Bayesian Unification in the History of Dark Matter

Confirmation by Restructuring Old Evidence

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

December 15, 2025

Abstract. The history of dark matter is often glossed as a linear progression of accumulating evidence confirming its existence. The actual history, however, is far more complex and philosophically nuanced. In this paper, I revisit key moments in the establishment and justification of dark matter as a viable hypothesis and offer a Bayesian representation of the core inferences made at those moments. I argue, contrary to the textbook narrative, that certain evidential reclassifications played a pivotal confirmatory role even in the absence of new data. In particular, I show how unifying previously separate anomalous phenomena under a single modeling framework provided novel support for the dark matter hypothesis.

1 Introduction

The nature of eighty percent of the mass-energy content of the universe is unknown. Its presence is inferred from observed gravitational effects on ordinary matter across a range of cosmic scales. This mass is known as *dark matter*, and its existence forms a central tenet of the current standard model of cosmology — the ΛCDM (concordance) model.[1] Dark matter is used to explain the motions of galaxies, the dynamics of galaxy clusters, and, crucially, the overall distribution of matter in the universe. The dark matter problem is

*simon.allzen@philosophy.su.se ORCID: 0000-0002-6124-8152

[1]Λ denotes dark energy, and CDM stands for Cold Dark Matter, where 'cold' refers to dark matter's negligible free-streaming length.

one of the most fundamental open problems in current physics — making it fertile ground for philosophical analysis.[2]

So, how was this problem established? There is a typical portrayal of dark matter's history that one finds in textbooks and popular accounts. In the 1930s, Swiss astronomer Fritz Zwicky (1933) discovered that galaxies in the Coma cluster were moving too fast for the cluster to remain bound given the visible mass. To account for the missing mass, he posited the presence of "*dunkle materie*" (dark matter). This evidence remained largely ignored for decades. In the 1970s, astronomers Vera Rubin and Kent Ford, observing the rotation speeds of stars in spiral galaxies (notably Andromeda), found that rotation curves remained flat out to large radii instead of falling off as expected, implying the presence of unseen mass in galactic halos (Rubin & Ford Jr, 1970). According to the textbook narrative, Rubin and Ford's observations *confirmed* Zwicky's dark matter and finally convinced the community of dark matter's reality. The story continues with subsequent evidence (e.g., the Bullet Cluster's gravitational lensing or large-scale structure formation) accumulating to firmly establish the dark matter hypothesis.

Some variety of the above is the popular or 'textbook' history often presented (see, e.g., Ryden (2017), Sparke & Gallagher (2000)). This linear narrative of steadily accumulating evidence is not an entirely faithful account of how dark matter actually became an accepted idea. As Thomas Kuhn famously noted, textbooks tend to smooth over the intricacies of scientific development in favor of didactic storytelling. In the case of dark matter, a closer look at the historical details reveals a less straightforward and more philosophically interesting process.

This paper provides a philosophical analysis of a more accurate historical trajectory of dark matter's acceptance and examine what it can teach us about theory assessment and evidential reasoning. I build on historical work by de Swart et al. (2017) and J. de Swart (2020) who show that the establishment of the dark matter hypothesis pivots on two central 1974 papers, which took a novel approach to demonstrate how previously anomalous phenomena could jointly be seen as evidence for missing mass. From this episode, I extract (at least) two philosophically significant aspects of the reasoning involved.

The first is an intriguing discrepancy in how scientists from different communities (astronomy and cosmology) evaluated the evidence for dark matter in the mid–1970s. Both groups agreed on the empirical findings, yet cosmologists were far more receptive to the dark matter hypothesis at that time than astronomers. I argue that this divergence in evaluation can be explained by differences in background commitments: cosmologists carried certain extra-empirical assumptions that astronomers did not. I show that by incorporating these assumptions into the *priors* of a Bayesian analysis, we obtain different posterior probabilities for the dark matter hypothesis for cosmologists versus astronomers. This result demonstrates two things: (i) that a Bayesian framework can accurately model the

[2] See, e.g., Vanderburgh (2003, 2005, 2014b, 2014a), Sus (2014), Jacquart (2021), Merritt (2021a, 2021b), De Baerdemaeker & Boyd (2020), De Baerdemaeker (2021), De Baerdemaeker & Dawid (2022), Antoniou (2023, 2025), Duerr & Wolf (2023), Wolf & Read (2025), Allzén (2021, 2024), Martens (2022), Martens et al. (2022), Vaynberg (2024).

scientific reasoning exhibited in the dark matter case, and (ii) that extra-empirical beliefs (e.g., a preference for a 'closed' universe) influenced the perceived plausibility of the dark matter hypothesis from a cosmological perspective. This difference in priors can be understood as reflecting the different background assumptions (b) and modeling frameworks of the two communities, rather than being attributed to epistemic luck or arbitrariness. The second key aspect is the pivotal role played by *unification* in the 1974 papers. The epistemic merits of unification are a contentious topic in the philosophy of science, so it is worth examining whether the unificatory aspects in 1974 were genuinely confirmatory from a Bayesian perspective. To that end, I employ Myrvold (2003, 2017)'s Bayesian treatment of unification as an epistemic virtue. In the pair of papers, dark matter's ability to unify two seemingly disparate phenomena had a powerful effect on scientists' appraisal of the hypothesis. In particular, the 1974 studies introduced a unified *halo* model in which each spiral galaxy was posited to have an extended dark-matter halo characterized by a common parameter (for example, a roughly tenfold increase over the visible mass). This single modeling assumption tied together the previously separate anomalies of galaxy rotation curves and cluster dynamics into one piece of evidence. Such a reclassification of the observational data provided *confirmatory support* for the dark matter hypothesis. Typically, Bayesianism is considered unfit to accommodate cases of unification as confirmatory, but in this case, Myrvold's Bayesian account of unification characterizes the reconfiguration of observational data as confirmatory of the dark matter hypothesis because it *raises the mutual information between these phenomena* under the dark matter hypothesis — even though the evidence itself was not temporally novel.

In summary, the paper argues that the 1974 dark matter episode both (a) vindicates a Bayesian picture on which extra-empirical commitments legitimately show up as different priors across scientific communities, and (b) provides a concrete, historically important instance where unification of old evidence generates genuinely new confirmation in the strong Myrvoldian sense. This case both supports a specific way of modifying Bayesian confirmation theory (via mutual information unification and structural novelty) and clarifies how such modifications can be historically grounded rather than purely formal.

There are three main contributions of the paper: (i) a demonstration of how extra-empirical commitments ($\Omega = 1$, the cosmological principle, Machian preferences) can be encoded as priors in a way that respects differences across communities without reducing them to 'irrational bias'; (ii) the first detailed, historically grounded applications of Myrvold's MIU framework to a major case in cosmology, showing that the extra evidential 'boost' attributed to unification aligns with actual historical scientific practice; (iii) a refinement of the old evidence discussion by considering the dark matter case as a paradigmatic example of structural novelty without temporal novelty, and how MIU gives us a principled reason to treat such structural novelty as confirmatory.

2 The textbook history of dark matter

In the 1930s, Fritz Zwicky — often described as a maverick astronomer — studied the motions of galaxies in the Coma Cluster (a gravitationally bound collection of galaxies). Zwicky applied the virial theorem to the Coma Cluster and found that the cluster's visible matter (as inferred from galaxy luminosities) was significantly insufficient to account for the galaxies' velocities. Considering its luminous mass, the galaxies should have flown apart rather than remained bound as a cluster. In hindsight, this is often taken to be the first concrete evidence for some form of non-luminous matter in the universe.[3]

The next major step in the standard story jumps from the 1930s to the 1970s. As noted above, Vera Rubin (together with W.K. Ford) measured the rotation curve of the Andromeda galaxy (M31) out to large radii (Rubin & Ford Jr, 1970). Subsequent observations by Roberts & Rots (1973) , Bosma (1978) , and others confirmed that many spiral galaxies exhibited *flat* rotation curves: stars orbiting far from the galactic center were moving as quickly as those closer in, rather than slowing down (a flat curve versus a Keplerian declining curve). This behavior contradicted the expectations from visible matter alone (by analogy, planets in our solar system orbit slower at larger distances, and general relativity predicts a similar drop-off on galactic scales). In the popular narrative, Rubin's findings are portrayed as the moment when Zwicky's conjecture of dark matter was vindicated. In reality, Rubin (2004) herself refrained from speculating about the cause of the flat rotation curves, and many astronomers at first hoped to explain the anomaly without invoking exotic new mass. Nevertheless, over the following years additional evidence accumulated (e.g., the mass discrepancy in the Bullet Cluster's colliding gas and galaxies Clowe et al. (2006), gravitational lensing measurements, and consistency with large-scale structure formation models) that gradually solidified the case for dark matter in the eyes of most of the scientific community.

In summary, the textbook history presents dark matter's acceptance as the result of a linear progression of accumulated evidence: Zwicky's cluster observations; Rubin's galaxy rotation curves; large-scale structure formation; the Bullet Cluster observation; all adding up to confirmation. However, this streamlined linear account obscures important complications. In what follows, I delve into some of the historical sequences in more detail. In particular, I focus on the mid-1970s period, when disparate evidence was first *unified* under the dark matter hypothesis, as identified by de Swart et al. (2017) and J. de Swart (2020). We will see that this unification introduced a novel evidential structure that affected how the hypothesis was evaluated.

[3] At the time, Zwicky (1933, 1937) did use the term "*dunkle materie*" (i.e., dark matter) for this unseen mass, but, keeping in coherence with his contemporary astronomical community, the concept of *dark matter* bears little resemblance to our current concept (see Allzén (2024) for a historical overview).

3 Mass budget deficit?

To understand how the dark matter hypothesis gained traction in both cosmology and astronomy, we must consider developments in cosmology as a discipline leading up to the 1970s. In the early and mid-20th century, cosmology was a somewhat speculative cousin of observation-driven astronomy, where debates often invoked theory-driven arguments (J. de Swart, 2020). One example is the early acceptance of the so-called 'cosmological principle' (formulated by Milne (1932)), which asserts that the universe should appear homogeneous and isotropic on large scales to all observers. Today this principle has received empirical support, but in the 1930s it was adopted partly on philosophical (or 'rationalist') grounds. Coupling the cosmological principle with general relativity led to the field of *relativistic cosmology*, which by the 1960s became the predominant cosmological framework, based on theoretical work by Alexander Friedmann and Georges Lemaître (who provided solutions to Einstein's field equations for the universe) and observational work by Edwin Hubble (J. de Swart, 2020).[4] By the 1950s and 1960s, two philosophical desiderata shaped cosmologists' thinking:

- The cosmological principle (homogeneity/isotropy)
- A 'Machian' insistence that local inertia should be determined by the total mass of the universe (which, in Einstein's equations, suggests a *closed* finite universe).

Together these principles fostered a theory-first approach wherein empirical data were interpreted through the lens of general relativity's cosmological models. This idiosyncratic mode of reasoning led to internal tensions, here articulated by Steven Weinberg:

> If one tentatively accepts the result that ρ_0 is of order unity [$\Omega \geq 1$], then one is forced to the conclusion that a mass density of about 2×10^{-29}~g/cm^3 must be found somewhere outside the normal galaxies. But where? (Weinberg, 1972, p. 478)

Notably, the near-critical density estimate ($\rho_0 \approx 1$) from Weinberg (1972) was based on observational considerations rather than a theory-first commitment. At the same time, the idea of a closed universe was a philosophically motivated desideratum for many cosmologists — an inclination inherited from Mach and Einstein that provided an additional motivating aspect for the missing mass. In other words, by the early 1970s it was both an empirical *expectation* and a philosophical preference that the universe's total matter density Ω should be 1 (the "just closed" universe).

Against this background, two independent groups of researchers arrived at the same answer to Weinberg's question in 1974. Ostriker et al. (1974) and Einasto et al. (1974) looked to galactic astronomy for the mass needed to make $\Omega \approx 1$. Both groups compiled and reanalyzed existing data from two classes of phenomena:

[4] See Bondi & Gold (1948), who argued that one cannot assume *a priori* that general relativity applies globally without empirical justification — an objection to the early philosophically driven cosmological assumptions.

1. *Galaxy dynamics.* Thanks to radio observations (21-cm hydrogen line), astronomers could measure the rotation speeds of spiral galaxies out to radii well beyond the visible stellar disk. By the early 1970s, abundant galaxy rotation curve data were available. Expectations, based on Keplerian dynamics, were that rotation speeds should decrease with radius (a sloping curve) once beyond most of the galaxy's light (see Figure 1 adapted from Roberts & Rots (1973)). Empirically, however, many spiral galaxies showed approximately flat rotation curves — indicating that mass continues to increase with radius (or at least does not drop off), as if each galaxy sits inside an extended massive halo.

2. *Galaxy cluster dynamics.* Here, too, the data were long-standing: since Zwicky's time, various studies (e.g., Shapiro (1971)) had measured the velocities and mass-to-light ratios of galaxies in clusters, consistently finding that clusters appeared to contain far too little visible mass to be gravitationally bound. In other words, clusters exhibited a 'missing mass' problem analogous to galaxies.

Ostriker et al. (1974) and Einasto et al. (1974) took these two formerly separate sets of anomalous observations and *combined* them into a single hypothesis-driven analysis. They argued, essentially, that if each galaxy is embedded in a massive dark halo extending beyond its visible edge, then: (i) the galaxy's own rotation curve will be flat (explaining 1); (ii) a cluster of such galaxies will contain sufficient total dark mass (in the sum of all the extended halos) to account for the high galaxy velocities (explaining 2); (iii) the collective mass in galaxy halos could close the universe ($\Omega \approx 1$), achieving the cosmologists' desired mass density. At the start of their paper, Ostriker et al. (1974) write that "observations may be consistent with a Universe which is 'just closed' ($\Omega = 1$) — a conclusion believed strongly by some (cf. Wheeler 1973) for essentially non-experimental reasons." In other words, the central claim was that there are "reasons, increasing in number and quality, to believe that the masses of ordinary galaxies have been underestimated by a factor of 10 or more" (Ostriker et al., 1974, p. L1).

Each galaxy's mass might be *ten times* higher than its visible mass suggests, due to a dark halo. This unification of galaxy and cluster observations implied that a very large amount of unseen matter exists on galactic scales, and that this could solve both the cluster 'missing mass' problem and at the same time satisfy cosmologists' closure preference.

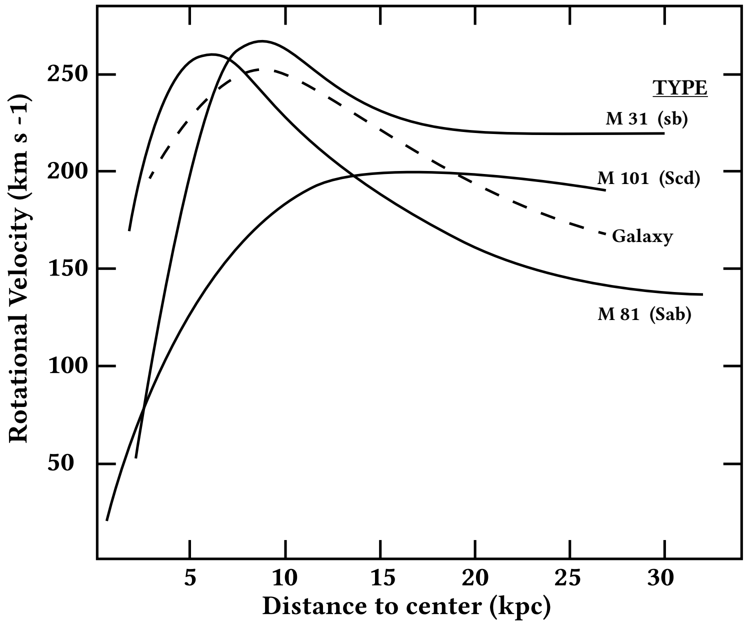

Figure 1: Rotation curves of three spiral galaxies (solid lines) as a function of distance from the center. (The Milky Way's expected Keplerian decline is shown for reference as the dotted line.) The flatness of the curves at large radii indicates the presence of unseen mass.

4 Extra-empirical beliefs and theory assessment

I will now turn to a Bayesian reconstruction of two philosophically significant aspects of the 1974 dark matter episode. First, I examine the divergence in assessment between astronomers and cosmologists regarding the dark matter hypothesis in the mid-1970s. Second, I analyze whether the unification of phenomena under a single modeling framework provided confirmatory support for the dark matter hypothesis.

4.1 Diverging assessments

There is an aspect of the history that warrants closer attention: astronomers and cosmologists in the 1970s (and into the early 1980s) differed in their willingness to accept the dark matter hypothesis, despite having access to the same evidence. In her retrospective *A Brief History of Dark Matter*, Rubin recalls that many astronomers were initially reluctant to embrace the idea that rotation curves stayed flat:

> In 1977, many astronomers hoped that dark matter might be avoided [...] there were still non-believers. One eminent astronomer said to me, 'When you observe low luminosity galaxies, you'll find Keplerian falling rotation curves.' Rubin (2004, p. 6)

And:

> Kalnajs' (1983) insistence that dark matter is not required, at least for a few galaxies with spatially limited data, convinced a few astronomers that dark matter could be avoided. In retrospect, we think it is fair to say that many astronomers hoped that Kalnajs was right; dark matter was to be avoided, if at all possible. Rubin (2004, p. 7)

Indeed, astronomers' aversion toward dark matter persisted for nearly a decade after 1974. This is striking, given the declaration by Ostriker, Peebles, and Yahil in 1974 that "there are reasons, increasing in number and quality, to believe that the masses of ordinary galaxies may have been underestimated by a factor of 10 or more." Why were astronomers less inclined than cosmologists to believe in the existence of this additional unseen mass?

One may attribute this divergence in assessment to differences in methodological orientation between the two communities. Astronomers in the 1970s generally insisted on strong direct empirical support before embracing new hypotheses — not all observed galaxy rotation curves at the time appeared sufficiently flat to necessitate the missing mass. Cosmologists, by contrast, were more guided by broad theoretical considerations (such as the cosmological principle or a preference for a closed universe), as long as those considerations remained consistent with the available data. This distinction between 'observation-minded' astronomers and 'theory-minded' cosmologists is echoed in Rubin, when she notes about the 1974 results that "the theorists had their eyes wide open" (Rubin, 2004, p. 3). This extra consideration that cosmologists had (the *philosophical* motivation to accept the missing mass) made the dark matter hypothesis prima facie attractive to cosmologists in a way it wouldn't have been for many astronomers.

4.2 Cosmologists' boosted priors

In Bayesian terms, we can characterize the situation as a difference in *posterior* belief $Pr(h \mid e)$ between the two groups, even though both had the same evidence e. For completion and clarity, we explicate the full Bayes' theorem (with an explicit term for the updated posterior $Pr_{\text{new}}(h)$), as:

$$\begin{aligned} Pr_{\text{new}}(h) \;=\; Pr(h \mid e) &= \frac{Pr(e \mid h)\, Pr(h)}{Pr(e)} \\ &= \frac{Pr(e \mid h)\, Pr(h)}{Pr(h)\, Pr(e \mid h) + Pr(\neg h)\, Pr(e \mid \neg h)} \end{aligned} \tag{4.1}$$

We take h as the dark matter hypothesis and e represents the relevant evidence (e.g., the set of observations of flat rotation curves and cluster dynamics). For Bayes' theorem to result in different *posterior probabilities* $Pr(h \mid e)$ for different individuals (astronomers vs. cosmologists), they must either assess the likelihoods $Pr(e \mid h)$ differently or assign different prior probabilities $Pr(h)$.

In this case, the evidence e (e.g., 'galaxy rotation curves are flat and cluster dynamics are inconsistent with visible mass') was equally known to both groups. It's reasonable to suppose both groups would agree on $Pr(e)$ and on the likelihood $Pr(e \mid h)$ — especially once the unified dark matter explanation was articulated. Both astronomers and cosmologists could acknowledge that if h (dark matter halos) is true, the evidence is expected so that $Pr(e \mid h)$ is high. Thus, the *empirical* evidence component $Pr(e)$ was effectively the same for everyone. The divergence in posterior, then, must come from the prior $Pr(h)$.

Cosmologists, due to their theoretical commitments, effectively had a higher prior $Pr(h)$ for the dark matter hypothesis than astronomers did. In a subjective Bayesian sense, cosmologists assigned more credence to the idea that there was extra mass in the universe even *before* considering the new evidence, because of their inclination toward a $\Omega = 1$ universe. Astronomers, lacking that inclination (and some perhaps hoping for other solutions like modified gravity or noisy data), had a lower prior for h. To illustrate: suppose (for the sake of argument) that a cosmologist's prior for "there exist substantial dark halos around galaxies" was $Pr(h) = 0.4$ in 1974, whereas an astronomer's prior was $Pr(h) = 0.15$. The evidence of flat rotation curves and cluster mass deficits (e) would update both, but the cosmologist's posterior $Pr(h \mid e)$ might then be much higher than the astronomer's. In essence, the cosmologists' extra-empirical beliefs *boosted the priors* for h.

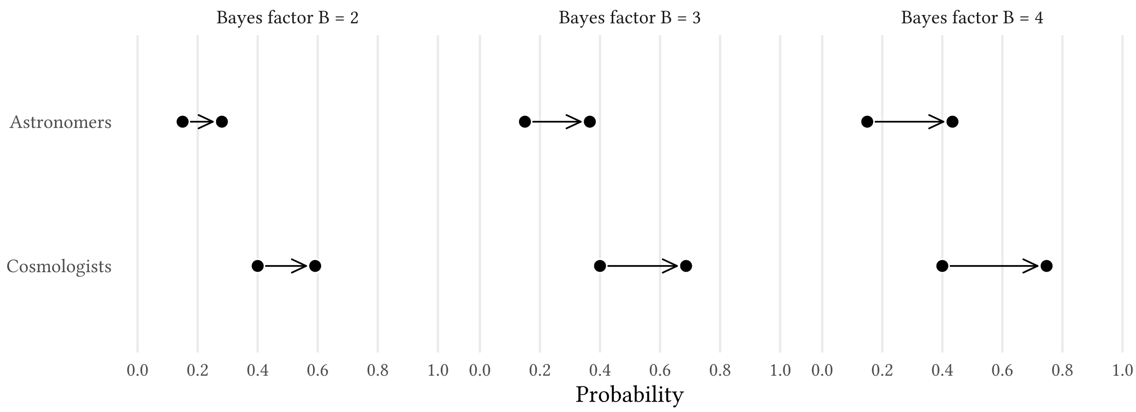

Figure 2: Updating on the same evidence e can yield different posteriors $Pr(h \mid e)$ if the prior $Pr(h)$ differs. Here we illustrate two groups (Astronomers and Cosmologists) with different priors (0.15 vs. 0.40) updating on the same evidence with varying Bayes factors (B = 2, 3, 4). The arrows indicate the update from prior to posterior for each group.

We can formalize this insight. If $Pr(e \mid h)$ is close to 1 (dark matter *predicts* the phenomena) and $Pr(e \mid \neg h)$ is small, then the Bayes factor is large. But if e was already known, one might set $Pr(e) = 1$, making the usual Bayes' update inert.[5] In such a scenario, it's the prior $Pr(h)$ that distinguishes agents. The $\Omega = 1$ preference, a theoretical, non-empirical consideration, gave cosmologists' h a non-zero prior weight that was not anchored in data but was nevertheless ultimately *correct.* The fact that this prior boost turned out to align with our current understanding of the composition of the universe means it wasn't simply

[5]That the evidence was already known is known as the *old evidence problem*, which I will address in a later section.

washed out or rendered insignificant. In hindsight, the cosmologists' non-empirical prior was on the right track.

This prior divergence clearly played a role in the historical reception of the dark matter hypothesis. The astronomers and cosmologists were effectively operating under different background contexts when evaluating the evidence, leading to different Bayesian updates. The lesson here is that scientists' theoretical and methodological commitments can translate into different prior probabilities, which in turn affect how evidence is weighed in practice. Normatively, the case described supports a moderately permissive but constrained view of extra-empirical considerations in scientific reasoning. Cosmologists' boosted prior for h was not an unconstrained preference for mathematical elegance, but was anchored in an already successful relativistic hot Big Bang framework, in the cosmological principle and related Machian ideas, and in theoretical arguments favoring an Einstein–de Sitter, $\Omega = 1$, universe — even though the best dynamical estimates at the time still placed the actual matter density at only $\approx$ 10–20% of the critical value. In that sense their stance was a risky but epistemically permissible heuristic rather than sheer bias.

5 Was the unification of phenomena confirmatory?

One of the central contributions of the 1974 papers by Ostriker et al. (1974) and Einasto et al. (1974) was precisely to *unify* several disparate data sets under a single explanatory model. By and large, the scientific community treated this unification as supporting the viability of the dark matter hypothesis. But was the act of unifying the phenomena itself *confirmatory*?

Traditionally, unification has been considered an *extra-empirical* theoretical virtue — at best guiding theory choice but not itself considered evidence in the Bayesian sense. For example, Psillos (1999, 2009) notes that if two competing theories H_1 and H_2 both fit the empirical data e, but H_1 can explain e with a single unified assumption while H_2 requires separate ad hoc assumptions for different parts of e, then scientists tend to view H_1 as better supported. The virtue of unification here is structural: it's about the relationship between evidence and hypothesis, not about new empirical content. A unified theory is epistemically appealing, yet by classical Bayesian lights, confirmation does not amount to more than the probability of evidence given the hypothesis.

Some philosophers (e.g., Kitcher (1989); Psillos (1999, 2009); Janssen (2002); Lipton (2003)) argue that explanatory power or unification can indeed contribute to a theory's support. Others — especially strict Bayesians like Howson & Urbach (2006) or Hartmann & Sprenger (2011) — maintain that a theory's epistemic merit is fully captured by the formal probability relation with evidence. From the latter perspective, unification shouldn't boost a hypothesis's probability unless it leads to a predictive difference. Given this debate, one might think that this turning point in the history of dark matter poses a challenge to Bayesian confirmation theory: unifying galaxy rotation curves and cluster dynamics was treated as providing (non-zero) epistemic support for dark matter, yet a staunch Bayesian might say no *new* evidence was introduced. Rubin, for instance, reflecting on

1974, wrote:

> Science often advances when ideas, formerly very disparate, are united. In retrospect, it took a long time for astronomers to relate Zwicky's dark matter to the flat rotation curves for some galaxies that were beginning to attract attention. (Rubin, 2004, p. 3)

Bringing Zwicky's cluster observations together with the rotation curve observations was described as a 'watershed moment' in our understanding of galaxies and cosmology Tremaine (1999, p. 1223). A Bayesian might retort that: either rotation curves and cluster dynamics were already known (in which case 'learning' them again adds no confirmation), or if they were independent and already confirm h individually, then together they *do* confirm h more strongly — but no more than the sum of their parts. Where, then, does the supposed extra boost from unification come from?

5.1 Bayesian unification

If the Bayesian wishes to accommodate the intuition that unification is confirmatory, there is a Bayesian framework they can resort to: Wayne Myrvold's *Mutual Information Unification* account. Myrvold (2003, 2017) proposes that a theory h is confirmed not only by how well it predicts each piece of evidence individually, but also by how much it makes the different pieces of evidence informatively *relevant* to each other. In simpler terms, if knowing one fact (p_1) gives you information about another fact (p_2) only when you assume h, then discovering both facts supports h above and beyond their separate evidential contributions. Myrvold defines a measure of unification in terms of the mutual information between pieces of evidence under h. The idea can be summarized as follows:

> Let $\mathbf{e} = \{p_1, p_2, \ldots, p_n\}$ be a body of evidence.
>
> Let $I(p, q \mid X)$ denote the *mutual information* between propositions p and q given background X.[6]
>
> If the phenomena are independent given background b, $I(p_1, p_2 \mid b) = 0$, but becomes informatively linked under hypothesis h and background b, $I(p_1, p_2 \mid h \wedge b) > 0$, then h makes p_1 and p_2 informatively connected.

The degree of unification $U(p_1, p_2; h \mid b)$ is defined as the increase in mutual information when h is assumed:

$$U(p_1, p_2; h \mid b) = I(p_1, p_2 \mid h \wedge b) - I(p_1, p_2 \mid b) \tag{5.1}$$

If p_1 and p_2 are independent without h ($I(p_1, p_2 \mid b) = 0$) but not independent with h, then $U = I(p_1, p_2 \mid h \wedge b)$. A positive U means h unified the phenomena by making them informatively connected. Myrvold proves that if you incorporate this notion into Bayesian

[6] Intuitively, how much knowing p reduces uncertainty about q, given X.

confirmation, the total support provided by e_1 and e_2 together can be decomposed into three parts: support from e_1 alone, support from e_2 alone, and an extra term corresponding to U. As a formula:

$$I(h; e_1 \wedge e_2 \mid b) = I(h; e_1 \mid b) + I(h; e_2 \mid b) + U(e_1, e_2; h \mid b) \tag{5.2}$$

where $I(h; e \mid b)$ denotes the information gain (the log of Bayes factor) from evidence e. This U term is precisely the mutual information unification measure. The question, then, is whether Myrvold's account applies to the dark matter case and whether unification indeed contributed epistemic support. To apply it, we identify:

h : the dark matter hypothesis (substantial unseen mass in extended halos).
p_1 : the galaxy cluster dynamics discrepancy.
p_2 : flat galaxy rotation curves.
b : background theory (general relativity, ordinary matter, etc.).
θ : halo parameter (e.g., halo-to-luminous mass ratio).

Let the background context b include at least general relativity and the existence of ordinary matter. Under the dark matter hypothesis h, both p_1 and p_2 are accounted for by a family of halo parameters, often summarized by an order-of-magnitude halo-to-luminous mass fraction (θ). Given only the background b, there is no particular connection between p_1 and p_2. Knowledge of galaxy rotation curves in the 1970s (flat vs. sloping) did not by itself inform the expectations about cluster masses, or vice versa, absent a linking hypothesis. Formally, $I(p_1, p_2 \mid b) = 0$; the phenomena were treated as independent puzzles.[7]

Introducing h changes this. Under h, both p_1 and p_2 are explained by the same physical assumption: extended dark matter halos around galaxies. The dark matter hypothesis therefore makes these phenomena informatively relevant to each other. If you learn (or assume) that galaxy rotation curves are flat, then given h you can predict that clusters made of such galaxies will show excess mass (because each galaxy contributes a halo to the cluster). Conversely, learning that clusters exhibit missing mass, under h, implies individual galaxies likely have massive halos (which would cause flat rotation curves). In Myrvold's terms, $I(p_1, p_2 \mid h \wedge b) > 0$; h generates a mutual informational relevance between the two phenomena. Thus, $U(p_1, p_2; h \mid b) > 0$ after 1974. The common halo-based parametrization effectively synthesized the two data sets, which had previously been viewed as unrelated. To stress this point, let θ denote a (schematic) halo mass–normalization parameter — for example an effective halo-to-luminous mass ratio of order ~ 10. Under h, both p_1 and p_2 are probabilistically constrained by θ; learning either

[7] For present purposes, I idealize the pre-1974 background b as comprising general relativity plus standard astrophysical practice, but excluding any explicit linking hypothesis between galaxy rotation data and cluster dynamics; in this sense, p_1 and p_2 were at the time treated as evidentially independent in the actual reasoning of astronomers and cosmologists.

therefore constrains θ and raises expectations about the other, generating $I(p_1; p_2 \mid h \wedge b) > 0$.

Before 1974, an astronomer might have considered flat rotation curves as one puzzle and cluster dynamics as another puzzle, without linking them to a common cause. After 1974, under h, these became manifestations of a single puzzle, with a single solution. Learning about one aspect gave information about the other. For instance, given h, the degree of cluster missing mass and the typical flatness of rotation curves are related quantitatively (both reflecting the halo mass fraction). Indeed, Ostriker et al. (1974) attempted to estimate the cosmic density parameter by assuming that galaxy halos account for cluster mass — showing how p_2 informs p_1 under h. This mutual informativity is the hallmark of confirmatory unification in Myrvold's framework.

(a) Pre-1974: treated as individual anomalies.

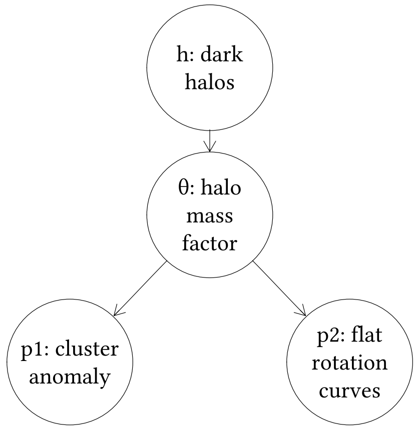

(b) Post-1974 under $h \wedge b$: unified via shared parameter θ.

Figure 3: Directed acyclic graph (DAG) illustrating structural reinterpretation/unification in 1974. Pre-1974: p_1 (cluster anomaly) and p_2 (flat rotation curves) were treated as unrelated. Post-1974: under h, both depend on a shared halo-parameter θ (an effective halo-to-luminous mass ratio), making p_1 and p_2 mutually informative. Arrows indicate a dependence/generative structure, not that observations causally affect θ.

Bayesian unification, therefore, provides a way to say that the 1974 reconfiguration of evidence was confirmatory. Even though p_1 and p_2 were not temporally novel data, the fact that they were unified by h meant the conjunction $e = \{p_1, p_2\}$ carried more evidential weight than the sum of its parts. Intuitively, the *structure* of the evidence changed: two facts that were previously viewed as independent (and hence unsurprising to learn together) became correlated under the hypothesis. This is why scientists experienced the unification as increasing their confidence in h. Myrvold's MIU formalism captures this by assigning a positive unification measure U.

To be explicit: a Bayesian can characterize the situation by claiming that, effectively, a *new prediction* was tested. Ostriker et al. (1974) and Einasto et al. (1974) predicted that if galaxies have massive halos, then combining galactic rotation data with cluster data should result in a consistent mass density Ω. They could be taken as treating the unified perspective as a new hypothesis h and 'predict' the observational consistency of p_1 and p_2 (which, retrospectively, was found to hold). On that view, the confirmation was not rendered by p_1 or p_2 individually (both were old), but by the fact that they *fit together* under h. This represents a use-novelty perspective, rather than temporal novelty.

The unification of galaxy and cluster phenomena in 1974 did provide confirmation for h — not by adding extra data points, but by changing the conceptual perception of the evidence. The mutual information between the two pieces of evidence increased given h, and that increase acted as an extra increment of support for h.

5.2 Summary of the Bayesian analysis

We can summarize the Bayesian perspective formally: Prior to unification we had evidence p_1 (cluster mass deficits) and p_2 (flat rotation curves) which were treated as unrelated. A Bayesian would update on each:

$$Pr(h \mid p_1 \wedge b) \text{ and } Pr(h \mid p_2 \wedge b) \tag{5.3}$$

Since p_1 was already known ($Pr(p_1) = 1$), it didn't update h when Rubin's p_2 came. After unification scientists effectively introduced a new composite evidence $e = \{p_1, p_2\}$ framed by h. Now, the evaluation concerned $Pr(h \mid p_1 \wedge p_2 \wedge b)$. Myrvold's framework gives us:

$$I(h; p_1 \wedge p_2 \mid b) = I(h; p_1 \mid b) + I(h; p_2 \mid b) + U(p_1, p_2; h \mid b) \tag{5.4}$$

with $U > 0$. The degree of support provided by the unified evidence is more than the sum of the supports from each piece alone, where the surplus is exactly the contribution of unification. This surplus can be seen as resolving an otherwise looming paradox: the 'old evidence' problem. Although p_1 and p_2 were old observations, the fact that they fit together under h was epistemically new. As long as one didn't already fully believe h or fully realize the connection, there was confirmation to be had.

In conclusion, from a Bayesian standpoint the 1974 unification can indeed be considered confirmatory, which aligns with the intuitions of scientists at the time. The dark matter hypothesis not only explained each phenomenon but also explained why the two phenomena occur together. This non-trivial linking of facts is what elevated confidence in the hypothesis, even though neither fact was temporally novel.

6 Old evidence and prior probabilities

As we saw, Myrvold's notion of unificatory confirmation revolves around an *increase* in mutual information within a body of evidence, $\mathbf{e} = \{p_1, p_2, \dots\}$, when a hypothesis h is introduced. However, an important caveat is that the standard Bayesian framework struggles when $Pr(\mathbf{e}) = 1$ or 0 (i.e., when the evidence is already known or logically entailed). This is the classic *problem of old evidence*, courtesy of Glymour (1981). In essence, if evidence e is fully known to an agent, then learning e does not change their credences: $Pr(h \mid e) = Pr(h)$. This means no confirmation, even if h was devised to explain e.

In our dark matter case, by the time massive halos was proposed in 1974, both p_1 and p_2 were *already observed* phenomena. Strictly speaking, $Pr(p_1) = Pr(p_2) = 1$ for the epistemic state of the proposers (i.e. Ostriker et al. (1974) and Einasto et al. (1974)). So a naive Bayesian should default to Bayes' theorem:

$$Pr(h \mid e) = \frac{Pr(e \mid h)\, Pr(h)}{Pr(e)} \tag{6.1}$$

However, since $Pr(e) = 1$ (because e is known), then (6.1) yields $Pr(h \mid e) = Pr(h)$. We cannot get any traction in our update from known evidence. We can see this by plugging in some arbitrarily set hypothetical numbers: take our previous prior $Pr(h) = 0.4$ for cosmologists who suspected extra mass. If h *predicts* the evidence, we might set $Pr(e \mid h) \approx 1$. And since e is known to occur, $Pr(e) = 1$. Then Bayes' theorem gives the unchanged:

$$Pr(h \mid e) = \frac{1 \cdot 0.4}{1} = 0.4 \tag{6.2}$$

The dark matter hypothesis would appear *not* to gain any confirmation from the evidence it was built to accommodate. This is the vexing issue pointed out by Glymour (1981): how can a hypothesis be confirmed by facts that were already known? The standard Bayesian workaround for old evidence is to adjust the formalism. One approach (advocated by Howson & Urbach (2006)) is to introduce a so-called *forgetful function*: pretend you didn't already know e, and update as if discovering it. Essentially, define a new conditional probability that 'forgets' the evidential import of e, then apply Bayes. While this solution is somewhat ad hoc (and arguably just pushes the issue into the priors), it allows one to say that if h entailed e and one hadn't taken e into account before, then observing e would raise $Pr(h)$.

In our case, one could argue that prior to 1974, nobody had yet fully treated p_1 and p_2 as combined evidence for a single hypothesis. Thus, one might *stipulate* a prior state in which h was considered without using p_1 and p_2. Then discovering that p_1 and p_2 (taken together) follow naturally from h would amount to confirmation. This is effectively what happened psychologically: before 1974, scientists did not recognize p_1 and p_2 as evidence for dark matter; after 1974, they did. The Bayesian can model this by saying the agents updated on e only after 1974, even though empirically they had recorded p_1 and p_2 earlier. As already foreshadowed, there is another way to think about this, involving the notion of *novelty*. We may distinguish between *temporal novelty* which denotes evidence not known before a theory is formulated, and *use novelty* denoting evidence not used in constructing the theory. Earman (1992) offers the following definitions:

> **Temporal novelty**: If T logically implies evidence statement E, and E was already known to be true before T was articulated, then E is not confirmatory for T.

> **Use novelty**: If T_1 and T_2 both imply E, but E was used in constructing T_1 (and not T_2), then E supports T_2 more than T_1. Earman (1992)

In the dark matter case, the evidence of flat rotation curves and cluster dynamics was *neither* temporally novel (it was known already) *nor* use-novel (it was explicitly used to formulate h). So, by these classical accounts, it should not confirm h at all: it was both 'old' and 'used'. This is why, to a Bayesian purist, the 1974 pivot might seem problematic.

Consulting Myrvold's account provides a resolution: the key is that the *structure* of the evidence was novel. The unificatory hypothesis h created a new evidential architecture in which p_1 and p_2 were connected. Although p_1 and p_2 themselves were not new facts, the realization that they could be understood as a unified package under one hypothesis was new. In Bayesian terms, even if $Pr(p_1) = Pr(p_2) = 1$, one might say $Pr(p_1 \wedge p_2) = 1$, but that does not capture the mutual information aspect. Myrvold would point out that what matters is $Pr(p_1 \wedge p_2 \mid h)$ in contrast with $Pr(p_1 \wedge p_2 \mid b)$. If the latter is just the product $Pr(p_1 \mid b)Pr(p_2 \mid b)$ (assuming independence under b) while the former is effectively $Pr(p_1 \mid b) \approx Pr(p_2 \mid b) \approx 1$ but also indicates a common origin, then we have $U > 0$. Even if both pieces of evidence were already available, we can consider the act of unifying them as generating a kind of *quasi-novelty*: the hypothesis made a successful connection that was not recognized before. We may refer to this as *structural novelty*: a change in the pattern of dependence among pieces of evidence (captured by mutual information), not a change in which propositions are known to be true. The crucial point is that the 1974 unification introduced a new evidential *structure* rather than simply reusing old data. By recasting the cluster and galaxy observations as manifestations of a single cause, the combined phenomenon explained by one hypothesis, the evidential situation was transformed in a way that standard Bayesian conditionalizing on each fact separately can not capture. So while the 'forgetful' strategy shows how one can, in principle, model confirmation by old evidence via a re-specified epistemic state, the MIU/structural novelty approach has the advantage of tracking the way scientists themselves reconfigured the evidential structure without needing to simulate ignorance.

Although the chronological order of hypothesis and evidence can matter psychologically, from a confirmation perspective, what really matters is whether the evidence was already taken for granted in the context of that hypothesis or not. The dark matter unification case shows that 'old' data can acquire fresh confirmatory import when placed in a new theoretical setting. Bayesianism can accommodate this by leveraging Myrvold's MIU measure to quantify the new structural relationship.

In conclusion, the unification in 1974 sidesteps the old evidence problem because it was not just plugging in previously known numbers; it was discovering a relationship between those numbers under a new hypothesis. This effectively provided *use novelty*, although not in method: the data were effectively 'new' evidence for h because h had not been explicitly constructed by hand to account for their unification (even if it was motivated by explaining each). The posterior probability of h ended up higher than its prior, not because new measurements were made, but because the evidential landscape was reconfigured. This is a nuanced point, but it highlights how scientific reasoning often goes beyond a simplistic Bayesian input-output model and involves an often side-lined aspect of science: scientists creative ability to reimagine and conceptually reconfigure the relation between theory and evidence.

7 Conclusion

In this paper, I have argued that an epistemically significant period in the history of dark matter can be illuminated by a Bayesian analysis once we include appropriate subtleties. Both of the main issues raised by this period – the diverging assessments by cosmologists and astronomers, and the confirmatory role of unification – find coherent explanations in Bayesian terms.

For the former, I showed that cosmologists' greater confidence in the dark matter hypothesis (circa 1974–1983) can be modeled as a higher prior probability $Pr(h)$ resulting from their extra-empirical background beliefs. Cosmologists, influenced by theoretical ideals like the $\Omega = 1$ and guided by relativistic cosmology's expectations, effectively assigned more credence to the missing mass hypothesis than astronomers did. This prior difference was not mere luck but reflected substantive differences in background context b. When formalized, plugging in a higher $Pr(h)$ for cosmologists resulted in a higher posterior given the same evidence — a result which came to fruition in practice. Bayesianism thus proves capable of representing how non-empirical (or extra-empirical) considerations shaped the evaluation of evidence, without endorsing those considerations as necessarily rational or not.

For the latter issue, the same extra-empirical commitments that raised cosmologists' priors for h also made them particularly receptive to the unificatory evidential structure that emerged in 1974. I used Myrvold's account of Bayesian confirmation by unification to demonstrate that the unificatory reconfiguration of galaxy and cluster phenomena in 1974 provided genuine confirmatory support for the dark-matter hypothesis. Although the data were not temporally novel, the hypothesis h induced structural novelty in the evidence by making the two sets of observations mutually informative. Qualitatively, I argued that this constitutes a new piece of evidence in its own right: the fact that two previously unrelated phenomena can be explained together. Using mutual information and the unification measure U, we saw that after 1974, evidence contributed more to confirming h than the sum of its parts did before. Thus, the much-debated 'extra epistemic value' of unification is captured in the formalism of Bayesian confirmation: unification can raise $Pr(h \mid e)$ above what it would otherwise have been. The dark matter case validates Myrvold's contention that the unificatory power, in this instance, was evidential.

Finally, I addressed technical concerns such as the problem of old evidence and how the notion of structural (as opposed to temporal) novelty can reconcile the Bayesian account with the historical fact that scientists felt more confirmed about h post-1974. We saw that while no new experiment was performed at that time, the conceptual move of unification functioned much like the prediction of a previously unobserved correlation, and in that sense it can be treated as 'new evidence' for the hypothesis. The Bayesian machinery, supplemented with a 'forgetful' update or framed in terms of innovation of use, can accommodate the situation.

In conclusion, the establishment of dark matter as a scientific hypothesis illustrates a rich dynamic of evidence, theory, and background belief. A Bayesian lens helps make sense of this dynamic by quantifying how evidence is assessed relative to varying prior commitments and by formalizing how combining pieces of evidence under a single hypothesis can amount in confirmation. This analysis underscores that confirmation in science is not purely about predicting new facts but often about finding new ways to interpret and connect facts that we already have. The dark matter case, far from being a linear progression of accumulating evidence, was a nuanced process involving shifts in perspective. Bayesian reasoning, when appropriately augmented to handle mutual information and context, proves to be a valuable tool for explicating these aspects of scientific reasoning.

References

Allzén, S. (2021). Scientific realism and empirical confirmation: A puzzle. *Studies in History and Philosophy of Science Part A*, *90*, 153–159.

Allzén, S. (2024). *Dark Matter: Explanatory Unification and Historical Continuity*. arXiv. https://doi.org/10.48550/arXiv.2412.13404

Antoniou, A. (2023). Robustness and Dark-Matter Observation. *Philosophy of Science*, *90*(3), 629–647.

Antoniou, A. (2025). Why did the dark matter hypothesis supersede modified gravity in the 1980s? *Studies in History and Philosophy of Science*, *112*, 141–152.

Bondi, H., & Gold, T. (1948). The steady-state theory of the expanding universe. *Monthly Notices of the Royal Astronomical Society*, *108*(3), 252–270.

Bosma, A. (1978). *The distribution and kinematics of neutral hydrogen in spiral galaxies of various morphological types* [{PhD} {Thesis}]. Rijksuniversiteit te Groningen.

Clowe, D., Bradač, M., Gonzalez, A. H., Markevitch, M., Randall, S. W., Jones, C., & Zaritsky, D. (2006). A Direct Empirical Proof of the Existence of Dark Matter. *The Astrophysical Journal Letters*, *648*(2), L109–L113. https://doi.org/10.1086/508162

De Baerdemaeker, S. (2021). Method-driven experiments and the search for dark matter. *Philosophy of Science*, *88*(1), 124–144.

De Baerdemaeker, S., & Boyd, N. M. (2020). Jump ship, shift gears, or just keep on chugging: Assessing the responses to tensions between theory and evidence in contemporary cosmology. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *72*, 205–216.

De Baerdemaeker, S., & Dawid, R. (2022). MOND and meta-empirical theory assessment. *Synthese*, *200*(5), 1–28.

Duerr, P. M., & Wolf, W. J. (2023). Methodological reflections on the MOND/dark matter debate. *Studies in History and Philosophy of Science*, *101*, 1–23. https://doi.org/https://doi.org/10.1016/j.shpsa.2023.07.001

Earman, J. (1992). *Bayes or Bust? A Critical Examination of Bayesian Confirmation Theory.* Cambridge, ma: MIT Press.

Einasto, J., Kaasik, A., & Saar, E. (1974). Dynamic evidence on massive coronas of galaxies. *Nature*, *250*(5464), 309–310.

Glymour, C. (1981). Theory and evidence. *Philosophy of Science*, *48*(3).

Hartmann, S., & Sprenger, J. (2011). Bayesian epistemology. In *The Routledge companion to epistemology* (pp. 609–619). Routledge.

Howson, C., & Urbach, P. (2006). *Scientific reasoning: The Bayesian approach.* Open Court Publishing.

Jacquart, M. (2021). ΛCDM and MOND: A debate about models or theory? *Studies in History and Philosophy of Science Part A*, *89*, 226–234.

Janssen, M. (2002). COI stories: Explanation and evidence in the history of science. *Perspectives on Science*, *10*(4), 457–522.

Kitcher, P. (1989). Explanatory Unification and the Causal Structure of the World. In P. Kitcher & W. Salmon (Eds.), *Scientific Explanation* (pp. 410–505). Minneapolis: University of Minnesota Press.

Lipton, P. (2003). *Inference to the best explanation.* Routledge.

Martens, N. C. (2022). Dark matter realism. *Foundations of Physics*, *52*(1), 16.

Martens, N. C., Sahuquillo, M. A. C., Scholz, E., Lehmkuhl, D., & Krämer, M. (2022). Integrating dark matter, modified gravity, and the humanities. *Studies in History and Philosophy of Science*, *91*, A1–A5.

Merritt, D. (2021a). Cosmological realism. *Studies in History and Philosophy of Science Part A*, *88*, 193–208.

Merritt, D. (2021b). MOND and Methodology. In *Karl Popper's Science and Philosophy* (pp. 69–96). Springer.

Milne, E. A. (1932). World structure and the expansion of the universe. *Nature*, *130*(3270), 9–10.

Myrvold, W. C. (2003). A Bayesian account of the virtue of unification. *Philosophy of Science*, *70*(2), 399–423.

Myrvold, W. C. (2017). On the evidential import of unification. *Philosophy of Science*, *84*(1), 92–114.

Ostriker, J. P., Peebles, P. J. E., & Yahil, A. (1974). The Size and Mass of Galaxies, and the Mass of the Universe. *The Astrophysical Journal*, *193*, L1–L4. https://doi.org/10.1086/181617

Psillos, S. (1999). *Scientific realism: How science tracks truth.* Routledge.

Psillos, S. (2009). *Knowing the structure of nature: Essays on realism and explanation.* Palgrave Macmillan.

Roberts, M., & Rots, A. (1973). Comparison of rotation curves of different galaxy types. *Astronomy and Astrophysics*, *26*, 483–485.

Rubin, V. C. (2004). A brief history of dark matter. In M. Livio (Ed.), *The Dark Universe: Matter, Energy and Gravity* (pp. 1–13). Cambridge University Press. https://doi.org/10.1017/CBO9780511536298.003

Rubin, V. C., & Ford Jr, W. K. (1970). Rotation of the Andromeda nebula from a spectro-

scopic survey of emission regions. *The Astrophysical Journal*, *159*, 379.
Ryden, B. (2017). *Introduction to Cosmology* (2nd ed.). Cambridge University Press.
Shapiro, S. L. (1971). The density of matter in the form of galaxies. *The Astronomical Journal*, *76*, 291.
Sparke, L. S., & Gallagher, J. S. (2000). *Galaxies in the universe: An introduction*. Cambridge University Press.
Sus, A. (2014). Dark matter, the equivalence principle and modified gravity. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *45*, 66–71.
Swart, de, Bertone, G., & Dongen, J. van. (2017). How dark matter came to matter. *Nature Astronomy*, *1*(3), 1–9.
Swart, J. de. (2020). Closing in on the Cosmos: Cosmology's Rebirth and the Rise of the Dark Matter Problem. In A. Blum, R. Lalli, & J. Renn (Eds.), *The Renaissance of General Relativity in Context. Einstein Studies, vol 16* (pp. 257–284). Birkhäuser, Cham. https://doi.org/10.1007/978-3-030-50754-1_8
Tremaine, S. (1999). Comments on Ostriker, Peebles, & Yahil (1974)" The Size and Mass of Galaxies, and the Mass of the Universe". *The Astrophysical Journal*, *525*, 1223.
Vanderburgh, W. L. (2003). The dark matter double bind: Astrophysical aspects of the evidential warrant for general relativity. *Philosophy of Science*, *70*(4), 812–832.
Vanderburgh, W. L. (2005). The methodological value of coincidences: Further remarks on dark matter and the astrophysical warrant for general relativity. *Philosophy of Science*, *72*(5), 1324–1335.
Vanderburgh, W. L. (2014a). On the interpretive role of theories of gravity and 'ugly'solutions to the total evidence for dark matter. *Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics*, *47*, 62–67.
Vanderburgh, W. L. (2014b). Quantitative Parsimony, Explanatory Power and Dark Matter. *Journal for General Philosophy of Science / Zeitschrift Für Allgemeine Wissenschaftstheorie*, *45*(2), 317–327. https://doi.org/10.1007/s10838-014-9261-9
Vaynberg, E. (2024). Realism and the detection of dark matter. *Synthese*, *204*(3), 82.
Weinberg, S. (1972). *Gravitation and Cosmology: Principles and Applications of the General Theory of Relativity*. John Wiley & Sons.
Wolf, W. J., & Read, J. (2025). *Navigating permanent underdetermination in dark energy and inflationary cosmology*. https://arxiv.org/abs/2501.13521
Zwicky, F. (1933). How Far Do Cosmic Rays Travel? *Phys. Rev.*, *43*(2), 147–148. https://doi.org/10.1103/PhysRev.43.147
Zwicky, F. (1937). On the Masses of Nebulae and of Clusters of Nebulae. *The Astrophysical Journal*, *86*, 217.