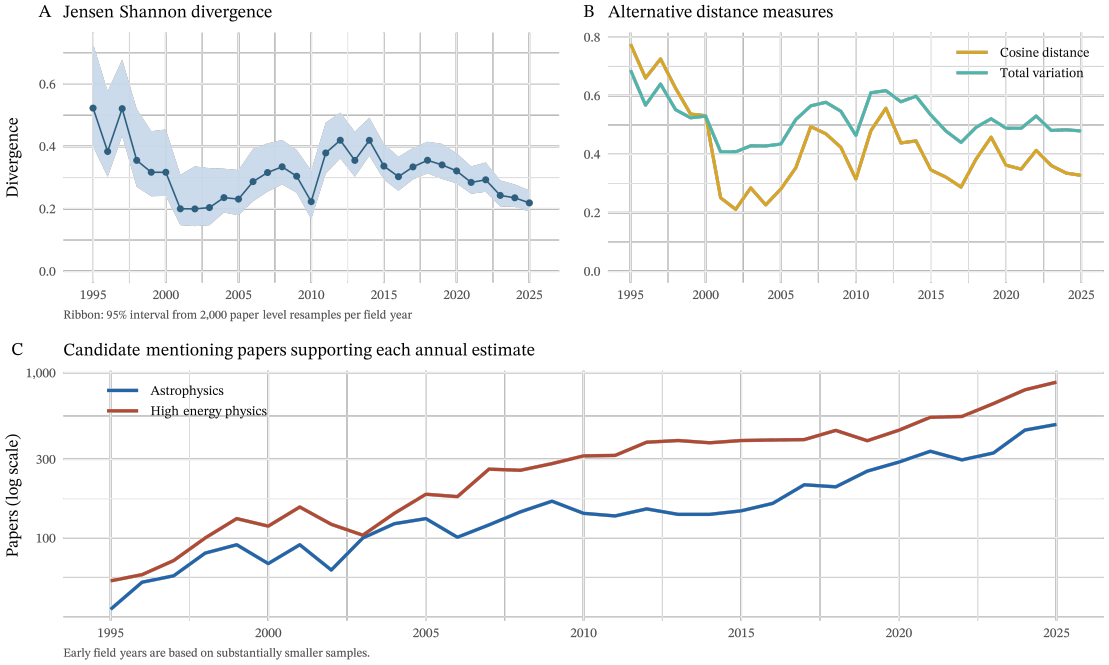


Figure S3: **Candidate-composition divergence and supporting sample sizes.** (**A**) Annual JSD with 95% intervals from 2,000 paper-level bootstrap replicates per field-year, with shaded windows marking 1995–2004 and 2016–2025. (**B**) Total-variation and cosine-distance summaries. (**C**) Candidate-mentioning papers supporting each annual field estimate on a logarithmic scale.

therefore removes evidence for a precise decline but supplies no evidence of an increase.

In a separate 1,000-replicate null, field labels were permuted within each year while field sample sizes were preserved. Observed JSD exceeds the 97.5th percentile of the null in all 31 years, showing persistent compositional differences beyond sparse sampling, and observed-minus-null divergence is 0.286 early and 0.282 late and does not increase.

Restricting the corpus to articles and eprints produces the same directional comparison: mean JSD is 0.354 in 1995–2004 and 0.294 in 2016–2025. Abstract coverage over 2010–2025 has a median of 1.000 in both fields and minima of 0.971 in astrophysics and 0.964 in high-energy physics. Whereas candidate and lexical analyses remain conditional on abstract availability, volume and classification results do not.

S5 Lexical robustness

Lexical data are annual paper-level document frequencies after normalization and masking of mathematical spans. A paper contributes at most once to a term-year count. The field comparison is recomputed after removing 16 predefined candidate-name unigrams: *adm*, *alp*, *axion*, *axino*, *bino*, *gravitino*, *higgsino*, *macho*, *neutralino*, *pbh*, *sneutrino*, *sterile*, *superwimp*, *ula*, *wimp*, and *wino*. This check prevents those explicit names from mechanically generating the lexical difference, but it does not remove every word that could occur in a candidate description.

Illustrative trajectories are selected algorithmically rather than by hand. After applying the released nuisance-term exclusion list, a degree-one loess curve with span 0.75 estimates each field–term trajectory at 1995 and 2025. Eligible terms have at least 40 paper–term occurrences, a smoothed endpoint frequency of at least 0.5%, and an absolute endpoint $\log_2$-ratio of at least 0.3. The five largest increases and five largest decreases are shown for each field, with

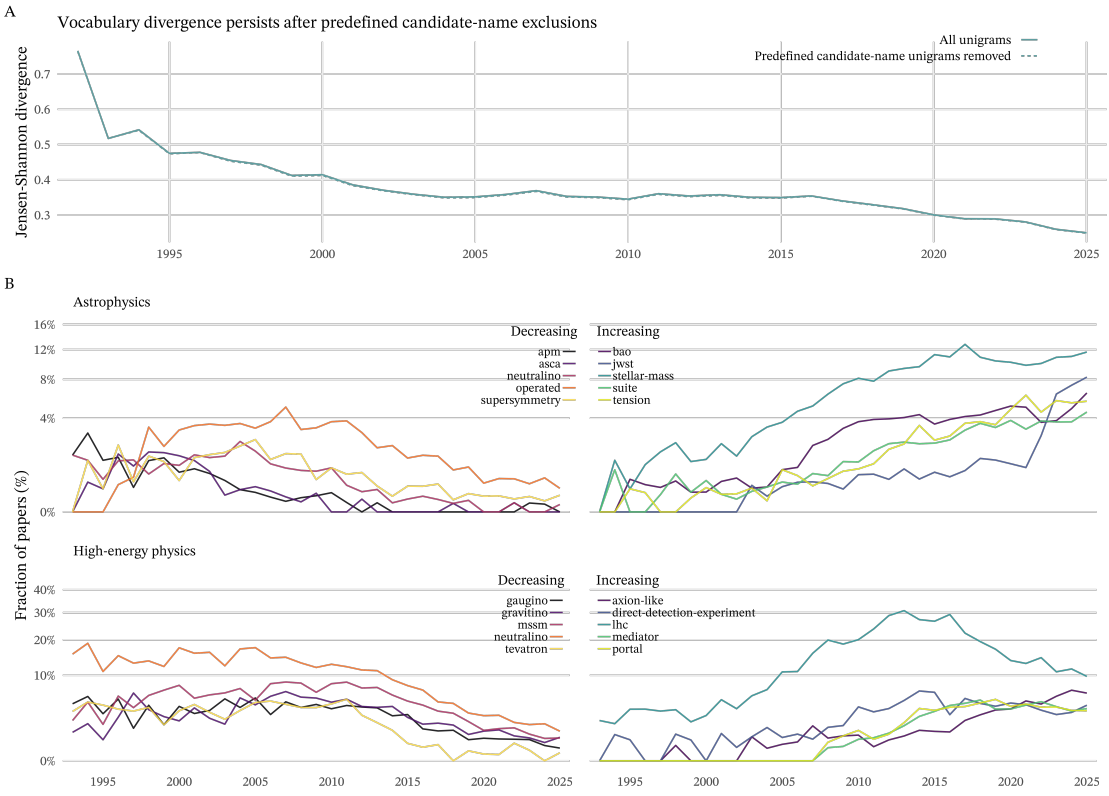


Figure S4: **Lexical divergence and selected term trajectories.** (**A**) Annual vocabulary JSD using all eligible unigrams and after removal of the 16 predefined candidate-name unigrams. (**B**) The five strongest data-selected declining and five strongest rising term trajectories in each field after the documented exclusions and minimum-frequency thresholds.

field-years below 40 papers omitted. The nonzero divergence persists under the predefined candidate-name exclusion, while the broad term trajectories remain dependent on preprocessing and vocabulary thresholds (Fig. S4).

S6 Concept-network analysis

The network analysis asks two questions that candidate shares cannot answer: whether WIMP-related research remains central to the area-wide discourse, and whether cross-field connection remains concentrated around it. It uses records with primary astrophysics or high-energy-physics assignments in two non-overlapping windows, 2006–2015 and 2016–2025. Abstract entities are normalized to one node per candidate, theory, phenomenon, problem, method, apparatus, or mechanism. Co-occurrence is paper-binary in an undirected network: an edge records that two nodes appear in the same abstract, regardless of repeated mentions. Candidate definitions are applied before aggregation so that overlapping generic and named WIMP mentions cannot create self-edges or duplicate paper-level support.

The area-wide importance network is bipartite, connecting candidate nodes to non-candidate concept nodes when they co-occur in an abstract. Its primary edge weight is the number of supporting papers divided by the number of candidate-bearing papers in the pooled focal-field corpus. Role-specific spectral centrality is calculated from the leading singular vectors of the weighted candidate–concept matrix, equivalent here to the candidate and concept roles measured by the Hyperlink-Induced Topic Search (HITS) algorithm. Candidate percentiles are calculated within each period and WIMP definition. Under the published umbrella definition,

the WIMP-related node ranks first among candidates in both periods even though its paper count and total prevalence strength decline.

The cross-field network contains one node per normalized concept rather than separate copies for the two fields. For an edge joining concepts i and j, the primary weight is

$$w_{ij} = \frac{1}{2}\left(\frac{n_{ij,A}}{N_A} + \frac{n_{ij,H}}{N_H}\right), \tag{2}$$

where $n_{ij,A}$ and $n_{ij,H}$ are the numbers of astrophysics and high-energy-physics papers containing both concepts, and N_A and N_H are the corresponding numbers of eligible papers. An edge is retained when it has support from at least five papers in either field. This arithmetic mean gives the fields equal influence rather than allowing the larger field corpus to determine the layout and shortest paths. Weighted distance is $d_{ij} = 1/w_{ij}$, so more prevalent conceptual connections are shorter.

For node i, field orientation uses the smoothed prevalences $\widetilde{p}_{i,A} = (n_{i,A} + 0.5)/(N_A + 1)$ and $\widetilde{p}_{i,H} = (n_{i,H} + 0.5)/(N_H + 1)$, where A and H denote astrophysics and high-energy physics. The 0.5 correction prevents undefined ratios. Orientation is $\log_{10}(\widetilde{p}_{i,A}/\widetilde{p}_{i,H})$. For display color only, it is clipped to $[-1, 1]$, corresponding to at most a tenfold orientation toward either field. Non-candidate nodes at least twice as prevalent in one field as in the other, evaluated from the unclipped smoothed ratio, define the two endpoint sets.

Cross-field betweenness is NetworkX's normalized weighted subset betweenness between these sets. It should therefore be read as a normalized positional centrality, not as a literal percentage of all astrophysics–HEP shortest paths. Under the primary published-umbrella definition, the WIMP-related node moves from rank 1 of 23 candidate nodes in 2006–2015 to rank 2 of 28 in 2016–2025, while its normalized value falls from 0.073 to 0.031. The precise late-period rank is definition-sensitive (Table S4).

The node-removal analysis supplies a distinct sensitivity test (Fig. S5). Cross-field efficiency is the mean reciprocal weighted distance between the field-oriented endpoint sets. For a tested node v, the removal effect is

$$R(v) = \frac{E(G) - E(G - v)}{E(G)}, \tag{3}$$

where $E(G)$ is baseline cross-field efficiency. Removal is evaluated for all candidate nodes and the 25 highest-betweenness non-anchor nodes, while untested nodes remain visible as network context at the minimum size. Under the primary WIMP umbrella, $R(v)$ falls from 41.5% to 6.2%. In 2016–2025, Axion + ALP has a 20.1% removal effect and PBH a 10.1% effect. As properties of the measured discourse network, these values are not estimates of evidential dependence and do not support a historical counterfactual in which the corresponding research was absent.

All three definitions preserve the main directional result: WIMP-related research remains highly central area-wide, while its normalized cross-field betweenness and removal effect are smaller in the later period. Because the precise late-period candidate rank is definition-sensitive, the main argument emphasizes the change from concentrated to more distributed connectivity rather than the displacement of one uniquely specified WIMP node.

The field-oriented endpoint sets are estimated separately in each period and therefore also change with the discourse. The earlier network contains 25 astrophysics-oriented and 124 high-energy-physics-oriented endpoints, compared with 55 and 181, respectively, in the later network. Fifteen astrophysics endpoints and 99 high-energy-physics endpoints occur in both sets. Recalculating both networks using only these common endpoints preserves the result: WIMP-related betweenness falls from 0.035 to 0.006, its candidate rank changes from 1/23 to 2/28, and its removal effect falls from 36.1% to 4.1%. The decline is therefore not produced by

Tbl S4: Sensitivity to the WIMP node definition. Area-wide HITS values are candidate-relative percentiles. Cross-field rank and normalized betweenness use the field-oriented network, and removal effect is the relative loss of cross-field efficiency. The published umbrella, used in the main text, combines generic WIMP, neutralino, LSP, Higgsino, gravitino, and sneutrino mentions. The canonical core omits gravitino.

Period	HITS percentile	Candidate rank	Normalized betweenness	Removal effect
Generic WIMP only				
2006–2015	0.862	1/28	0.039	21.6%
2016–2025	0.909	5/33	0.010	3.4%
Canonical WIMP core				
2006–2015	1.000	1/24	0.069	39.1%
2016–2025	1.000	2/29	0.026	5.7%
Published WIMP umbrella				
2006–2015	1.000	1/23	0.073	41.5%
2016–2025	1.000	2/28	0.031	6.2%

Tbl S5: Rolling-window sensitivity of the published WIMP umbrella. Five-year windows advance by one year. Rolling networks retain edges supported by at least three papers in either field. HITS is the candidate-relative area-wide spectral percentile, whereas betweenness and removal effect are computed in the cross-field network.

Window	HITS percentile	Normalized betweenness	Removal effect
2006–2010	1.000	0.093	39.6%
2007–2011	1.000	0.088	42.5%
2008–2012	1.000	0.089	40.7%
2009–2013	1.000	0.076	39.6%
2010–2014	1.000	0.077	35.1%
2011–2015	1.000	0.076	28.5%
2012–2016	1.000	0.066	26.3%
2013–2017	1.000	0.067	18.7%
2014–2018	1.000	0.057	16.4%
2015–2019	1.000	0.055	13.0%
2016–2020	1.000	0.046	8.5%
2017–2021	1.000	0.041	6.7%
2018–2022	1.000	0.038	5.4%
2019–2023	1.000	0.041	6.4%
2020–2024	0.964	0.023	3.9%
2021–2025	0.964	0.021	3.5%

period-specific endpoint turnover. Endpoint membership, overlap, efficiency, reachability, and unrounded values are reported in the released `common_endpoint_sensitivity.csv`.

Across the 16 rolling windows, the WIMP umbrella's normalized betweenness declines from 0.093 to 0.021 and its removal effect from 39.6% to 3.5%, notwithstanding small local reversals. Its area-wide HITS percentile remains 1.000 through 2019–2023 and is 0.964 in the final two windows. The decade-window contrast therefore does not depend on placing the boundary specifically between 2015 and 2016.

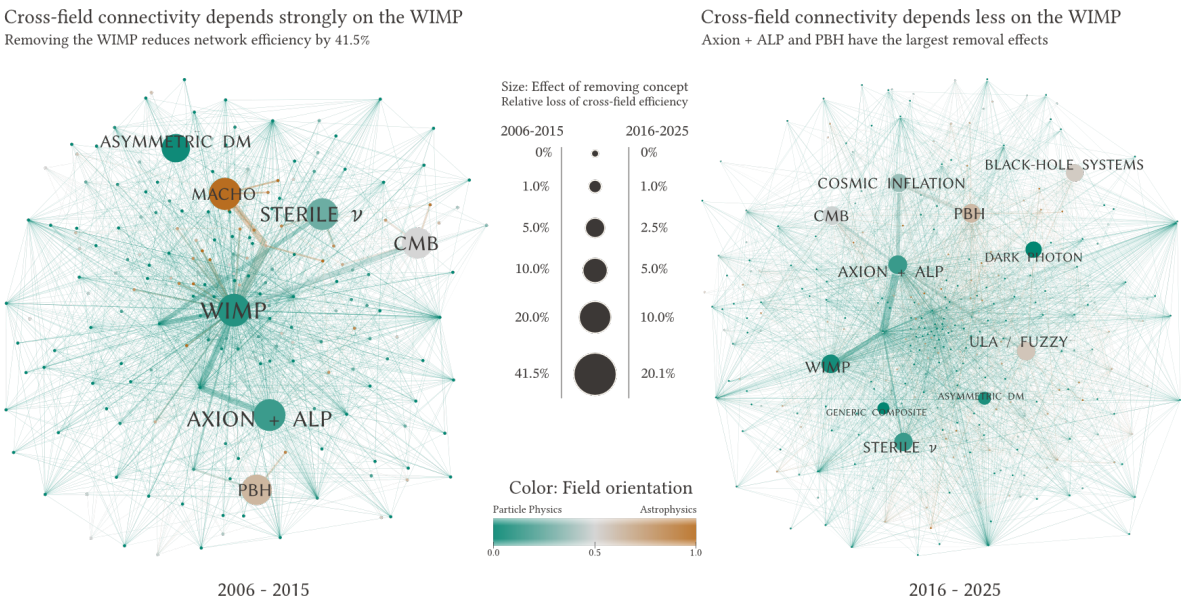


Figure S5: **The cross-field discourse network becomes less dependent on the WIMP-related node.** Node size encodes the relative loss of cross-field efficiency after removal, while color and edge weight follow the cross-field network in main-text Fig. 4. The cyan outline identifies the published WIMP umbrella. Removing this node reduces cross-field efficiency by 41.5% in 2006–2015 and 6.2% in 2016–2025. Axion + ALP and PBH have larger effects in the later period. Nodes without a calculated removal effect are shown at the minimum size.

S7 Reproducibility

The reproduction bundle contains the retained ADS snapshot, analysis manifest, primary-order audit, yearly source tables, figure data, and the six-stage pipeline. The concept-network extension documents its node and edge definitions, filters, centrality measures, WIMP mappings, and canonical metric tables, while the Gephi project files retain the layouts and display mappings used in Figs. 4 and S5. Deterministic tests enforce one primary classification per paper, fractional and equal-paper weights summing to one, reproducible samples and bootstraps, distance values within their mathematical bounds, and candidate-ontology invariants. The repository is available at https://github.com/SALLZEN/dark-matter-plurality, and its archived releases are linked through the repository's Zenodo record.