PNAS Nexus, 2026, , 1–9
Research Report

Diverging paths to dark-matter discovery

Simon Allzén[1,2,3,*]

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

Abstract

Laboratory detection of a new particle and astronomical identification of the gravitating component called dark matter are often treated as routes toward one discovery. The WIMP made that expectation scientifically credible by connecting a microphysical candidate, a cosmological production history and abundance, laboratory search strategies, and explanations of astronomical phenomena. I analyse 201,627 NASA Astrophysics Data System records through 2025, including 143,390 with arXiv classifications, to examine what followed the decline of WIMP prominence. Astrophysics and high-energy physics remain jointly represented, with substantial cross-listing between them. WIMP-related candidate mentions lose relative share in both fields, while successor families receive unequal emphasis. Yet candidate-mention profiles across all 13 tracked families do not become progressively more divergent. This result is robust to small-sample, weighting, candidate-family hierarchy, and cutoff sensitivities, while the larger broad-field astrophysics share persists under alternative field assignments. These bibliometric indicators measure the organisation of attention, not scientists' discovery standards. I therefore interpret them through a conceptual framework distinguishing constitutive identification, cosmic attribution, explanatory adequacy, and the cross-domain identification that makes separately warranted results parts of one dark-matter claim. A particle, compact-object population, or gravitational theory can resolve different parts of the inherited problem, and a partial or multicomponent discovery need not close it as a whole. "Diverging paths" accordingly describes branching routes to resolution after the WIMP's decline, not increasing statistical separation between the fields' candidate-mention profiles.

Keywords dark matter, WIMP, astroparticle physics, bibliometrics, scientific discovery

Subject areas Astronomy (primary), Physics (secondary)

A panacea

Particle physics can identify a new particle without establishing what role it plays on astronomical scales. Astronomy can infer the distribution and behaviour of gravitating matter without identifying its microscopic constitution. These are independently warrantable achievements. Yet for much of late twentieth-century dark-matter research they could be treated as routes toward the same discovery. The weakly interacting massive particle (WIMP) made that expectation scientifically credible: cosmology required cold, non-baryonic matter, while weak-scale extensions of the Standard Model supplied stable relics whose freeze-out abundance could be of the order required for the astronomically inferred dark component [10–14].

The WIMP connected questions that neither astronomy nor particle physics could answer alone. 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. Particle physics supplied stable weak-scale candidates and methods for detecting their non-gravitational interactions [8, 9]. A successful WIMP program could therefore do more than add a particle to the inventory of nature: it could establish that an experimentally identified particle was the same component whose effects had been inferred at galactic and cosmological scales. Contemporary reviews made this prospective convergence explicit by treating relic abundance, halo properties, and direct and indirect detection as connected parts of the WIMP program [13, 14]. The WIMP's earlier cross-disciplinary prominence has since declined. Direct searches have progressively restricted canonical WIMP parameter space [15–17]. As the bibliometric analysis below demonstrates, WIMP-related candidate mentions lose share while axion and ALP work and PBH work expand (Fig. 3). Contemporary reviews place this shift within a broader diversification that also includes dark sectors and modified-gravity approaches [18, 27]. The consequential question is not only which candidate, if any, will succeed the WIMP. It is whether any successor can connect the same evidential achievements.

This question is inseparable from the organisation of dark-matter research. The particle turn gave particle physics substantial standing within a problem whose empirical basis remained astronomical and cosmological. That publishing and classification structure persists: particle phenomenology is a principal location of dark-matter research, aggregated astrophysics remains larger, and cross-listing between the focal fields remains substantial. This article makes two distinct contributions. First, I establish four bibliometric regularities using an ADS snapshot covering records through 2025: persistent cross-disciplinary participation, declining WIMP prominence, unequal successor mixtures, and no increase in the divergence of the 13-family candidate-mention profiles. These indicators measure organised attention; they do not directly measure epistemic complementarity, discovery standards, or communal judgements. Second, I use a conceptual reconstruction of candidate-specific evidence to argue that continued participation and vocabulary overlap need not make locally decisive results parts of one discovery. Here and throughout, *centrality* means relative prominence in the publication and candidate-mention record, not authority, prestige, funding, or network position.

The argument therefore does not infer a change in epistemic organisation from bibliometric divergence. It asks how the measured reorganisation of attention should be interpreted once the evidential achievements associated with different candidates are distinguished. The following section reconstructs the WIMP as an evidential bridge and the basis of a reciprocal division of epistemic labour. The empirical sections then establish the four bounded regularities just stated. I next distinguish constitutive identification, cosmic attribution, explanatory adequacy, and the cross-domain identification that connects them. This framework explains how a particle, a compact-object population, or a gravitational theory can resolve different parts of the inherited problem without yielding the same kind of discovery. The discussion returns to the difference between local closure, epistemic closure of the joint problem, and communal ratification.

The WIMP as an evidential bridge

Dark matter first presents a role-based problem. Historical studies show that several missing-mass problems were consolidated before they were coupled to particle candidates [19, 20]. A range of observations supports the existence of something that supplies additional gravitating mass, possesses an appropriate abundance and distribution, and contributes to the formation of cosmic structure. This claim can be represented schematically as

$$\exists x \, R(x) \tag{1}$$

where $R(x)$ denotes the cluster of dynamical, cosmological, and explanatory roles assigned to dark matter. Evidence for the existence of a realiser of R does not by itself identify what that realiser is. The inference can become increasingly robust while its constitutive interpretation remains open. The particle turn did not create the dark-matter problem; it reconfigured an existing, historically heterogeneous set of astronomical concerns as a question of microphysical constitution [42]. Cosmological mass bounds first made massive neutrinos candidates for the missing component [1, 2, 6, 7]. Axion and supersymmetric proposals then connected the same astronomical role to particle properties, production mechanisms, and non-gravitational searches [10–13, 21]. The question was no longer only whether unseen matter was required, but which microscopic entity realised the inferred role. Particle physics thereby introduced two further ambitions: *constitutive identification* and non-gravitational detection. A dark-matter discovery could now be expected to reveal not only the existence of additional gravitating matter, but also its nature.

This reformulation created a specific upshot to astronomy and cosmology. A particle candidate could turn an inferential placeholder into a physically specified entity, explain its stability and interactions, and provide evidence independent of the gravitational phenomena that originally motivated the problem. Direct detection made that promise even more concrete by making it possible, in principle, to investigate the same entity in a controlled terrestrial setting [8, 9]. The evidential chain would then run from astronomical effects to a cosmic component and from that component to a laboratory-identifiable constituent.

The arrival of particle physics was not only beneficial to astronomy and cosmology. Dark matter gave theories of physics beyond the Standard Model an empirically established cosmic role to fill. Cosmological abundance and structure formation constrained candidate properties, while the local halo supplied a target for experiments. Supersymmetric relics were therefore not merely speculative additions to particle physics: they could be presented as prospective constituents of matter already required on independent grounds [10–13]. Dark matter made new microphysics empirically actionable, just as new microphysics promised to determine the nature of dark matter. The WIMP coordinated these reciprocal promises through an implicit division of epistemic labour. Astronomy established the relevant phenomena and constrained the spatial distribution of the gravitating component. Cosmology constrained its abundance, production history, and role in structure formation. Particle physics proposed its microscopic constitution and developed laboratory detection strategies. The WIMP supplied a reason to regard these different activities as investigations of the same entity. Their contributions could thus be considered epistemically complementary without using the same methods, addressing the same scales, or producing the same kinds of evidence.

No single element of this arrangement was sufficient for discovery. A relic-density calculation did not establish that WIMPs existed; a laboratory excess would not by itself show that the detected particle supplied the cosmic abundance; and an astronomical fit could not determine a unique microscopic constitution. The force of the WIMP program lay instead in the prospect that independently produced results would converge upon a common candidate. Its historical importance is not exhausted by its prominence in a catalogue of possibilities. I reconstruct the WIMP as an evidential bridge: it specified how otherwise decoupled scientific advancements might become parts of a single coherent discovery. This is an analytical account of the program's evidential structure, not a claim that every historical actor articulated the structure in these terms. The convergence at issue is functional rather than distributional. A productive division of labour gives participating fields different immediate objects, vocabularies, and evidential responsibilities; it therefore does not predict identical candidate distributions. What matters is whether their outputs can be connected through a warranted identification of the entity studied in each context. Astrophysics and high-energy physics need not become increasingly dissimilar for their candidate programs to offer different routes to discovery.

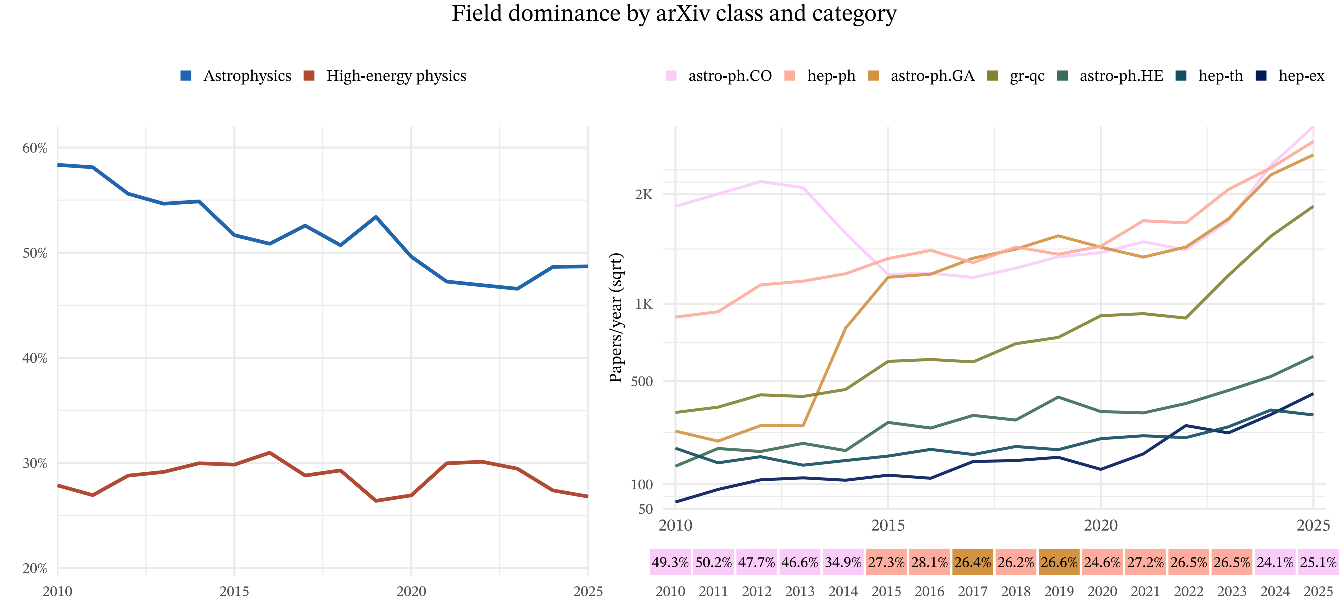

Figure 1 The disciplinary location of dark-matter research. (A) Annual shares of classified papers under primary assignments, aggregated into broad fields; astrophysics remains larger than high-energy physics through 2025. (B) Annual paper counts on a square-root vertical scale for the seven primary arXiv classes with the largest cumulative counts in 2010–2025; *hep-ph* first exceeds the single class *astro-ph.CO* in 2015, while the class ordering varies thereafter. The strip beneath panel B identifies the leading displayed class each year and its share of papers across those seven classes. Publication volume, alternative assignment rules, and cross-listing are reported in Fig. S1.

A jointly occupied problem

The WIMP's coordinating role matters only because dark matter became jointly occupied in the publication and classification structure. I locate that limited form of organisation through arXiv classifications attached to the ADS corpus. Of 201,627 unique records published no later than 2025, 143,390 have between one and nine arXiv classes. These metadata identify where research is published and how it is classified; they do not transparently encode disciplinary identity, confer epistemic authority, or show that every author classified in a field shares the same aims.

The scale and class results must be read together. Under primary assignments aggregated into broad fields, astrophysics remains larger throughout 2010–2025 and accounts for 48.7% of classified papers in 2025, compared with 26.8% for high-energy physics (Fig. 1A). At the level of individual primary classes, however, *hep-ph* first exceeds *astro-ph.CO* in 2015 (Fig. 1B). This is a narrow comparison between two classes, not evidence that high-energy physics displaced astrophysics as a whole. The ordering among leading classes also changes after 2015. The result is therefore substantial particle-phenomenology prominence within a literature that remains larger on the astrophysical side when classes are aggregated.

That broad-field direction is unchanged under any-class and fractional assignments, although the estimated shares depend on how multi-classified papers are counted (Fig. S1C). Between 16.8% and 20.2% of classified papers each year from 2010 through 2025 belong to both broad fields under any-class membership (Fig. S1B). 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). This supports the primary-order convention for the sampled records without making the alternative assignments redundant.

Together, these results show institutional differentiation without separation. Particle phenomenology is central to dark-matter research, aggregated astrophysics remains larger, and cross-listing between the focal fields remains substantial. In this limited sense, dark matter is a jointly occupied problem: several fields retain standing to formulate it and contribute to its resolution. Publication classifications alone do not establish how discovery claims would be evaluated or whether any field possesses a veto over them.

The loss of a shared path to closure

Joint occupation does not imply that the participating fields retain one route to resolution. Candidate families receive persistently unequal attention within the two broad literatures. Figure 2 compares 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, whereas PBHs, MACHOs, and ultralight or fuzzy candidates are relatively more prominent in astrophysics. These are relative alignments, not disciplinary ownership. A family can be studied in both fields while occupying a different place in each field's candidate mixture.

Mentions track organised attention, not acceptance. A paper can introduce, compare, constrain, or exclude a candidate it names. With that limitation, the within-field composition supplies the clearest temporal result (Fig. 3). WIMP-related candidates lose share in both literatures. Axion and ALP work expands especially in high-energy physics, supported by haloscope and helioscope programs [21–24]; PBH work

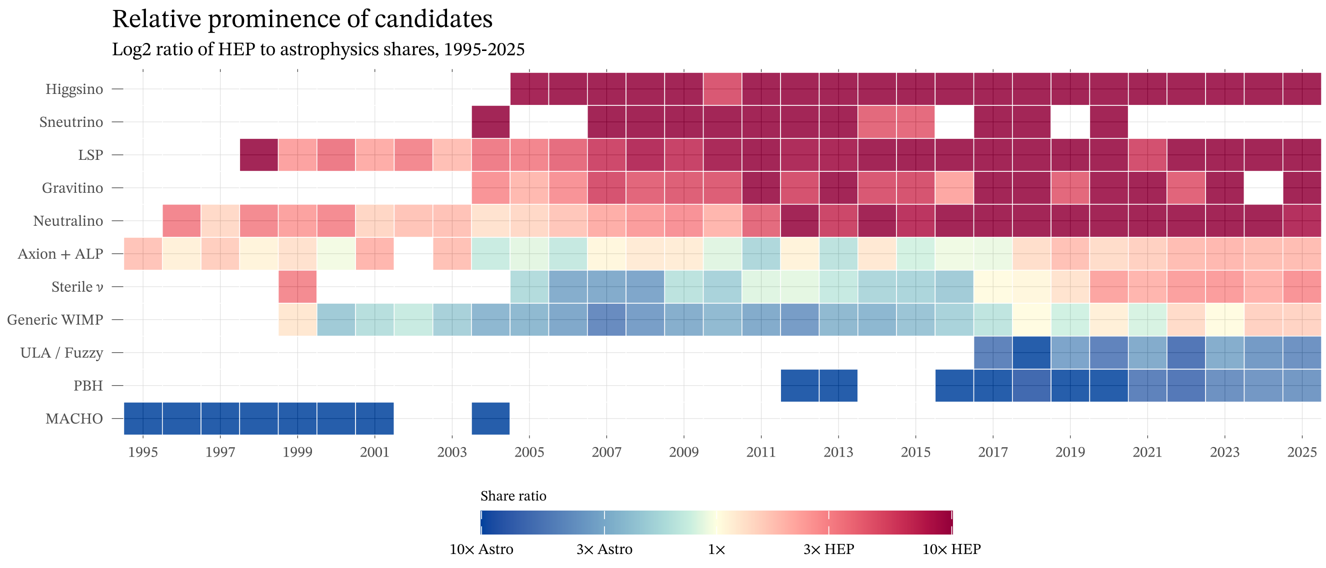

Figure 2 Candidate prominence differs between astrophysics and high-energy physics. Each cell shows the base-2 logarithm of a candidate family's share of tracked mentions in high-energy physics divided by its share in astrophysics for that year. Positive values indicate greater relative prominence in 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.

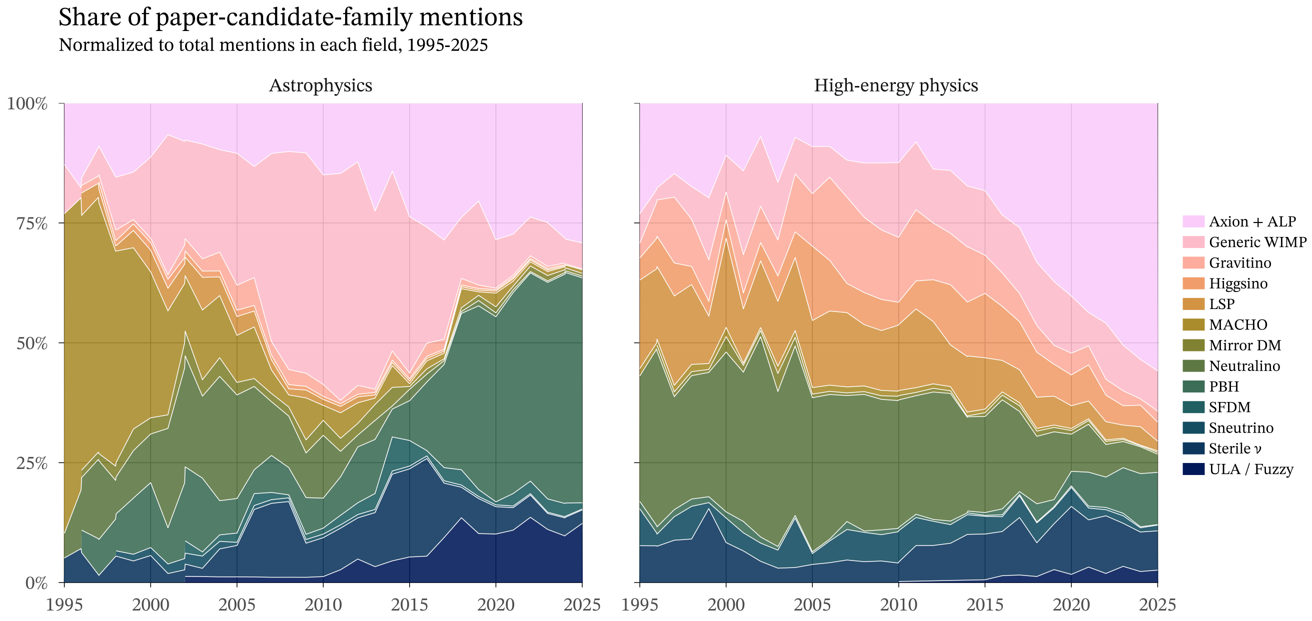

Figure 3 Candidate composition follows different field-specific trajectories. Annual paper–candidate-family mentions are normalised within astrophysics and high-energy physics. A paper can mention more than one family, and mentions do not imply endorsement. Raw counts, equal-paper weighting, and formal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

expands especially in astrophysics, where lensing, compact-object dynamics, and gravitational-wave observations constrain candidate populations [25–27]. The successor mixtures differ without becoming exclusive to either field.

Different successor mixtures do not produce evidence of 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 (Fig. S3). The predefined mean Jensen–Shannon divergence is 0.332 in 1995–2004 and 0.291 in 2016–2025, a late-minus-early contrast of -0.041. Early annual values are also supported by far fewer candidate-mentioning papers. Lexical trajectories separately show rising high-energy-physics terms such as *mediator* and *portal*. Lexical divergence after the predefined candidate-name exclusions remains an SI robustness check because its magnitude depends more strongly on preprocessing and term selection.

The negative divergence result constrains rather than establishes the conceptual argument. It rules out a simple reading of "diverging paths" as steadily increasing bibliometric dissimilarity. Distributional similarity asks whether fields allocate attention across comparable candidate sets; functional complementarity asks whether their different results can be connected into one discovery claim. The first does not guarantee the second, but neither does the candidate-mention profile directly measure the second. The conceptual comparison below shows instead that leading successor programs can connect constitutive identification, cosmic attribution, and astronomical explanation in different ways. The fields have neither separated bibliometrically nor converged on one replacement with WIMP-comparable prominence. The post-WIMP condition is therefore interpreted as several asymmetrically developed paths through a still-shared problem.

Relative paths to discovery

These empirical patterns matter because candidate programs offer different routes through the inherited dark-matter problem. For a candidate c, four achievements must be distinguished. *Constitutive identification*, $I(c)$, establishes what kind of particle, field, object, population, or physical structure constitutes the proposed component. *Cosmic attribution*, $A(c)$, establishes that c exists with the abundance, production history, stability, and distribution required of that component. *Explanatory adequacy*, $E(c)$, establishes that c accounts for the relevant astronomical and cosmological phenomena at the required scales. *Cross-domain identification*, $L(c)$, warrants that the entity identified in one evidential setting is the same component invoked in the attribution and explanatory claims. A strong dark-matter discovery claim therefore has the structure

$$D_{\mathrm{DM}}(c) = I(c) \wedge A(c) \wedge E(c) \wedge L(c). \tag{2}$$

This expression is a schema for the strong claim that c constitutes the relevant cosmic dark component; it does not define discovery of c simpliciter. If c supplies only a fraction f of the abundance, $A_f(c)$ records that partial attribution. A multicomponent claim replaces c with a set $C = \{c_1, \ldots, c_n\}$ and asks whether the attributed fractions jointly supply the relevant abundance and explanatory roles. Identifying one member of C can be a genuine particle or component discovery without closing the residual dark-matter problem.

The conjunction does not require its achievements to occur simultaneously, in one discipline, or in a fixed order. Nor is $L(c)$ a fourth independent observation. Although written with one argument for compactness, it is relational: it warrants that the entity identified through I is the same component to which the abundance and explanatory claims in A and E refer. For example, laboratory evidence might identify a particle with specified properties while cosmological evidence independently supports an otherwise compatible population; without evidence that the laboratory entity constitutes that population, the individual claims can be warranted while L remains open. This is the content not already supplied by I, A, or E alone.

The WIMP made these conditions appear jointly attainable. A robustly interpreted recoil signal could contribute to constitutive identification, but the signal would not by itself establish the candidate's production history, cosmic abundance, or explanatory reach. Those claims require halo, cosmological, and astronomical evidence, together with the inference that the locally detected entity is the component responsible for the cosmic phenomena. The distinctive promise of the WIMP program was that convergent evidence could potentially warrant all four conditions. No single WIMP-era result completed that chain automatically.

Axion and ALP searches illustrate a different evidential profile. A haloscope such as ADMX searches directly 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 new particle or field, thereby addressing $I(c)$, but the route to $A(c)$ would differ. A haloscope signal would provide candidate-specific evidence of local halo occupation; a helioscope signal would not by itself identify the particle as the cosmic dark component. Neither result alone would establish the required cosmological abundance, production history, distribution, explanatory adequacy, and cross-domain identification. A particle or field can therefore be discovered before it is discovered as dark matter.

The PBH profile reverses the particle-centred expectation. Evidence for an abundant, appropriately distributed PBH population could identify the kind of object that constitutes the dark component, establish its cosmic attribution, and account for relevant gravitational phenomena [26, 27]. Cross-domain identification would still require evidence that the characterised population is the same population responsible for those phenomena. If that connection were warranted, dark matter could be identified without fulfilling the beyond-Standard-Model particle project historically attached to its discovery. This is why constitutive identification must remain candidate-neutral.

Modified-gravity approaches differ more fundamentally. A successful theory could achieve explanatory adequacy for phenomena attributed to dark matter while removing the need for a material dark component in that domain [28–30]. In that case constitutive identification, cosmic attribution, and cross-domain identification are not incomplete stages of the same discovery; the claim that there is a component to identify has instead been revised. Such a result could resolve or dissolve part of the inherited problem without constituting a discovery of dark matter.

These profiles show why dark-matter discovery need not have a single evidential form. A result can be decisive within one candidate program while leaving closure of the jointly occupied problem open. A discovery claim must therefore state what has been identified, whether it constitutes all or part of the cosmic component, which phenomena it explains, and how those achievements have been connected.

Discussion

The post-WIMP transition is often described as diversification: more candidates, broader mass ranges, and new experiments. The bibliometric evidence establishes a bounded reorganisation of attention. Dark-matter research remains jointly occupied in publication and classification, its candidate vocabularies

Candidate-specific paths to dark-matter closure

A strong discovery claim requires I(c), A(c), E(c), and cross-domain identification L(c).

	Constitutive identification	Cosmic attribution	Explanatory adequacy	Cross-domain identification
Laboratory axion/ALP signal	Potentially New particle or field identified	Not by itself Abundance and distribution remain open	Not by itself Cosmic explanatory reach remains open	Not by itself Signal must be linked to cosmic component
WIMP recoil with convergent evidence	Potentially Requires robust signal interpretation	Potentially Requires halo and cosmological inference	Potentially Could account for attributed phenomena	Potentially Convergence could connect all evidence
Cosmologically significant PBH population	Potentially Population identified categorically	Potentially Population and abundance constrained	Potentially Could account for gravitational evidence	Potentially Same population must close the roles
Modified-gravity success	No component No dark constituent identified	No component No dark constituent attributed	Potentially May explain attributed phenomena	No component Resolves or dissolves the composition claim

Assessments indicate what each stipulated result could establish; they are not claims about candidate viability.

Figure 4 Candidate-specific paths to dark-matter closure. Columns distinguish constitutive identification, $I(c)$; cosmic attribution, $A(c)$; explanatory adequacy, $E(c)$; and the relational cross-domain identification, $L(c)$. Cells summarise what each stipulated result could establish; they are conceptual assessments rather than claims about candidate viability. A WIMP result could potentially connect all four only through convergent evidence. A laboratory axion or ALP signal leaves cosmic attribution and cross-domain identification open; a cosmologically significant PBH population could close the dark-matter problem without a new particle; and modified-gravity success may explain attributed phenomena while removing the need for a dark component.

overlap, and the 13-family candidate-mention profiles do not become progressively more divergent. These observations do not measure a change in discovery standards. Interpreted through the candidate profiles, however, declining WIMP prominence removes a former reason to presume that results produced in different parts of the enterprise will be cumulative by default. The WIMP connected a particle kind, a production mechanism, a cosmological abundance, a laboratory strategy, and an explanation of astronomical phenomena. Its decline makes the separability of those achievements analytically visible.

This distinction separates local closure from closure of the jointly occupied problem. A result achieves local closure when it settles the question posed within a candidate program: whether a signal identifies a new particle or field, whether a population supplies a specified abundance fraction, or whether a theory explains a specified class of phenomena. A strong claim to have identified the relevant cosmic dark component requires more. Constitutive identification, cosmic attribution, explanatory adequacy, and their cross-domain identification must be warranted together. The candidate profiles in Fig. 4 are therefore not stages in a hierarchy. They show how a locally decisive result can leave other parts of the inherited problem open.

Epistemic closure must also be distinguished from communal ratification. The first concerns whether the available evidence warrants the relevant identification, attribution, explanatory, and connecting claims. The second concerns whether the participating research communities accept that those claims jointly resolve the problem. This study identifies conditions under which the two may come apart, but it does not measure ratification. Classifications and candidate mentions cannot show how researchers would evaluate a future claim, whether disagreement would persist, or whether any community would exercise legitimate authority over the result.

Evidential practices make this distinction consequential. Particle physics commonly uses a 5σ excess as the convention for calling an effect a discovery [31]. Historical and methodological analyses show how this field-specific threshold emerged from shifting experimental standards and why its application requires case-specific judgement rather than the mechanical use of a statistic [32, 33]. Astronomical and cosmological attribution may instead require consistency across independent observations, population inferences, and physical scales; PBH assessments, for example, combine abundance with lensing, compact-object, dynamical, and gravitational-wave constraints [25–27]. These practices are neither uniform within fields nor competing universal rules. Cross-domain identification requires the relevant warrants to be connected, not one field's convention to displace the others.

The disciplinary payoffs are asymmetric. An axion or ALP result could vindicate a particle-search program by identifying a new entity while leaving its dark-matter attribution unresolved. A sufficiently comprehensive PBH result could identify the cosmic component while providing no new beyond-Standard-Model particle. Modified-gravity success could explain or reclassify part of the empirical problem while removing the need for a dark component in that domain. Cross-disciplinary participation can therefore persist even when success in one program no longer automatically advances the principal ambitions pursued through the others.

The earlier history of "dark matter" illustrates the contingency of such unity. Historical studies show that problems later treated as one dark-matter anomaly were assembled from previously distinct concerns [19, 20, 42]. Kelvin and Poincaré discussed unseen or extinguished stellar bodies [34, 35]; Kapteyn and Jeans connected additional matter to stellar dynamics [36, 37]; and Trumpler used "dark matter" for light-obscuring material [38]. Lundmark, Pannekoek, and Holmberg went further by joining absorption, dynamical mass deficits, and mass-to-light estimates within the same problem [39–41]. When interstellar absorption became a separate topic, the obscuration strand separated from the dynamical one without a single decisive redefinition [42]. This analogy does not predict another semantic split. It shows more modestly that one term can coordinate several problems only while the evidential connections among them remain productive.

The interdisciplinary lesson is limited. Shared candidates can coordinate communities that investigate different objects with different methods, but the resulting institutions may outlast the candidate that supplied a reason to treat their work as complementary. Persistent participation does not ensure cumulative results. The loss of presumed complementarity also need not produce institutional fragmentation. Cross-listing remains substantial, aggregated astrophysics remains larger, and paths to discovery can branch without increasing distributional divergence.

The analysis does not establish that researchers endorse the candidates they mention or that metadata transparently represents disciplinary identity. ADS and arXiv omit some venues; abstract availability varies across time and 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; none captures every source of measurement, historical, or interpretive uncertainty. The alternative classification rules, coverage tables, lexical exclusions, and document-type sensitivities bound these inferences. They do not convert publication patterns into evidence of authority, conflict, or communal closure.

Conclusion

The WIMP made laboratory particle detection and identification of the cosmic dark component appear to be stages of one discovery. Its decline makes clear that they are independently warrantable achievements. Dark-matter research remains jointly occupied in the bibliometric sense, but the record does not identify a successor with comparable prominence across the participating fields. The conceptual comparison further shows that leading programs organise different combinations of constitutive identification, cosmic attribution, explanatory adequacy, and cross-domain identification. It does not establish that no successor could ultimately connect them. "Diverging paths" describes this branching of possible resolutions, not increasing statistical separation between the fields' candidate-mention profiles.

The WIMP's decline therefore does more than enlarge the candidate list. It removes 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 remain open.

Materials and methods

Canonical corpus and reproducibility

The retained ADS snapshot was queried with `full:dark matter` and indexed on May 21, 2026. Records dated after 2025 are excluded before every analysis. Deduplication uses ADS bibcodes. A machine-readable `analysis_manifest.json` records the query, indexing timestamp, cutoff, exclusions, input and output row counts, SHA-256 hashes, software versions, and random seeds. Earlier duplicate tables are not active inputs: a clean rebuild from the retained ADS archive regenerates the canonical tables, and the manifest identifies and hashes the canonical inputs used by the analysis. The manifest also provides the headline sample sizes.

Classifications and primary-order audit

For each paper, the first listed ADS arXiv class defines the primary assignment. Any-class membership counts a paper in every represented broad field. Fractional allocation distributes weight $1/k$ across its k listed classes before aggregation. Astro–HEP cross-listing is the annual share assigned to both broad fields. The primary-order audit used a fixed seed to draw 500 multi-classified papers, stratified by year from 2010–2025, and compared ADS order with arXiv primary metadata. Agreement uncertainty uses a Wilson 95% interval. The SI reports the deterministic class-to-field mapping and classification coverage by year and document type.

Candidate extraction

Dictionary matching operates on normalised abstracts and reports paper-level mentions. The candidate ontology maps explicit names and abbreviations to canonical families; expressions with substantial non-dark-matter uses require nearby dark-matter context where appropriate. Each family contributes at most once per paper. Mentions can occur in statements of support, comparison, constraint, or exclusion and therefore do not imply endorsement. Implicit references and formulations 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 but do not constitute a performance estimate.

Candidate composition and robustness

For each year and field, candidate counts are normalised over the 13 tracked dictionary families. The SI 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 is the number of resamples, not 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; the exclusion list, rankings, thresholds, and selection rules are released in full.

Additional sensitivities

The SI reports abstract and classification coverage, articles/eprints-only results versus the broader eligible-document definition, raw versus normalised 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`. Its Zenodo concept DOI is `https://doi.org/10.5281/zenodo.20241725` [43]; the version-specific DOI for the article's archived release will be added when that release is deposited. 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 material is available in the accompanying SI PDF and reproduction bundle.

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”. *Zenodo* (2026). URL `https://doi.org/10.5281/zenodo.20241725`.

PNAS Nexus, 2026, , 1–4
Supplementary Information

Supplementary Information for “Diverging paths to dark-matter discovery”

Simon Allzén[1,2,3,*]

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

Abstract

Robustness analyses for classification, candidate measurement and composition, and lexical change supporting the accompanying Research Report.

Keywords dark matter, robustness, candidate extraction, supplementary information

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 records 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; the released document-flow table reports every type and count.

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; it does not make the alternative assignment rules redundant. Annual classification coverage ranges from 71.7% to 85.3% during 2010–2025 and does not decline monotonically; coverage is 85.3% in 2025.

Table 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; 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

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; `arxiv_class_mapping.csv` enumerates all 149 observed classes and their corpus counts.

Candidate measurement

Candidate extraction uses a frozen dictionary applied to normalised abstracts and counts paper-level mentions. The ontology maps

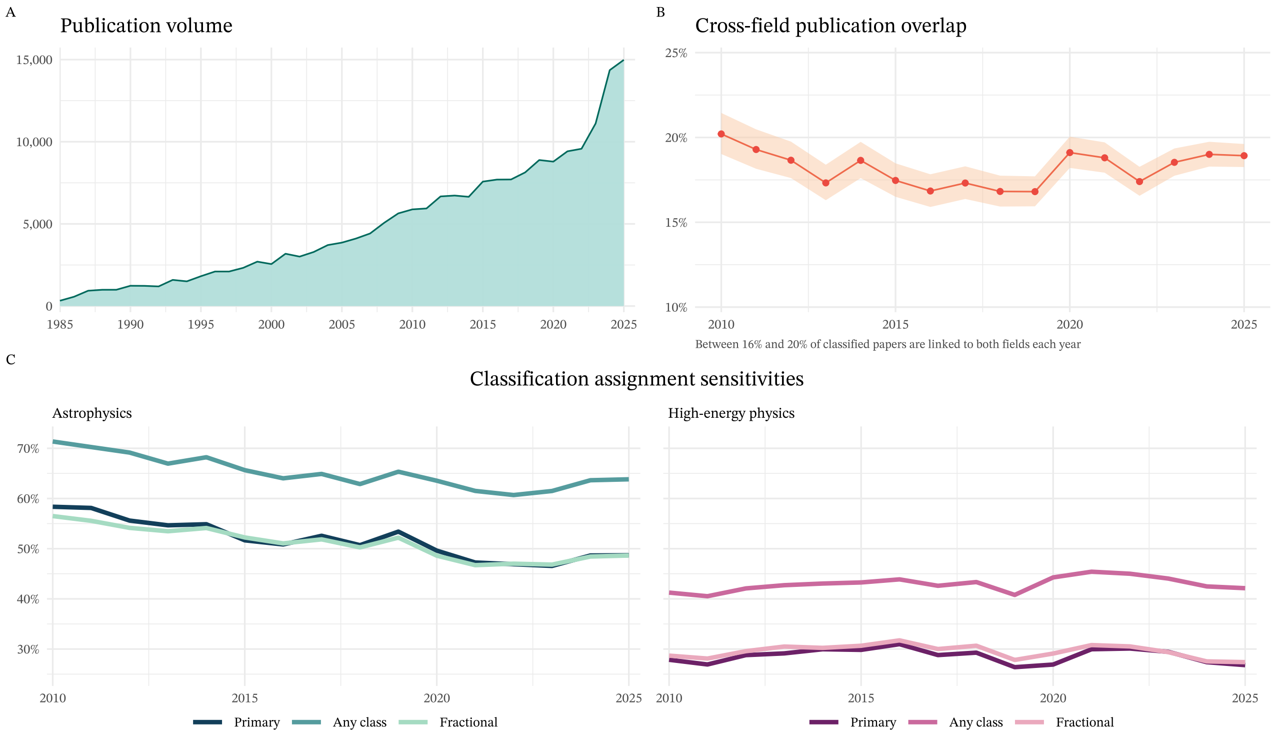

Figure S1 Corpus scale and classification sensitivity. (**A**) Annual publication volume in the frozen ADS corpus. (**B**) Annual share of papers linked to both astrophysics and high-energy physics under any-class membership, with a Wilson 95% interval. (**C**) Broad-field shares under primary, any-class,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

explicit names and abbreviations to canonical candidate families; expressions with substantial non-dark-matter uses require nearby dark-matter context where appropriate. Each family contributes at most once per paper. 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 in Table S2. The complete regular expressions, the 60-character context window on each side of context-dependent expressions, the wider ontology, and the canonical mappings are released with the analysis. Deterministic tests exercise canonical mentions, ambiguous abbreviations, experiment-name exclusions, and family collapse; these tests check implementation integrity rather than estimate population-level precision or recall.

Table 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`
Higgsino	`higgsino_dm`
LSP	`lsp`
MACHO	`macho`
Mirror DM	`mirror_dm`
Neutralino	`neutralino`
PBH	`pbh`
SFDM	`scalar_field_dm`
Sneutrino	`sneutrino_dm`
Sterile ν	`sterile_neutrino`
ULA / Fuzzy	`ula`, `fuzzy_dm`

Candidate-composition robustness

The main text normalises 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 normalising 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.

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

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

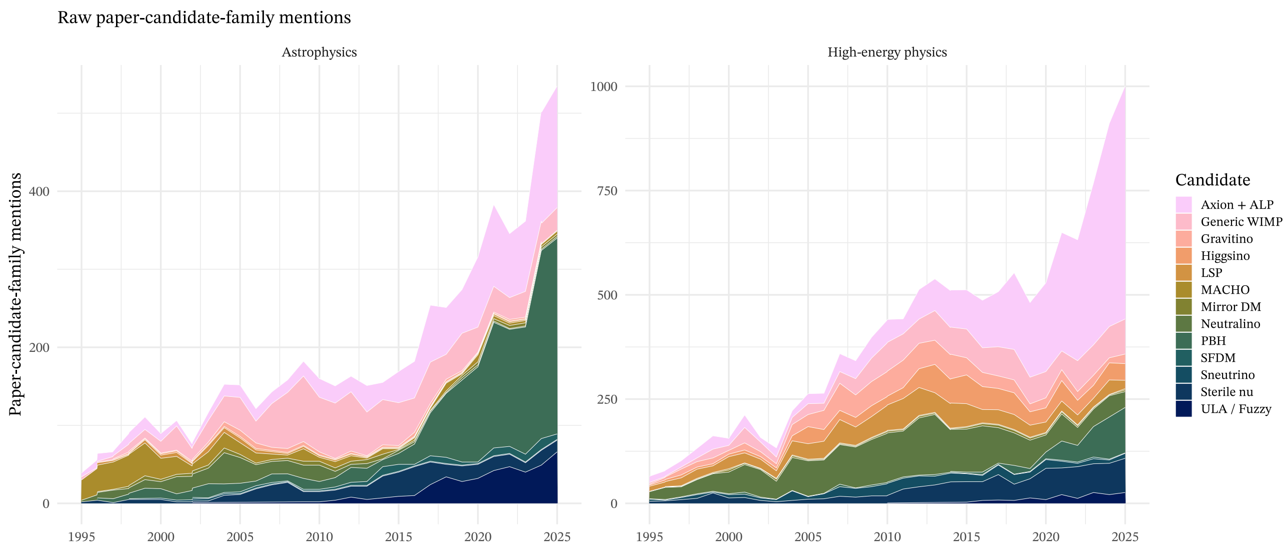

Figure S2 Raw annual paper–candidate-family mentions. Counts are shown separately for astrophysics and high-energy physics; the panels use separate vertical scales.

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 58 high-energy-physics papers mention a tracked candidate in 1995, compared with 487 and 880, respectively, in 2025.

The predefined comparison is mean annual JSD in 2016–2025 minus mean annual JSD in 1995–2004. Mean JSD declines from 0.332 to 0.291, giving a contrast of -0.041 (95% bootstrap interval -0.102 to -0.027). None of the three distance summaries supports increasing distributional divergence (Fig. S3). 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 low-cost specification checks. Giving every paper total weight one across all families it mentions strengthens rather than reverses the negative contrast. Of 2,566 papers with a generic-WIMP mention, 296 (11.5%) 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 downsampled without replacement to the minimum annual support (37 astrophysics and 58 high-energy-physics papers; 1,000 replicates). Mean rarefied JSD is 0.362 in 1995–2004 and 0.334 in 2016–2025, a contrast of -0.028 (simulation interval -0.078 to 0.024). Equalising support 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; observed-minus-null divergence is 0.292 early and 0.280 late and does not increase.

Table S3. Candidate-composition specification checks. Except for the final 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
Paper–candidate-family mentions	0.332	0.291	-0.041
Equal total weight per paper	0.352	0.282	-0.071
Overlapping generic WIMP suppressed	0.344	0.294	-0.050
Tracked WIMP families collapsed	0.263	0.200	-0.063
Corpus ending in 2024	0.332	0.302	-0.030

Restricting the corpus to articles and eprints produces the same directional comparison: mean JSD is 0.359 in 1995–2004 and 0.292 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. Candidate and lexical analyses remain conditional on abstract availability; volume and classification results do not.

Lexical robustness

Lexical data are annual paper-level document frequencies after normalisation 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; 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

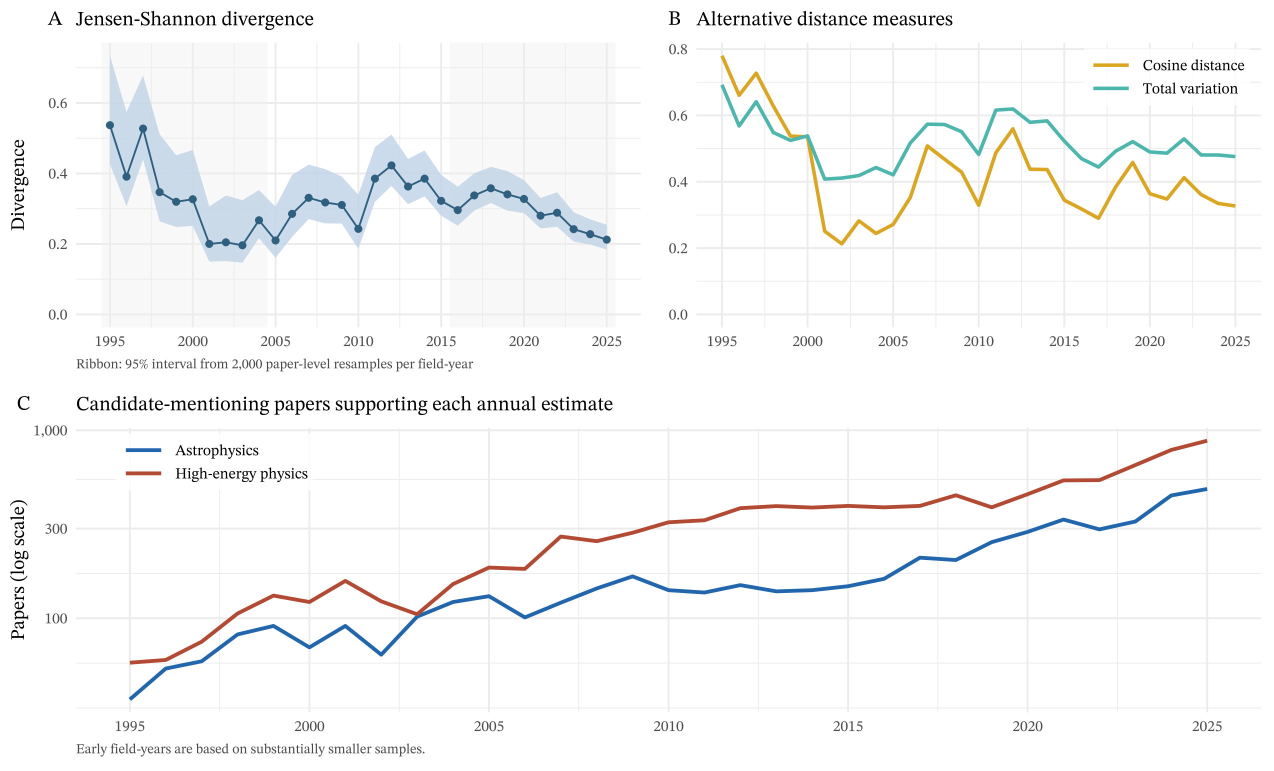

Figure S3 Candidate-composition divergence and supporting sample sizes. (A) Annual JSD with 95% intervals from 2,000 paper-level bootstrap replicates per field-year; shaded windows mark 1995–2004 and 2016–2025. (B) Total-variation and cosine-distance summaries. (C) Candidate-mentioning papers supporting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

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 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).

Reproducibility

The reproduction bundle contains the retained ADS snapshot, analysis manifest, primary-order audit, yearly source tables, figure data, and the six-stage pipeline. 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`; its archived releases are linked through the repository's Zenodo record.

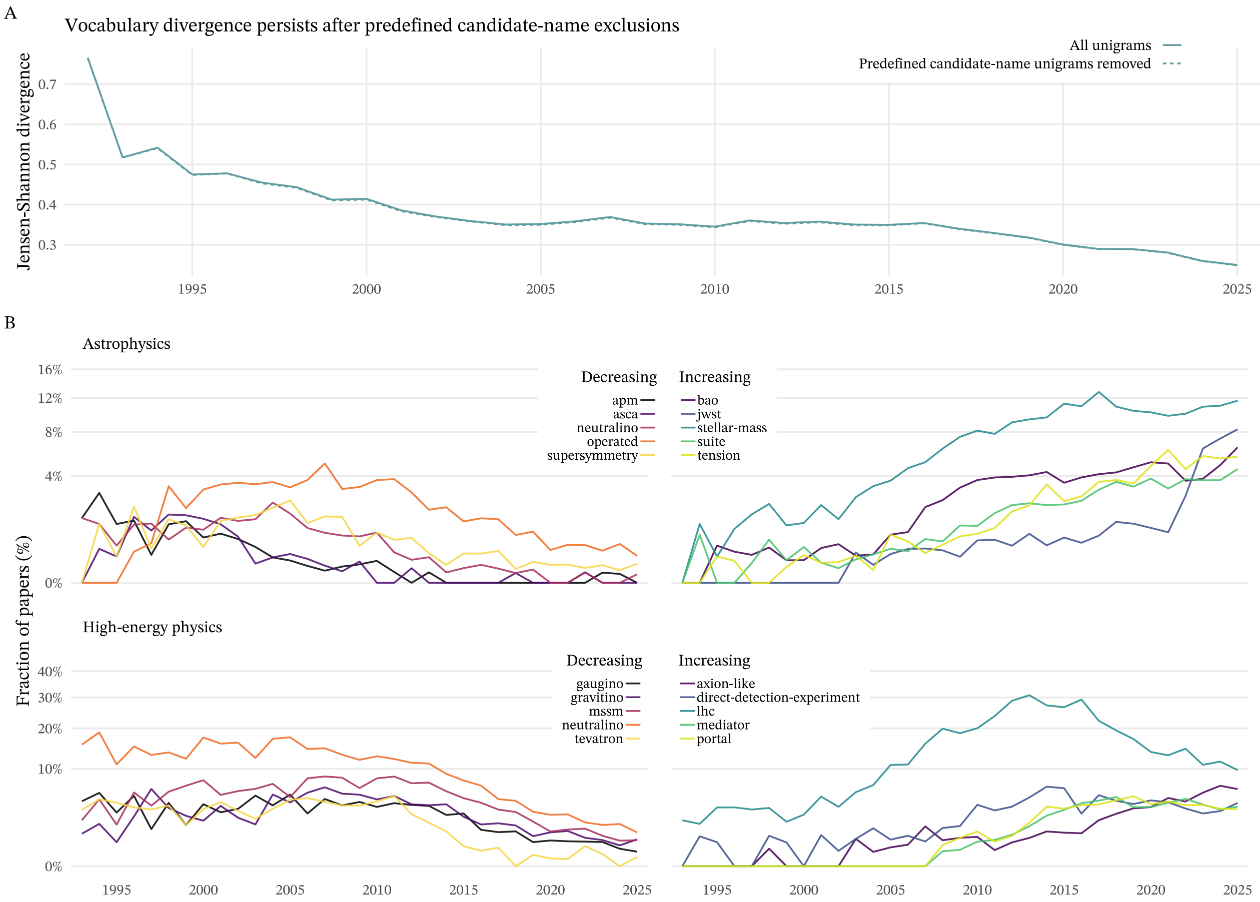

Figure S4 Lexical divergence and selected term trajectories. (**A**) Annual vocabulary JSD using all eligible unigrams and after removal of the 16 predefined candidate-name unigrams. (**B**) The five strongest data-selected declining and five strongest rising term trajectories in each field after the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