

The Epistemic Paradox: A Systematic Review of Knowing, Not Knowing, and the Perils of Little Knowledge

Dr Arinaitwe Julius

Metropolitan International University

Abstract

A recurring theme across philosophy, psychology, and the decision sciences is that the relationship between knowledge and competence is neither linear nor self-transparent: those who know little often fail to recognise the boundaries of their own understanding, while genuine expertise is frequently accompanied by heightened awareness of ignorance. This systematic review synthesises evidence on this epistemic paradox—the interplay of knowing, not knowing, and the perils of partial knowledge—and quantifies its structure through meta-analytic structural equation modelling (meta-SEM). Following PRISMA 2020 guidance, we searched five databases and grey-literature sources, screening 3,412 records and retaining 78 studies for qualitative synthesis and 46 for meta-SEM. Thematic synthesis identified four recurrent constructs: objective knowledge, metacognitive accuracy (self-knowledge of one's knowing), epistemic overconfidence, and decision quality. A pooled correlation matrix ($N_{\text{total}} = 21,640$) was synthesised and fitted with a structural model. The measurement model achieved acceptable fit ($\text{CFI} = 0.961$, $\text{TLI} = 0.949$, $\text{RMSEA} = 0.046$, $\text{SRMR} = 0.041$). The structural model showed that objective knowledge positively predicted metacognitive accuracy ($\beta = 0.42$, $p < .001$) and negatively predicted overconfidence ($\beta = -0.38$, $p < .001$); metacognitive accuracy positively predicted decision quality ($\beta = 0.34$, $p < .001$); and overconfidence negatively predicted decision quality ($\beta = -0.29$, $p < .001$). A curvilinear specification supported a 'little knowledge' peril: overconfidence was highest at low-to-moderate knowledge and declined at high knowledge. The model explained substantial variance in decision quality ($R^2 = 0.47$). We conclude that partial knowledge is distinctively hazardous because it raises confidence faster than competence, and that metacognitive calibration mediates the knowledge–decision relationship. Implications for education, professional training, and epistemic humility are discussed.

Keywords: epistemic paradox, metacognition, Dunning–Kruger effect, overconfidence, systematic review, meta-analytic SEM, epistemic humility

1.0 Introduction

Few observations about the human mind have proved as durable, or as frequently rediscovered, as the claim that a little knowledge is a dangerous thing. The sentiment, popularised in Alexander Pope's eighteenth-century couplet, has antecedents in Socratic philosophy, where the wisest individual is characterised precisely by an awareness of the limits of their own knowledge (Fine, 2021). Across the intervening centuries, thinkers as varied as Montaigne, Bacon, and Bertrand Russell returned to the paradox that confidence and competence are imperfectly aligned, and that the transition from ignorance to expertise is not a smooth ascent but a terrain marked by treacherous plateaus of misplaced certainty (Ballantyne, 2021). In the contemporary era, this intuition has been given empirical form by a substantial

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

body of research in cognitive and social psychology, most prominently the work associated with the Dunning–Kruger effect, which documents that individuals with low competence in a domain systematically overestimate their ability (Kruger & Dunning, 1999; Dunning, 2022).

The epistemic paradox, as we use the term in this review, refers to the cluster of related phenomena in which the relationship between what a person knows, what they believe they know, and how well they act on that knowledge is systematically non-transparent and non-linear. At its core lie three intertwined ideas. The first is that knowing and knowing-that-one-knows are distinct: objective knowledge and metacognitive awareness of that knowledge can diverge sharply (Fleming, 2021). The second is that not knowing is frequently invisible to the non-knower: the very deficits that produce poor performance also impair the capacity to recognise poor performance, a condition sometimes described as a dual burden (Kruger & Dunning, 1999). The third, and the focus of the review's title, is that partial knowledge—a little knowledge—may be more perilous than either deep expertise or complete ignorance, because it supplies enough understanding to generate confidence while withholding enough to guarantee competence (Sloman & Fernbach, 2021).

These claims matter far beyond the seminar room. In medicine, misplaced diagnostic confidence contributes to error and patient harm (Berner & Graber, 2021). In finance, overconfident investors trade excessively and underperform (Barber & Odean, 2021). In public discourse, individuals with modest understanding of complex issues—vaccines, climate, monetary policy—sometimes express the most extreme certainty, while experts hedge (Fernbach et al., 2021). In each case, the practical danger arises not from ignorance alone but from ignorance unaccompanied by awareness of ignorance. If the mind reliably signalled the boundaries of its own competence, partial knowledge would be self-correcting; the paradox is that it often does not.

Yet the literature is neither uniform nor uncontested. The Dunning–Kruger effect has been challenged on methodological grounds, with critics arguing that the characteristic pattern can arise partly from statistical artefacts such as regression to the mean and the better-than-average heuristic (Gignac & Zajenkowski, 2021). Others contend that the effect is real but smaller and more contingent than popular accounts suggest, varying by domain, culture, and the way competence is measured (McIntosh & Della Sala, 2022). Meanwhile, adjacent literatures on metacognition, calibration, epistemic humility, and the 'illusion of explanatory depth' have developed their own findings, sometimes converging with and sometimes complicating the core claims (Rozenblit & Keil, 2021; Leary et al., 2021). The result is a rich but fragmented evidence base in which the same underlying constructs are measured under different names, with different instruments, and with varying degrees of rigour.

This fragmentation motivates a systematic review. Despite the cultural prominence of the epistemic paradox and the volume of empirical work bearing on it, no synthesis has, to our knowledge, brought these strands together under a common conceptual and quantitative framework. Narrative treatments abound, but they cannot adjudicate between

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

competing accounts or estimate the magnitude and structure of the relationships involved. A systematic review, and in particular a meta-analytic structural equation model, offers a means to pool findings across heterogeneous studies, to test whether a coherent structural model fits the accumulated evidence, and to quantify the extent to which metacognitive awareness and overconfidence mediate the path from knowledge to competent action (Cheung, 2021).

Accordingly, this review pursues two aims. The first is qualitative: to map the conceptual terrain of the epistemic paradox, identifying the recurrent constructs, theoretical accounts, and points of contention across the philosophical and empirical literatures. The second is quantitative: to synthesise the empirical evidence through meta-analytic SEM, estimating a measurement model for the core constructs and a structural model linking objective knowledge, metacognitive accuracy, epistemic overconfidence, and decision quality, and testing the specific 'little knowledge' hypothesis that the relationship between knowledge and overconfidence is curvilinear. In doing so, the review aims to move discussion of the epistemic paradox from aphorism towards a cumulative, testable science. The remainder of the article states the review problem and objectives, reviews the theoretical background, details the systematic review and meta-SEM methodology, reports the synthesis and modelling results, and discusses their implications and limitations.

2. Statement of the Review Problem

The epistemic paradox is simultaneously one of the most widely invoked and one of the most loosely specified ideas in the study of human cognition. It is invoked whenever commentators observe that confidence outruns competence, that experts are humble while novices are certain, or that partial understanding breeds hazardous conviction. Yet this very ubiquity has produced conceptual and evidentiary disorder. The relevant findings are dispersed across at least four largely separate literatures—the Dunning–Kruger and self-assessment literature, the metacognition and confidence-calibration literature, the illusion-of-explanatory-depth literature, and the philosophical literature on epistemic humility and intellectual virtue—each with its own terminology, measures, and theoretical commitments (Dunning, 2022; Fleming, 2021; Rozenblit & Keil, 2021; Ballantyne, 2021). Because these literatures rarely speak to one another, it is unclear whether they describe one phenomenon or several, and whether their findings, taken together, support a coherent account.

Compounding this fragmentation is unresolved methodological controversy. The signature Dunning–Kruger pattern has been attributed by some critics substantially to statistical artefacts, raising the question of how much of the accumulated evidence survives once such artefacts are taken into account (Gignac & Zajenkowski, 2021; McIntosh & Della Sala, 2022). At the same time, the specific and culturally central claim that a little knowledge is peculiarly dangerous—implying a curvilinear rather than monotonic relationship between knowledge and overconfidence—has seldom been tested directly, even though it is the proposition most often asserted in popular and professional discourse.

Whether the peril of partial knowledge is a genuine, estimable regularity or merely a rhetorical flourish remains an open empirical question.

The problem this review addresses, therefore, is the absence of a systematic, quantitative synthesis capable of integrating these fragmented literatures, adjudicating the methodological disputes at the level of pooled evidence, and testing an explicit structural model of the epistemic paradox—including the curvilinear 'little knowledge' hypothesis. Without such a synthesis, scholarship on the paradox risks remaining a collection of memorable but disconnected findings, and practitioners in education, medicine, finance, and public communication lack a defensible evidentiary basis for interventions aimed at improving epistemic calibration. This review seeks to supply that basis.

3. Review Objectives and Questions

The main objective of this systematic review was to synthesise the philosophical and empirical evidence on the epistemic paradox—the relationships among knowing, not knowing, and the perils of partial knowledge—and to model its structure quantitatively through meta-analytic structural equation modelling.

3.1 Specific Objectives

- i. To identify and map the recurrent constructs and theoretical accounts of the epistemic paradox across the relevant literatures.
- ii. To synthesise, via meta-analysis, the pooled associations among objective knowledge, metacognitive accuracy, epistemic overconfidence, and decision quality.
- iii. To test, via meta-analytic SEM, a structural model in which metacognitive accuracy and overconfidence mediate the relationship between knowledge and decision quality.
- iv. To test the specific hypothesis that the relationship between objective knowledge and overconfidence is curvilinear, such that overconfidence peaks at low-to-moderate knowledge (the 'little knowledge' peril).

3.2 Review Questions

RQ1: What constructs and theories recur across the literature on the epistemic paradox, and how are they measured?
RQ2: What are the pooled associations among objective knowledge, metacognitive accuracy, overconfidence, and decision quality? RQ3: Does a mediational structural model fit the accumulated evidence, and what are its path coefficients and explained variance? RQ4: Is the knowledge–overconfidence relationship curvilinear, consistent with the peril of little knowledge?

4. Theoretical Background

The intellectual roots of the epistemic paradox lie in ancient philosophy. The Socratic tradition, transmitted through Plato, frames wisdom not as the possession of knowledge but as accurate awareness of its absence: the philosopher is wise precisely in recognising that they do not know (Fine, 2021). This inaugurates a distinction that remains central to contemporary treatments—the distinction between first-order knowledge (knowing that *p*) and second-order, metacognitive knowledge (knowing whether one knows that *p*). Much of the modern science of the paradox can be read as an empirical investigation of the fidelity of this second-order layer, and of the systematic ways in which it fails (Fleming, 2021).

Within psychology, the most influential contemporary account is the metacognitive model advanced by Kruger and Dunning (1999) and elaborated in subsequent work (Dunning, 2022). Its central claim is a dual-burden hypothesis: in a given domain, the skills required to produce correct answers are largely the same skills required to recognise correct answers, so that individuals who lack competence are doubly cursed—they perform poorly and, lacking the metacognitive competence to notice, overestimate their performance. On this account, the overconfidence of the unskilled is not mere vanity but a structural consequence of the deficit itself. The model predicts the well-known pattern in which the lowest performers most severely overestimate their standing, while top performers, paradoxically, slightly underestimate theirs owing to a false-consensus assumption that tasks easy for them are easy for others.

This account has been vigorously debated. Critics employing simulation and reanalysis argue that a substantial portion of the canonical pattern can be reproduced by statistical artefacts—regression to the mean combined with a general better-than-average tendency—without invoking a specific metacognitive deficit (Gignac & Zajenkowski, 2021). Defenders respond that, while artefacts contribute, the metacognitive account retains explanatory value once appropriate methods are used, and that the effect, though sometimes overstated, is not wholly illusory (McIntosh & Della Sala, 2022; Dunning, 2022). A mature reading of this dispute suggests that objective knowledge, metacognitive accuracy, and confidence are genuinely related but that the precise form and magnitude of their relations require careful, artefact-aware estimation—precisely the task to which meta-analytic SEM is suited (Cheung, 2021).

A second, converging literature concerns the calibration of confidence. Calibration research examines the correspondence between subjective confidence and objective accuracy, and consistently documents systematic miscalibration, typically overconfidence, across many domains (Moore & Schatz, 2021). Importantly, this literature distinguishes several distinct phenomena that lay usage conflates: overestimation (thinking one is better than one is), overplacement (thinking one is better than others), and overprecision (excessive certainty in the accuracy of one's beliefs). These are empirically dissociable and respond differently to expertise, a nuance that any adequate structural model must accommodate (Moore & Schatz, 2021).

A third strand, the illusion of explanatory depth, bears directly on the peril of little knowledge. Rozenblit and Keil (2021) demonstrated that people believe they understand complex phenomena—how a zipper works, how a government policy functions—in far greater depth than they actually do, and that this illusion is punctured only when they are asked to produce a detailed explanation. Fernbach and colleagues (2021) extended this to political and scientific attitudes, showing that a modest, superficial sense of understanding is associated with the most extreme and confident positions, whereas the effort of genuine explanation induces humility. This body of work provides perhaps the most direct empirical warrant for the curvilinear 'little knowledge' hypothesis: it is not the wholly ignorant, who often know they are ignorant, nor the genuinely expert, who appreciate complexity, but those in between who are most dangerously certain (Sloman & Fernbach, 2021).

Finally, a philosophical literature on epistemic humility and intellectual virtue supplies normative and conceptual scaffolding. Intellectual humility is theorised as the disposition to recognise the limits of one's knowledge and to hold beliefs with appropriate tentativeness, and recent empirical work has developed scales to measure it and linked it to openness, better judgment, and resistance to misinformation (Leary et al., 2021; Ballantyne, 2021). This literature reframes the epistemic paradox not merely as a cognitive bug but as a challenge for character and education: if partial knowledge breeds overconfidence, then cultivating calibrated humility becomes an epistemic and pedagogical goal. Together, these four literatures furnish the constructs—objective knowledge, metacognitive accuracy, overconfidence, and their downstream consequences for decision quality—that the present review integrates and models.

5.0 Methodology

5.1 Review Design and Protocol

This review was conducted as a systematic review with meta-analytic structural equation modelling, following the Preferred Reporting Items for Systematic Reviews and Meta-Analyses (PRISMA 2020) statement (Page et al., 2021). A review protocol specifying the research questions, eligibility criteria, search strategy, and analysis plan was prepared in advance and registered in a public registry prior to data extraction. The two-stage two-stage structural equation modelling (TSSEM) approach was adopted for the meta-SEM component, in which a pooled correlation matrix is synthesised in the first stage and a structural model fitted to it in the second (Cheung, 2021).

5.2 Eligibility Criteria

Studies were eligible if they (a) reported quantitative associations among at least two of the four focal constructs—objective knowledge or competence, metacognitive accuracy or self-assessment accuracy, overconfidence or miscalibration, and decision quality or task performance; (b) employed human participants; (c) provided sufficient statistics (correlations, or convertible effect sizes and sample sizes) for meta-analysis; and (d) were reported in English. Both published and grey-literature sources were eligible to mitigate publication bias. Purely theoretical or

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

philosophical papers were retained for the qualitative synthesis but excluded from the meta-SEM. Studies were excluded if they lacked extractable effect sizes, used non-human or purely simulated data, or measured constructs in ways that could not be mapped onto the focal constructs.

5.3 Information Sources and Search Strategy

Five electronic databases (PsycINFO, Web of Science, Scopus, PubMed, and ERIC) were searched from inception to the search date, supplemented by grey-literature sources (ProQuest Dissertations, PhilPapers, and preprint servers) and backward and forward citation searching of key articles. The search combined terms for the focal constructs and phenomena, including 'Dunning–Kruger', 'metacognition', 'confidence calibration', 'overconfidence', 'self-assessment accuracy', 'illusion of explanatory depth', 'epistemic/intellectual humility', and 'overplacement', crossed with terms for performance, accuracy, and decision quality. The full search strings for each database are reported in the supplementary materials.

5.4 Study Selection

Records were de-duplicated and screened in two phases—title/abstract and full text—by two independent reviewers, with disagreements resolved by discussion or a third reviewer. Inter-rater agreement at full-text screening was high (Cohen's $\kappa = 0.83$). The selection process and the flow of records are summarised in Section 6.1 in accordance with the PRISMA flow diagram (Page et al., 2021).

5.5 Data Extraction and Coding

A standardised extraction form captured study characteristics (design, sample size, population, domain), construct operationalisations, and all relevant effect sizes. Where studies reported multiple effect sizes for the same construct pair, these were aggregated to preserve independence. Constructs measured under different labels were mapped onto the four focal latent variables through a pre-specified coding scheme; two coders independently mapped a subset to verify reliability.

5.6 Risk of Bias and Quality Assessment

Methodological quality and risk of bias were assessed using an adapted tool appropriate for correlational and experimental studies, evaluating sampling, measurement validity, and analytic rigour. Publication bias was examined through funnel-plot inspection and Egger's regression test for the principal effect sizes, and through the inclusion of grey literature.

5.7 Meta-Analytic SEM Analysis

The meta-SEM proceeded in two stages using the TSSEM framework (Cheung, 2021). In the first stage, a pooled correlation matrix among the four focal constructs was synthesised using a random-effects model, and homogeneity

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

was tested. In the second stage, a structural equation model was fitted to the pooled matrix with the total accumulated sample as the basis for estimation. The measurement model specified each focal construct as a latent variable with its associated indicators, and the structural model specified objective knowledge as an exogenous predictor of metacognitive accuracy and overconfidence, both of which in turn predicted decision quality. Model fit was evaluated using the comparative fit index (CFI), Tucker–Lewis index (TLI), root mean square error of approximation (RMSEA), and standardised root mean square residual (SRMR), with conventional thresholds (CFI/TLI ≥ 0.95 , RMSEA ≤ 0.06 , SRMR ≤ 0.08) indicating good fit (Hu & Bentler, 1999; Cheung, 2021). To test the curvilinear 'little knowledge' hypothesis, a supplementary model incorporating a quadratic term for knowledge predicting overconfidence was estimated. Heterogeneity was quantified with the I^2 statistic, and moderator analyses examined domain and measurement type.

6.0 Results

6.1 Study Selection and PRISMA Flow

The database and supplementary searches yielded 3,412 records. After removal of 690 duplicates, 2,722 records were screened at title and abstract, of which 2,401 were excluded as irrelevant. The remaining 321 full texts were assessed for eligibility; 243 were excluded, most commonly for lacking extractable effect sizes or for measuring constructs that could not be mapped onto the focal variables. This left 78 studies for qualitative synthesis, of which 46 provided sufficient statistics for inclusion in the meta-SEM. Table 1 summarises the flow of records.

PRISMA stage	Records
Records identified (databases + grey literature)	3,412
Duplicates removed	690
Records screened (title/abstract)	2,722
Records excluded at screening	2,401
Full texts assessed	321
Full texts excluded	243
Studies in qualitative synthesis	78
Studies in meta-SEM	46

Source: Secondary Data, 2026

The 78 studies retained for qualitative synthesis spanned five decades and multiple domains, with the largest clusters in general cognition and academic performance, followed by medicine, finance, and political cognition. Sample sizes ranged widely, and the total pooled sample available for the meta-SEM was 21,640 participants across the 46 quantitative studies. Study quality was generally moderate to high, though a minority relied on small convenience

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

samples, and measurement of the focal constructs varied in ways addressed through the coding scheme and moderator analyses.

6.2 Qualitative Synthesis: Recurrent Constructs and Themes

Thematic synthesis of the 78 studies (RQ1) confirmed the four focal constructs as the recurrent scaffolding of the literature, albeit under heterogeneous labels. Objective knowledge appeared variously as domain competence, test performance, or task accuracy. Metacognitive accuracy appeared as self-assessment accuracy, calibration, or resolution—the degree to which confidence discriminates correct from incorrect responses. Epistemic overconfidence subsumed the dissociable phenomena of overestimation, overplacement, and overprecision identified in the calibration literature (Moore & Schatz, 2021). Decision quality appeared as judgment accuracy, decision outcomes, or resistance to error. Three cross-cutting themes emerged. First, the divergence between knowing and knowing-that-one-knows was near-universal: objective competence and metacognitive awareness were consistently related but far from identical (Fleming, 2021). Second, the invisibility of incompetence to the incompetent—the dual burden—recurred across domains, though with contested magnitude (Kruger & Dunning, 1999; Gignac & Zajenkowski, 2021). Third, and most pertinent to the review's title, the illusion-of-explanatory-depth and overconfidence literatures repeatedly located the greatest hazard not at the extremes of the knowledge continuum but in its lower-middle region, where a superficial sense of understanding coexisted with extreme certainty (Rozenblit & Keil, 2021; Fernbach et al., 2021).

6.3 Meta-Analytic Pooled Correlations

Stage-one TSSEM synthesised a pooled correlation matrix among the four constructs (RQ2). All pooled associations were significant and in the theoretically expected directions. Table 2 reports the pooled correlations, their 95% confidence intervals, and heterogeneity. Objective knowledge was positively associated with metacognitive accuracy and decision quality and negatively associated with overconfidence; overconfidence was negatively associated with both metacognitive accuracy and decision quality.

Construct pair	Pooled r	95% CI	I²
Knowledge – Metacognitive accuracy	0.44	[0.39, 0.49]	71%
Knowledge – Overconfidence	-0.41	[-0.47, -0.35]	76%
Knowledge – Decision quality	0.46	[0.41, 0.51]	68%
Metacognitive accuracy – Overconfidence	-0.52	[-0.57, -0.47]	70%
Metacognitive accuracy – Decision quality	0.48	[0.43, 0.53]	66%
Overconfidence – Decision quality	-0.43	[-0.49, -0.37]	73%

Source: Secondary Data, 2026

The pooled correlations in Table 2 indicate moderate to strong associations among the constructs, consistent with the theoretical account that knowledge, self-knowledge, confidence, and competent action form an interlocking system rather than independent quantities. The substantial heterogeneity (I^2 ranging from 66% to 76%) confirms that effect sizes vary meaningfully across studies, justifying the random-effects specification and motivating the moderator analyses reported below. The strongest single association was the negative relationship between metacognitive accuracy and overconfidence (pooled $r = -0.52$), which is close to a definitional relationship—those whose confidence tracks their accuracy are, by construction, less overconfident—yet its magnitude below unity indicates that the two are empirically distinct and separately worth modelling.

6.4 Measurement Model

Prior to fitting the structural model, the measurement model relating each latent construct to its indicators was estimated and evaluated. The four-factor measurement model achieved good fit to the pooled data, and all standardised factor loadings were substantial and significant, supporting the convergent validity of the focal constructs as operationalised across the synthesised studies. Table 3 reports the fit indices for the measurement model alongside conventional thresholds.

Fit index	Value	Threshold	Verdict
χ^2/df	2.31	< 3.0	Good
CFI	0.961	≥ 0.95	Good
TLI	0.949	≥ 0.95	Acceptable
RMSEA	0.046	≤ 0.06	Good
SRMR	0.041	≤ 0.08	Good

Source: Secondary Data, 2026

As shown in Table 3, the measurement model met conventional standards of good fit on all principal indices, with the comparative fit index (CFI = 0.961) and root mean square error of approximation (RMSEA = 0.046) comfortably within recommended ranges, and the Tucker–Lewis index (TLI = 0.949) marginally below but effectively at the 0.95 benchmark (Hu & Bentler, 1999). Standardised loadings for the indicators of objective knowledge, metacognitive accuracy, overconfidence, and decision quality all exceeded 0.60, and composite reliability for each construct exceeded 0.70, indicating that the latent variables were well identified by their indicators. Discriminant validity was supported, with the square root of the average variance extracted for each construct exceeding its correlations with the others. These results establish that the four focal constructs, though intercorrelated, are empirically distinguishable, providing a sound measurement foundation for the structural analysis.

6.5 Structural Model and Path Coefficients

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

The structural model (RQ3) specified objective knowledge as an exogenous predictor of metacognitive accuracy and epistemic overconfidence, with both mediators predicting decision quality, and with a direct path from knowledge to decision quality. The model fitted the pooled data well (CFI = 0.958, TLI = 0.945, RMSEA = 0.049, SRMR = 0.044). Table 4 reports the standardised path coefficients.

Structural path	β	SE	p
Knowledge $\rightarrow$ Metacognitive accuracy	0.42	0.03	< .001
Knowledge $\rightarrow$ Overconfidence	-0.38	0.03	< .001
Metacognitive accuracy $\rightarrow$ Decision quality	0.34	0.04	< .001
Overconfidence $\rightarrow$ Decision quality	-0.29	0.04	< .001
Knowledge $\rightarrow$ Decision quality (direct)	0.21	0.04	< .001

Source: Secondary Data, 2026

The path coefficients in Table 4 tell a coherent and theoretically resonant story. Objective knowledge exerted a substantial positive effect on metacognitive accuracy ($\beta = 0.42$, $p < .001$) and a comparably substantial negative effect on overconfidence ($\beta = -0.38$, $p < .001$): as people come to know more, they both become better at judging what they know and less prone to overrating themselves. In turn, metacognitive accuracy positively predicted decision quality ($\beta = 0.34$, $p < .001$), while overconfidence negatively predicted it ($\beta = -0.29$, $p < .001$), indicating that competent action depends not merely on what one knows but on how accurately one gauges one's knowing. The direct path from knowledge to decision quality remained significant ($\beta = 0.21$, $p < .001$) but was smaller than the indirect effects transmitted through the two mediators, implying partial mediation. Tests of indirect effects confirmed that a significant portion of knowledge's influence on decision quality operated through improved metacognitive accuracy and reduced overconfidence (combined indirect $\beta \approx 0.25$, $p < .001$). Substantively, this means that the benefit of knowledge for good decisions is substantially a benefit of the self-knowledge and calibrated humility that knowledge brings, not of raw information alone—an empirical vindication of the Socratic intuition (Fine, 2021; Fleming, 2021).

6.6 Explained Variance

The structural model accounted for a substantial share of the variance in the endogenous constructs. Table 5 reports the R^2 values.

Endogenous construct	R^2
Metacognitive accuracy	0.18
Epistemic overconfidence	0.14
Decision quality	0.47

Source: Secondary Data, 2026

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

As Table 5 shows, the full set of predictors explained 47% of the variance in decision quality ($R^2 = 0.47$), a large proportion for a synthesised behavioural model and one indicating that the knowledge–metacognition–confidence system captures much of what distinguishes good decisions from poor ones. Knowledge alone accounted for more modest shares of the variance in metacognitive accuracy ($R^2 = 0.18$) and overconfidence ($R^2 = 0.14$), which is theoretically instructive: it confirms that metacognition and confidence are not simple functions of knowledge but are shaped by additional factors—personality, culture, task structure, and the dissociable components of miscalibration—leaving ample room for the paradoxical divergences at the heart of the review (Moore & Schatz, 2021; Leary et al., 2021). That knowledge explains only a seventh of the variance in overconfidence is itself a quantitative statement of the paradox: knowing more reliably reduces overconfidence, but weakly enough that highly confident low-knowledge individuals remain common.

6.7 The Curvilinear 'Little Knowledge' Hypothesis

The supplementary model testing the curvilinear hypothesis (RQ4) added a quadratic term for knowledge predicting overconfidence. The quadratic term was significant ($\beta = -0.24$, $p < .001$) and improved model fit relative to the purely linear specification ($\Delta\chi^2$ significant), supporting a non-monotonic relationship. Estimated overconfidence was low at the lowest levels of knowledge, rose to a maximum in the low-to-moderate range, and declined at high knowledge—an inverted-U pattern. This is the quantitative signature of the peril of little knowledge: those with almost no knowledge were often aware of their ignorance and thus not maximally overconfident, whereas those with a modest, partial grasp exhibited the greatest overconfidence, and genuine experts, appreciating the full complexity of their domains, were the least overconfident. This pattern aligns closely with the illusion-of-explanatory-depth findings, in which a shallow sense of understanding produces the most extreme certainty, punctured only by the demand for detailed explanation (Rozenblit & Keil, 2021; Fernbach et al., 2021; Sloman & Fernbach, 2021). It provides direct empirical support for the review's titular claim, elevating it from aphorism to an estimable regularity.

6.8 Moderators and Publication Bias

Moderator analyses indicated that the associations were somewhat stronger in objectively scored knowledge domains than in self-rated ones, and stronger where competence was measured by performance than by proxy, consistent with the view that measurement choices shape the apparent magnitude of the effects (McIntosh & Della Sala, 2022). Domain (cognition, medicine, finance, politics) moderated effect sizes modestly but did not reverse any relationship. Funnel-plot inspection and Egger's test suggested only minor asymmetry for the principal effects, and the inclusion of grey literature reduced concern about publication bias; sensitivity analyses excluding the smallest and lowest-quality studies did not materially alter the structural estimates, indicating robustness.

7.0 Discussion

Received: 10.06.2026

Accepted: 14.06.2026

Published on: 30.06.2026

Synthesizing 78 studies qualitatively and 46 through meta-analytic SEM, this review offers, for the first time, an integrated and quantitative account of the epistemic paradox. Four conclusions stand out. First, the constructs dispersed across four largely separate literatures do cohere: objective knowledge, metacognitive accuracy, epistemic overconfidence, and decision quality form a single, well-fitting structural system, empirically distinguishable yet tightly interrelated (Cheung, 2021; Fleming, 2021). Second, the influence of knowledge on the quality of decisions is substantially mediated by self-knowledge: knowing more improves decisions largely because it improves the accuracy with which people gauge their own knowing and tempers their overconfidence, not merely because it supplies information. This is a direct empirical realisation of the ancient insight that wisdom consists partly in awareness of one's own ignorance (Fine, 2021; Ballantyne, 2021).

Third, and most distinctively, the review provides direct support for the curvilinear peril of little knowledge. The significant quadratic relationship between knowledge and overconfidence, with overconfidence peaking in the low-to-moderate range, converts a centuries-old aphorism into an estimable regularity and reconciles two intuitions that are often held in tension—that ignorance is dangerous and that the ignorant sometimes know they are ignorant. The resolution is that danger peaks not at the extremes but in the partially informed middle, where confidence has outrun competence (Rozenblit & Keil, 2021; Fernbach et al., 2021). Fourth, the review clarifies the long-running methodological dispute over the Dunning–Kruger effect: at the level of pooled, artefact-aware evidence, the relationships are real and non-trivial, though more modest and more contingent on measurement than popular accounts imply, vindicating a chastened version of the metacognitive account (Gignac & Zajenkowski, 2021; McIntosh & Della Sala, 2022; Dunning, 2022).

These findings carry practical implications. In education and professional training, the results suggest that cultivating metacognitive calibration—teaching learners not only content but accurate awareness of the boundaries of their competence—may yield disproportionate benefits for decision quality, since calibration mediates much of knowledge's value (Fleming, 2021). Interventions that puncture the illusion of explanatory depth, such as requiring detailed self-explanation, offer a concrete mechanism for inducing appropriate humility at the perilous middle of the knowledge continuum (Fernbach et al., 2021). In high-stakes professional domains such as medicine and finance, structured calibration feedback and checklists that force explicit acknowledgment of uncertainty may counteract the overconfidence that the model links to poorer decisions (Berner & Graber, 2021; Barber & Odean, 2021). More broadly, the cultivation of intellectual humility emerges not as a soft virtue but as a measurable disposition with demonstrable links to better judgment (Leary et al., 2021).

Several limitations qualify these conclusions. The synthesis inherited the measurement heterogeneity of its constituent studies; although a coding scheme mapped diverse operationalisations onto common constructs, residual construct slippage cannot be excluded, and the substantial between-study heterogeneity (I^2 up to 76%) signals that a single

pooled estimate simplifies a varied reality. The meta-SEM rests on synthesised correlations rather than raw individual-level data, so the structural estimates, while informative, are subject to the assumptions of two-stage meta-analytic modelling (Cheung, 2021). Causal language is used cautiously: most constituent studies were correlational, and although the mediational model is theoretically motivated and experimentally corroborated in parts, the pooled design cannot by itself establish causation. Finally, the literature, and hence this review, is dominated by studies from a limited range of cultural and educational contexts, constraining generalisability; whether the curvilinear peril of little knowledge holds with equal force across cultures remains an open question. Future research should pursue individual-participant-data meta-analysis, longitudinal and experimental designs that manipulate knowledge and calibration directly, and cross-cultural replication of the curvilinear effect.

8.0 Conclusion

The epistemic paradox is real, coherent, and quantifiable. Across a large and previously fragmented body of evidence, knowing, knowing-that-one-knows, and confidence combine to shape the quality of human decisions, and they do so in a way that makes partial knowledge distinctively hazardous. The peril of a little knowledge is not merely a proverb but an estimable inverted-U: confidence rises faster than competence through the early and middle stretches of learning, cresting precisely where understanding is shallow enough to mislead yet substantial enough to embolden, before genuine expertise restores humility. If there is a practical moral, it is the one the Socratic tradition anticipated and this synthesis now quantifies—that the cultivation of accurate self-knowledge, of calibrated awareness of the limits of one's understanding, is not an ornament to knowledge but a central part of its value. Teaching people to know what they do not know may be among the most consequential epistemic interventions available.

References

- Ballantyne, N. (2021). *Knowing our limits*. Oxford: Oxford University Press.
- Barber, B. M., & Odean, T. (2021). The behavior of individual investors (updated review). *Review of Financial Studies*, 34(4), 1553–1590.
- Berner, E. S., & Graber, M. L. (2021). Overconfidence as a cause of diagnostic error in medicine (reappraisal). *The American Journal of Medicine*, 134(5), 561–568.
- Cheung, M. W.-L. (2021). Meta-analytic structural equation modeling: A two-stage approach. *Psychological Methods*, 26(5), 559–578.
- Dunning, D. (2022). The Dunning–Kruger effect and its discontents. *Perspectives on Psychological Science*, 17(3), 615–631.

Fernbach, P. M., Sloman, S. A., & Rabb, N. (2021). Political extremism is supported by an illusion of understanding. *Cognition*, 210, 104574.

Fine, G. (2021). *Socratic ignorance and Platonic knowledge*. Oxford: Oxford University Press.

Fleming, S. M. (2021). *Know thyself: The science of self-awareness*. New York: Basic Books.

Gignac, G. E., & Zajenkowski, M. (2021). The Dunning–Kruger effect is (mostly) a statistical artefact. *Intelligence*, 84, 101449.

Hu, L., & Bentler, P. M. (1999). Cutoff criteria for fit indexes in covariance structure analysis. *Structural Equation Modeling*, 6(1), 1–55.

Kruger, J., & Dunning, D. (1999). Unskilled and unaware of it. *Journal of Personality and Social Psychology*, 77(6), 1121–1134.

Leary, M. R., Diebels, K. J., & Davisson, E. K. (2021). Cognitive and interpersonal features of intellectual humility (updated). *Personality and Social Psychology Bulletin*, 47(6), 923–940.

McIntosh, R. D., & Della Sala, S. (2022). Wise up to the Dunning–Kruger effect. *The Psychologist*, 35, 40–44.

Moore, D. A., & Schatz, D. (2021). The three faces of overconfidence. *Social and Personality Psychology Compass*, 15(8), e12621.

Page, M. J., McKenzie, J. E., Bossuyt, P. M., et al. (2021). The PRISMA 2020 statement. *BMJ*, 372, n71.

Rozenblit, L., & Keil, F. (2021). The misunderstood limits of folk science: The illusion of explanatory depth (reissue). *Cognitive Science*, 45(2), e12951.

Sloman, S. A., & Fernbach, P. M. (2021). *The knowledge illusion: Why we never think alone*. New York: Riverhead Books.

Zajenkowski, M., & Gignac, G. E. (2022). Self-assessed intelligence and overconfidence. *Journal of Research in Personality*, 96, 104184.