

THE SAFETY RADIUS OF SKILL : MODAL ROBUSTNESS AND THE DEFENSE OF INTELLECTUALISM

Feinan YE, Xia LI

ABSTRACT: The gradability of know-how presents a persistent puzzle for theories of skill. While knowing that p is typically binary, knowing how to ϕ admits of degrees of excellence. Leading Intellectualist accounts (Pavese 2017; Zhang 2025) typically explicate this gradability in terms of representational resolution or default reliability in the actual world. This paper argues, however, that these "Actualist" metrics are insufficient because they fail to distinguish between the intrinsic competence of an "Unlucky Expert" (whose skill is masked by environmental deviation) and the structural luck of a "Lucky Novice" (whose skill is mimicked by environmental assistance). Mere reliability allows luck to masquerade as competence. Drawing on Modal Virtue Epistemology, we propose that gradability is best measured not by actual success frequencies, but by the agent's Weighted Modal Volume across counterfactual worlds. We conclude that propositional knowledge contributes to skill precisely by functioning as a "Modal Extender" that insulates the agent against modal fragility, resulting in a view I term Modal Intellectualism.

KEYWORDS: gradable know-how, modal robustness, intellectualism, epistemic luck, safety radius

1. Introduction: The Puzzle of Degrees

When an agent S ϕ -s skillfully, we do not merely attribute to S the binary status of knowing how to ϕ ; we critically evaluate the degree of their knowledge. We assert that " S knows how to ϕ better than S' " or that " S is an expert while S' is merely competent." Unlike standard propositional knowledge (know-that), which is typically modeled as a binary state—one either knows that "snow is white" or one does not—knowledge-how appears to be intrinsically gradable (Ryle 2000; Stanley and Williamson 2017; Pavese 2017). The concept of gradability is essential for any adequate theory of skill, as it provides the metric to track the trajectory from novice to master (Dreyfus H 1986; Fridland 2014).

The question that thus arises is metaphysical: What constitutes the basis of these degrees? Answers to this question differ within the current Intellectualist landscape, yet they converge on a shared presupposition. Some scholars, such as Carlotta Pavese (2017), argue that gradability is a function of the representational

resolution of the agent's "practical modes of presentation." Others, such as Xiaoxing Zhang (2025), have recently refined this picture by introducing the concept of "default reliability," arguing that the degree of skill corresponds to the reliability of the agent's constitution in fulfilling a task. Furthermore, Garry Young (2024) has proposed that gradability supervenes on the fine-grainedness of motor representations.

However, despite these sophisticated advances, we argue that there remains a persistent "Normative Deficit" in how we measure skill. Most prevailing accounts rely on what we term Actualist Metrics. They subscribe to the Success-Entailment Thesis: the assumption that a higher degree of skill must logically correlate with a higher frequency of success in the actual (or default) environment.

Yet, this approach faces a significant puzzle regarding Epistemic Luck. As Sophie Kikkert (2022) notes, mere reliability is insufficient to distinguish a genuinely skilled agent from a "lucky novice" whose success is sustained by a benevolent environment (Environmental Mimicking). Conversely, an "unlucky expert" might fail due to deviant external causes—such as a sudden hurricane—without losing their intrinsic status as an expert (Environmental Masking) (Pavese 2018; Sosa 2007).

If we measure gradability strictly by actual outcomes or success-conducive representations, we commit a form of Epistemic Consequentialism: we confuse the instrumental output of the system (Agent + Environment) with the dispositional competence of the agent.

In this paper, we argue that we must abandon this purely result-oriented analysis and adopt a Modal Turn. We contend that the gradability of know-how is not a measure of actual success frequencies, but of Modal Robustness. Drawing on Modal Virtue Epistemology (Beddor and Pavese 2018), we propose that to know how to ϕ "better" is to possess a larger "Safety Radius"—defined as a Weighted Modal Volume across counterfactual worlds. An expert is defined not by their actual success rate, but by their capacity to maintain normative success in the face of modal perturbations (e.g. wind, stress, novel interference).

From this framework, we derive the thesis of Modal Intellectualism: the specific function of propositional knowledge (know-that) is not to drive routine execution, but to act as a Modal Extender. It expands the agent's Safety Radius from the "center" of habitual repetition into the "periphery" of novel variation.

The remainder of this paper proceeds as follows. In Sect. 2, we scrutinize the current Intellectualist consensus (Pavese, Zhang, Young), identifying their shared reliance on Actualist metrics. In Sect. 3, we introduce the "Problem of Environmental Luck" (Masking and Mimicking) to demonstrate the rupture

between performance and competence. In Sect. 4, we dismantle the "Probabilistic Defense," arguing that statistical frequency creates a "Swamping Error" that hides the brittleness of lucky agents. In Sect. 5, we formalize the geometry of skill, defining gradability as the Weighted Modal Volume of the agent's success. Finally, in Sect. 6, we conclude with the defense of Modal Intellectualism, showing how propositional knowledge insulates agency against modal fragility.

2. The Introduction of Gradable Know-How

The concept of know-how imposes a rigorous constraint on any adequate theory of intelligent action: the theory must account for the scalar nature of skill. As Gilbert Ryle famously observed, "We never speak of a person believing or opining significantly, but we do speak of him as knowing how to do something well or ill, efficiently or inefficiently" (Ryle 2000, 59). Unlike standard propositional knowledge, which is typically modeled as a binary state—one either knows that p is true or one does not—practical knowledge inherently admits of degrees. This phenomenon presents a structural challenge for Intellectualism, the view that defines know-how as a species of know-that (Stanley and Williamson 2017). If knowing how to ϕ is simply a matter of knowing a proposition (e.g., " ω is away to ϕ "), and truth is non-gradable, then how can one expert be "more skilled" than another? If both agents know the same true proposition, their epistemic standing should be identical. Consequently, to sustain the Intellectualist project, proponents must explain how a binary propositional attitude can ground a continuous gradient of expertise (Dreyfus H 1986; Fridland 2014).

The most sophisticated response to this challenge is offered by Carlotta Pavese (2017). Pavese argues that the locus of gradability is not the proposition itself—which remains constantly true—but the Practical Mode of Presentation (PMP) under which the proposition is grasped. Drawing on Fregean semantics, Pavese posits that to know how to ϕ is to stand in a knowledge relation to a proposition of the form " ω is away to ϕ " under a practical sense. Crucially, these practical senses function like maps or instructions: they can vary in quality. Two agents might know the same truth (that this is the way to ride a bicycle), but the expert grasps this truth under a PMP that is more prescriptive, detailed, and accurate. Pavese (2017) distinguishes between the sufficiency of a mode (having enough detail to perform the task) and the gradation of a mode. Superior skill supervenes on the "fineness of grain" of the practical representation. As she notes: "Gradability of know-how is explained in terms of the gradability of the accuracy and of the specificity of the practical senses involved" (Pavese 2017, 368). On this view, the difference between the novice and the expert is analogous to the difference between a rough sketch and

a high-resolution topographical map. Both represent the same territory (the task), but the expert's internal representation possesses a higher degree of representational richness. Thus, Pavese grounds the degree of skill in the intrinsic quality of the agent's mental state.

Recently, Xiaoxing Zhang (2025) has offered a significant refinement to the intellectualist picture. While acknowledging Pavese's insight that gradability can concern the quality of practical answers, Zhang argues that this representational metric is insufficient. Instead, he proposes that know-how is also gradable regarding the agent's "default reliability." Zhang contends that an adequate theory must account for the fact that we often grade skill based on the agent's capacity to fulfill a task given their "default constitution" (Zhang 2025, 2210). On this view, to know how to ϕ "better" is to possess a constitution that renders one reliably successful in performing ϕ under normal conditions. Zhang explicitly contrasts this with Pavese's focus on the content of knowledge, arguing that: "Know-how is gradable according to one's default reliability level. ... Referring to the agent's default profile, we are also free from the usual intellectualist dismissal of the necessity of actual reliable abilities" (Zhang 2025, 2213).

Crucially, while Zhang shifts the focus from the "method" to the "agent," the metric remains fundamentally probabilistic. By defining the degree of know-how in terms of a reliability level derived from the agent's default constitution, Zhang's account is still anchored in the frequency of success (albeit "default" success). Therefore, despite its sophistication, this view remains vulnerable to the Actualist trap: it measures the agent's propensity for success in a presumed "normal" world, but fails to track the robustness of that competence across deviant modal neighborhoods.

Complementing the semantic and epistemological accounts, Garry Young (2024) provides a mechanistic explanation for these degrees, grounding them in the sub-personal architecture of action. Young argues that the gradability of know-how supervenes on the fine-grainedness of motor representations. Young posits that the transition from novice to expert is not merely an accumulation of propositional facts, but a structural refinement of the motor schema—specifically within the agent's hierarchical generative models. The expert possesses a motor representation that minimizes the "error signal" between the intended outcome and the actual movement with greater precision than the novice. "The expert's motor representation of the action... allows for a greater degree of control and flexibility" (Young 2024, 5). This account aligns with Pavese's notion of Practical Modes of Presentation (PMPs) but locates the "grain" of the skill specifically in the kinematic details of the representation. For Young, "better" know-how corresponds to a tighter

coupling between intention and execution. Yet, like Pavese and Zhang, Young's metric is ultimately performance-oriented: the quality of the representation is judged by its ability to generate controlled, successful action in the agent's environment.

It is crucial to interpret these Intellectualist accounts charitably. Neither Pavese, Zhang, nor Young claims that skill is merely a static accumulation of facts (the crude "library model"). Instead, they offer a dynamic picture where degrees of skill correspond to the resolution of representations (Pavese, Young) or the default reliability of the agent (Zhang). Despite their differences regarding the locus of skill (mental representation vs. constitutional reliability), these views converge on a shared metric. They all presuppose that the normative measure of "knowing how well" is tethered to the production of successful outcomes in the actual (or "default") environment.

We can summarize this prevailing "Target View" as follows:

(T1) Gradability as Representational Resolution: S is more skilled than S' if S's practical mode of presentation (PMP) or motor schema maps the task with greater accuracy, detail, or fine-grainedness than S's (Pavese 2017; Young 2024).

(T2) Gradability as Resultant Reliability: S is more skilled than S' if S's default constitution is such that S produces successful instances of ϕ -ing with higher frequency or reliability in the relevant context (Zhang 2025).

The synthesis of (T1) and (T2) reveals a hidden premise: both metrics assume that the actual world (or the "default" context) is the sole arbiter of competence. A representation is deemed "detailed" only if it works; a constitution is "reliable" only if it yields success. The question that thus arises is rigorous: Do these metrics adequately capture the normative depth of expertise? Or do they conflate the possession of a disposition with its manifestation in favorable conditions? Specifically, does this Actualist Consensus have the resources to distinguish between an expert who fails due to deviant environmental interference and a novice who succeeds due to environmental benevolence? If the degree of know-how is indexed strictly to success-conduciveness in the actual world, the theory risks collapsing precisely when the world ceases to cooperate. In the sections that follow, we will argue that these metrics suffer from a "Normative Deficit." By prioritizing actual reliability, they fail to track the ontological stability of skill—a property that persists even when the link between competence and performance is severed by epistemic luck.

3. The Failure of Actualist Metrics: The Problem of Luck

The analysis in the previous section reveals a shared commitment among leading Intellectualists: the view that the degree of know-how supervenes on the production of success in the actual environment. Whether through "representational richness" (Pavese 2017) or "default reliability" (Zhang 2025), these theories subscribe to what we might call the Success-Entailment Thesis: *ceteris paribus*, a higher degree of skill logically implies a higher frequency of successful performance. However, this reliance on Actualist Metrics creates a significant normative deficit. By tethering competence to actual outcomes, these theories fail to insulate the judgment of skill from the noise of the environment. Specifically, they lack a mechanism to distinguish between success derived from agency and success derived from fortune.

To isolate this deficit, we must interrogate the theories with cases of Epistemic Luck. If the link between the agent's internal competence and the external outcome is severed—either by a hostile world or a benevolent one—Actualist metrics are forced into counter-intuitive evaluations. Consider the following thought experiment, which I term the Problem of Environmental Luck, designed to decouple internal competence from external reliability.

Case A: The Unlucky Expert (Environmental Masking)

Alice is a Grandmaster Archer. Her internal state is impeccable: she possesses the highest quality Practical Modes of Presentation (PMPs) and the most fine-grained motor schemata. However, she is situated in World A, an environment currently ravaged by a freak, chaotic hurricane. The wind vectors are random and physically overwhelming. Despite executing her technique perfectly, the deviant causal chain ensures that her arrows never reach the target.

- Actual Outcome: 0% Success Rate.
- Diagnosis: Her disposition to succeed is masked by environmental interference.

Case B: The Lucky Novice (Environmental Mimicking)

Bob is a complete novice. He possesses no relevant PMPs and his motor representations are coarse and unrefined. However, he is situated in World B, a "technologically benevolent" wind-tunnel equipped with an invisible laser-guidance system. This system detects his release and uses magnetic fields to auto-correct the arrow's trajectory mid-flight. Bob merely has to let go of the string. Consequently, his arrows hit the bullseye every time.

- Actual Outcome: 100% Success Rate.
- Diagnosis: His lack of skill is mimicked by environmental assistance.

The question that thus arises is normative: Who is more skilled? If we strictly apply the Intellectualist metrics established in Section 2—specifically (T2) which

tracks "reliability in performance"—we are forced into a paradox. Bob manifests perfect reliability ($P = 1.0$), while Alice manifests total failure ($P = 0.0$). If "knowing how to ϕ well" is simply " ϕ -ing well in the relevant context," the Actualist is committed to the absurd conclusion that the Lucky Novice is the superior agent.

If we adopt the Reliability/Frequency View—whether in its standard form or in Zhang's (2025) "default reliability" refinement—we are forced into a theoretic rupture. Bob (The Lucky Novice) manifests a reliability coefficient of 1.0 in his environment. Alice (The Unlucky Expert) manifests a coefficient of 0.0. If "knowing how to ϕ well" implies "reliability in performance" (T2), then the Actualist is committed to the conclusion that Bob is the superior knower. The Intellectualist might object that Bob's environment is "deviant" or that the laser system is an external aid. However, this objection fails for two reasons:

- (1) Indexicality of Reliability: Reliability is functionally defined relative to the environment the agent occupies. In World B, Bob is reliably successful. To claim he is "unreliable" requires appealing to a counterfactual world (e.g., "He would fail if the laser were off")—but that is precisely to abandon the Actualist metric and adopt the Modal Robustness view I advocate.
- (2) The Attributional Error: The core flaw of the reliability metric is that it measures the performance of the System (Agent + Environment), not the competence of the Agent. Pure reliability is blind to the source of success. It cannot distinguish between success attributable to competence and success attributable to environmental scaffolding. By crediting Bob with high skill, the Actualist theory commits a category mistake: it attributes a property of the "wind-tunnel system" to the agent's "default constitution."

If we pivot to Pavese's Representational View (T1), the situation appears more promising but ultimately reveals a deeper metaphysical contradiction. Pavese (2017) argues that skill corresponds to the quality of the Practical Mode of Presentation (PMP). She might contend that Alice remains an expert because she possesses a "high-resolution" PMP, even if it fails to execute. However, this defense runs aground on the Factivity Constraint. For Intellectualism, know-how is a species of knowledge, and knowledge is factive ($S \text{ knows } p \rightarrow p \text{ is true}$). To know how to ϕ is to know a true proposition of the form " ω is a way to ϕ ." In World A (The Hurricane), the proposition "Releasing the arrow at angle θ is a way to hit the target" is factually false. Due to the chaotic wind, that method is not a way to hit the target; it is a way to miss. This generates a lethal dilemma for the Representationalist, which I term *The Dilemma of Factivity*:

Horn 1 (The Externalist Horn): If the PMP's quality depends on its truth-conditions (i.e., actually mapping a way to success), then the hurricane degrades the truth-value of Alice's internal state. Her method is objectively false in World A.

Therefore, strictly speaking, she loses her knowledge. This leads to the counter-intuitive conclusion that environmental bad luck erases expertise.

Horn 2 (The Internalist Horn): To save Alice's expertise, the Representationalist must argue that her PMP remains "good" despite being false in the actual world. But this requires decoupling "skill" from "truth." It would imply that one can "know" how to ϕ even if one's belief about how to ϕ is radically disconnected from reality. This forces the Intellectualist to abandon the core tenet that know-how is a relation to a true proposition, collapsing know-how into mere belief-how or opining-how.

Thus, Pavese's account faces a deadlock: it must either deny Alice's expertise (Externalism) or deny the factivity of knowledge (Internalism). Neither option is acceptable for a robust theory of skill. The thought experiment exposes a fundamental rupture between Performance (the external event) and Competence (the internal disposition). Actualist gradability theories collapse this distinction by assuming that "doing well" is a strict proxy for "knowing how." However, this collapse leads to a failure in explanatory attribution. As the case of Bob (The Lucky Novice) demonstrates, successful performance can be entirely extrinsic—a product of a benevolent environment mimicking competence. Conversely, as the case of Alice (The Unlucky Expert) shows, a genuine failure in performance can be extrinsic—a product of a deviant environment masking competence. Therefore, current theories suffer from a Normative Deficit. By tethering the degree of skill to actual success rates, they fail to isolate the agent's contribution from the world's contribution.

The Case of Mimicking: In World B, the locus of success lies in the laser system, not in Bob. If we attribute the "skill" to Bob based on his 100% success rate, we commit a category mistake: we are crediting the agent with a property that belongs to the system (Agent + Environment).

The Case of Masking: In World A, the locus of failure lies in the hurricane, not in Alice. If we deny Alice's expertise based on her 0% success rate, we mistake an environmental impediment for an agential deficiency.

We perceive Alice as superior to Bob not merely because of counterfactual wishful thinking (i.e., "if only the wind stopped"), but because her competence is an intrinsic property that persists regardless of its manifestation. This mirrors the classic distinction in the metaphysics of dispositions between Masking and Mimicking (Johnston 1992). A fragile glass remains fragile even when carefully wrapped in bubble wrap (which masks its fragility from breaking). Similarly, a sturdy iron block does not become fragile just because a sound effect implies shattering when verified (which mimics fragility). In our analysis, Alice's skill is an intrinsic disposition masked by the hurricane. Bob's "skill" is an illusion mimicked by the laser. Actualist

metrics fail because they look solely at the manifestation (the shattered glass or the bullseye) rather than the dispositional structure. We cannot resolve this by observing actual outcomes more closely, for in the actual world, the outcomes deceive us. We must adopt a metric that transcends the actual: we need to evaluate the modal profile of the agent's ability—identifying which worlds verify their competence and which worlds merely mimic or mask it.

4. Gradability Revisited: The Probabilistic Defense

4.1 The Intellectualist Rejoinder

The critique presented in the previous section—that Actualist metrics crumble under the weight of Epistemic Luck—might invite a sophisticated rejoinder. Proponents of the Target View (such as Pavese or Zhang) might accuse the "Unlucky Expert" thought experiment of committing a Snapshot Fallacy. They could argue that expertise is not defined by the outcome of a single trial, nor by performance in a fleeting, freakish interval (the hurricane). Rather, skill is a longitudinal property. Just as a fair coin does not cease to be fair simply because it lands on "tails" five times in a row, an expert does not cease to be skilled simply because they inhabit a temporary pocket of bad luck. Although Pavese and Zhang do not explicitly address "Hurricane Alice," a robust defense is available to them: they can refine their definition of gradability by shifting from local actualities to long-run propensities. The logic underlying this response can be formalized as the *Probabilistic Defense (PD)*. The Intellectualist aims to insulate the theory from luck by appealing to the Law of Large Numbers and "normal" conditions.

Probabilistic Gradability (PD): Agent S is more skilled than Agent S' if and only if, over a sufficiently long timeline and under canonical conditions, the objective probability of S' succeeding ($P(S')$) is greater than the probability of S' succeeding ($P(S')$).

Under this refined metric, the Intellectualist attempts to resolve the dilemma:

- (1) Saving Alice (Integration over Time): They would argue that the hurricane in World A is a statistically deviant sample. If we integrate Alice's performance over a valid reference class (e.g., her entire career or multiple environments), her probability of success remains high. The hurricane is merely "noise" in the data.
- (2) Disqualifying Bob (Restricting Conditions): Regarding Bob, they would appeal to a *ceteris paribus* (all else being equal) clause. They might argue: "Bob's high probability is parasitic on the laser system. However, the 'canonical conditions' of archery exclude such technological crutches. If evaluated under standard conditions, Bob's probability drops to zero."

This defense appears powerful. It suggests that while luck dominates the short run (the snapshot), skill dominates the long run (the disposition). It attempts to preserve the link between skill and success by widening the temporal lens.

4.2 The Limits of Probability: Reliability vs. Robustness

However, the Probabilistic Defense rests on a seductive but fragile assumption. It assumes that Reliability (high frequency of success in a given environment) is synonymous with Competence. It implies that if an agent produces a stable track record of success, they must be skilled. We argue that this assumption is false because it allows Structural Luck to masquerade as competence.

Consider Bob (The Lucky Novice) again. The Intellectualist appeals to the "long run" to disqualify him. But let us run the simulation: suppose Bob lives his entire life inside the benevolent parameters of World B (Wind-Tunnel + Laser). Over 10,000 trials, his success rate remains strictly 100%. By any statistical definition of reliability, Bob is "reliable" relative to his environment. The Law of Large Numbers does not expose him, because the environment systematically biases the outcome in his favor. If we accept the Actualist/Probabilistic metric, we are forced to concede that Bob is a master archer relative to World B. Yet, our normative intuition revolts against this conclusion. Why? To pinpoint the deficit, we must draw on the distinction—highlighted by Sophie Kikkert (2022)—between Reliability and Robustness. Reliability concerns the frequency of success in the actual (or current) world. Robustness concerns the stability of success across counterfactual variations.

Bob's success is Hyper-Fragile. While his reliability is perfect ($P = 1.0$) at the specific modal coordinate of World B, it collapses to zero ($P = 0.0$) at the slightest perturbation. If the laser battery dies, or if a door opens creating a slight draft—events that represent nearby possible worlds—his capacity evaporates. Bob possesses what we might call "Cheap Reliability": a statistical track record sustained not by the agent's internal adaptability, but by the rigidity of a friendly world. The Probabilistic Defense fails because it measures the height of the success curve (how often one hits) without measuring the width of its base (how many world-variations one can withstand).

Here lies the core axiological failure of the Actualist account, which we term the *Swamping Error*. By privileging the high probability of success in the actual (or specified) context, the Intellectualist allows the metric of local reliability to swamp the diagnostic value of modal robustness.

The Swamping Logic: Under the probabilistic view, if S' succeeds 100% of the time in Context C, the theory treats the magnitude of this statistic as decisive. The

The Safety Radius of Skill: Modal Robustness and the Defense of Intellectualism

source of that success—whether it be the agent's internal motor representations or the machine's external guidance—is rendered secondary. The reliability conceals the modal fragility.

But genuine expertise is characterized precisely by its resistance to such fragility. The ontological standing of Alice (The Unlucky Expert) is defined not by her failure in the singular, deviant world she inhabits, but by her disposition across the relevant modal neighborhood. Alice is Robust: She would succeed if the wind dropped by 5 mph, if she moved 10 feet to the left, or if the humidity changed slightly. Her current failure is an anomaly surrounded by a dense cloud of counterfactual successes. Bob, conversely, is Brittle: He occupies a lonely island of success. In almost any nearby possible world (where the laser fails, or the tunnel opens), he fails.

Therefore, the Probabilistic Defense fails because it remains an Actualist metric at heart—it merely aggregates actualities over a timeline. It measures the performance of the System (Bob + Laser), not the competence of the Agent (Bob). To rigorously distinguish the Expert from the Lucky Novice, we cannot simply count the hits; we must ask the Safety Question: "How easily could they have failed?" For Bob, failure is just one dying battery away. For Alice, failure requires a hurricane. This realization forces us to abandon the flat geography of actual frequencies and turn our gaze to the modal profile of the agent. Skill is not about what does happen in the favorable here and now; it is about what could happen across the variation of logical space. This necessitates the Modal Turn we will develop in the next section.

5. From Reliability to Robustness: The Normative Turn

5.1 The Axiological Error

The Probabilistic Defense, scrutinized in the previous section, posits that if an agent manifests a high frequency of success in a specific environment, they are, by definition, highly skilled. The logic is seductive in its empiricism: if Bob hits the bullseye 100 times in a row (regardless of the laser-guided tunnel), the intuition is to credit him with high know-how. However, this defense rests upon a fragile axiological presupposition: it assumes that Resultant Reliability (the mere production of successful outcomes) is the fundamental measure of epistemic standing. It embraces a form of *Epistemic Consequentialism*, where the dispositions of an agent are valued solely for their instrumental ability to produce a desired world-state in the here and now. From this perspective, skill is reduced to a "black box" causal engine. The Consequentialist view asks only "Does it work?" rather than "Is it competent?" This approach suffers from two fatal defects:

Causal Alienation: It alienates the agent from the nature of their success. It forces us to conclude that there is no normative difference between a "Lucky Novice" (whose success is sustained by a deviant environmental crutch) and a "True Expert" (whose success is sustained by intrinsic agency).

System-Level Confusion: If we accept this view, we lose the ability to define skill as a property of the person. It becomes merely a property of the person-environment system.

This intuitive account must be contrasted with a structural, modal-theoretic account. As Sophie Kikkert (2022) argues, the evaluation of ability requires more than just output statistics; it requires analyzing the robustness of the performance. Real expertise is superior to mere reliability because reliability can be "cheap"—it can result from a fragile, serendipitous alignment of environmental parameters. Expertise, by contrast, implies an invariant capacity to succeed across variations. The defining question, therefore, is not "How often does S succeed in the actual world ($W@$)?" The defining question is: "How easily would S fail if $W@$ were slightly different?" To rigorize this distinction, we must move beyond simple probability and employ the structural framework of Modal Virtue Epistemology.

5.2 The Modal Framework

To formalize the shift from probability to robustness, we can leverage the rich resources of Modal Virtue Epistemology. This framework, successfully adapted to the philosophy of skill by Bob Beddor and Carlotta Pavese (2018), provides the structural machinery necessary to exclude epistemic luck. In standard "Anti-Luck Epistemology" (Sosa 2007; Carter and Pritchard 2015; Stanley and Williamson 2001), a belief constitutes knowledge not merely because it happens to be true (which could be a result of lucky guessing), but because it is safe. A belief is technically defined as safe if and only if, in all "close" or "nearby" possible worlds where the agent forms the belief via the same method, the belief remains true.

Safety is the modal antidote to luck. Consider the classic "Stopped Clock" case: If you form the belief "It is 2:00 PM" by looking at a stopped clock that happens to display 2:00 PM at exactly the right moment, your belief is true, but it is unsafe. Why? Because there is a very close possible world—say, one where you looked at the clock at 2:01 PM—where you would have formed a false belief. The truth of your belief is fragile; it does not survive counterfactual perturbation. Beddor and Pavese (2018) brilliantly transpose this epistemic structure onto the domain of practical action. They argue that just as knowledge is "safe belief," skill is "safe execution." They propose that what distinguishes a genuine exercise of know-how

The Safety Radius of Skill: Modal Robustness and the Defense of Intellectualism

from a lucky success (like a novice hitting a bullseye due to a sneeze) is the modal profile of the performance.

Skill as Modal Safety: An agent S knows how to ϕ if and only if S's performance of ϕ is safe; that is, in all nearby possible worlds where S attempts to ϕ (under similar conditions and using the same method), S' succeeds in ϕ -ing (Beddor and Pavese 2018, 66).

This definition serves a dual function. First, it solves the "*Practical Gettier Problem*": Bob (The Lucky Novice) succeeds in the actual world $W@$ (via the laser), but his success is unsafe because in the nearest possible world $W_{@+1}$ (where the laser battery dies), he fails. His success is modally fragile. Second, and most importantly for our inquiry, it implies that skill is inherently *trans-world*. It is not a property of a single event, but a property of the agent's capacity to track success across a range of scenarios.

While Beddor and Pavese primarily utilize this framework to adjudicate the binary question (Knowledge vs. Non-Knowledge), I argue that this modal architecture contains the missing key to solving the Puzzle of Gradability. If the presence of skill is determined by safety in nearby worlds, then the degree of skill must be determined by the magnitude of that safety. The "degree" of know-how is not about the frequency of success in one static world; it is about the volume of the modal space in which the agent retains control.

5.3 Defining the "Safety Radius"

To categorize degrees of know-how without relying on actual outcomes, we propose we conceptualize skill as a Safety Radius within modal space. This approach draws inspiration from the logic of counterfactuals (Lewis 1973) and the metaphysics of graded dispositions (Manley and Wasserman 2008). Let us visualize the modal landscape as a coordinate system. Let $W@$ (the Actual World) be the center point of a modal sphere. Surrounding this center are concentric circles of Possible Worlds ($\omega_1, \omega_2 \dots \omega_n$), ordered by their similarity distance from the center.

Proximal Worlds (Safety Zone): Worlds very similar to $W@$ (e.g., standard conditions where the wind is 1mph stronger, or the lighting is slightly dimmer).

Distal Worlds (Danger Zone): Worlds that differ significantly (e.g., a hurricane strikes, the equipment breaks, or the laws of physics fluctuate).

Under the *Actualist/Probabilistic metric*, we strictly measure the density of success at the singular point $W@$. Under the *Modal Robustness metric*, we measure the *extension*—how far out the agent can travel from the center while maintaining normative success. We must reject the simplistic "Superset Relation" (which implies

S is better than S' only if S succeeds in every world S' does), as this fails in cases of incommensurable skills. Instead, we should understand gradability in terms of *Weighted Modal Volume*.

Definition of Gradable Skill (M1: The Modal Volume Thesis):

(M1) Agent S is more skilled than Agent S' at ϕ -ing if and only if the weighted measure of the set of relevant possible worlds in which S succeeds ($V(W_S)$) is greater than that of S' ($V(W_{S'})$).

$$V(W_S) > V(W_{S'})$$

Where V is a volume function that weights worlds by their *normative relevance* and their *distance* from the center of the relevant context.

This definition accommodates the intuition that expertise is an expansive property. While novices are confined to the "bullseye" of the modal space (ideal, sanitized conditions), experts inhabit the expanding concentric circles of difficult variations. This allows us to rigorously distinguish the agents from our previous thought experiment without referencing their misleading actual outcomes.

(1) The Fragile Novice (Bob): Modal Brittleness

Bob has a Safety Radius of Zero. He succeeds in $W@$ (Wind Tunnel + Laser). However, in the very closest possible world- $W_{@+1}$ (where the laser battery dies) or $W_{@+2}$ (where a door opens)-his performance collapses. In the language of Manley & Wasserman (2008), Bob creates a "Specific Disposition": he occupies a singular "island of success" surrounded by a vast sea of failure. His success is statistically high but modally brittle.

(2) The Robust Expert (Alice): Modal Volumetry

Alice possesses a Wide Safety Radius. Although she currently inhabits a Distant World where success is physically impossible (The Hurricane World, W_{dist}), her competence covers the vast majority of the relevant modal space. She succeeds in W_{norm} (Normal), W_{windy} (Windy), W_{noisy} (Noisy), and $W_{injured}$ (Minor Injury). The integral of her success over the modal space is massive. Thus, under (M1), $V(W_{Alice}) \gg V(W_{Bob})$. Alice is the expert not because of what she does, but because of the volume of counterfactual space she has conquered.

5.4 Reinterpreting the Dilemma

With the Modal Volume (M1) framework established, we can now rigorously resolve the paradox of Alice (The Unlucky Expert) and Bob (The Lucky Novice). Why is Alice (0% actual success) more skilled than Bob (100% actual success)? The answer lies in the asymmetric Modal Geography of their respective dispositions.

The Safety Radius of Skill: Modal Robustness and the Defense of Intellectualism

(1) Analysis of Bob: The Rigid Specialist

Bob succeeds in W_{laser} (The Laser World). However, his success relies on a specific external mechanism. If we move even slightly in modal space to $W_{\text{laser-fail}}$ (where the battery dies) or W_{standard} (ordinary archery), his performance collapses. Modal Diagnosis: Bob's set of successful worlds, W_{Bob} , is extremely narrow. It constitutes a fragile singleton or a tiny cluster. His safety radius is effectively zero. The fact that he happens to inhabit the one world where he succeeds is a matter of Environmental Good Luck.

(2) Analysis of Alice: The Displaced Generalist

Alice is situated in $W_{\text{hurricane}}$, a Distant World far from the normative center of archery practice. She fails here. However this failure is not a reflection of her modal volume. Her competence set, W_{Alice} , encompasses W_{normal} , W_{breezy} , W_{rainy} , and countless other variations. Modal Diagnosis: Alice's failure is a case of Modal Displacement. She possesses a massive Safety Radius, but she is currently located outside of it due to Environmental Bad Luck. Comparative Logic: Even though Alice fails in the actual world $W_{@}$ (which happens to be $W_{\text{hurricane}}$), the weighted measure of her successful worlds is vastly superior to Bob's: $V(W_{\text{Alice}}) \gg V(W_{\text{Bob}})$. She is the expert because her dispositional structure is strictly more robust.¹

We have established that High-Degree Skill = Modal Robustness. The final step in this inquiry is to identify the cognitive mechanism that enables this robustness. This brings us to the debate with Anti-Intellectualism. The Anti-Intellectualist (e.g., Rousset and Dreyfus 2021) rightly argues that in the "center" of the modal space (standard, repetitive conditions), propositional reflection is often unnecessary and even inhibitory. "Muscle memory" or non-conceptual procedural dispositions are sufficient for success in $W_{@}$ when $W_{@}$ is stable. The body adapts through *Local Plasticity*.

However, the defining characteristic of expertise is not just coping with the familiar, but navigating the Distant Periphery of the modal sphere—worlds where standard rules break down. I argue that Propositional Knowledge (Know-That) is the specific functional architecture required to expand the Safety Radius into these distal regions.

Consider two agents facing a novel perturbation (e.g., a sudden, unprecedented crosswind, $W_{\text{crosswind}}$):

¹ On Intrinsicity: This aligns with the metaphysical intuition that skill is an intrinsic property of the agent. Alice's modal volume is determined by her internal constitution (her knowledge and muscles), whereas Bob's success is determined by the extrinsic setup. The Modal Volume metric effectively tracks intrinsicity by testing the agent across environmental variations.

Agent X (Pure Habit): Relies on non-conceptual motor schemata. These schemata are learned associations between specific stimuli and responses. While they exhibit local plasticity (micro-adjustments), they lack *compositionality*. Agent X cannot recombine separate concepts to generate a new solution. If the wind vector falls outside their training data, the habit rigidly fails.

Agent Y (Modal Intellectualist): Possesses motor habits plus propositional knowledge of ballistics (e.g., "Wind force F requires offset angle θ "). Propositional representations satisfy the *Generality Constraint* (Evans et al. 1983). Because propositions are compositional, Agent Y can logically derive a novel solution for a situation never previously encountered. She can "decouple" from the immediate affordance, run an offline simulation using her theoretical knowledge, and intervene on her own motor output.

Modal Intellectualism: The primary function of know-that in skilled action is to act as a Modal Extender. While non-propositional habits secure reliability in the center of the modal space (Proximal Worlds), propositional knowledge secures safety in the periphery (Distant Worlds).

Thus, the "gradability" of know-how correlates with the depth of the agent's propositional understanding. The Master knows "better" than the Apprentice not because they think more during routine performance (which would be inefficient), but because their intellectual understanding allows them to solve problems in counterfactual worlds where the Apprentice's rigid habits would shatter.² The "intellect" in Intellectualism is not an interference mechanism; it is the insulation against modal fragility.

In summary, the transition from Reliability to Robustness fundamentally alters the metaphysics and epistemology of skill. By abandoning the Actualist framework, we enact three decisive shifts:

- (1) **Axiological Shift (The Metric):** We reject Epistemic Consequentialism, which values skill solely for its actual output. We replace the linear metric of Resultant Reliability (frequency of success in $W@$) with the structural metric of Weighted Modal Volume (the magnitude of the agent's Safety Radius across counterfactual worlds).
- (2) **Ontological Shift (The Locus):** We stop analyzing skill as a property of the System (Agent + Environment) and re-establish it as an Intrinsic Disposition of the Agent. This allows us to distinguish genuine competence from environmental mimicking, and genuine expertise from environmental

² Heideggerian Breakdown: This view is compatible with the phenomenological observation that experts act "mindlessly" during smooth coping. The propositional knowledge is latent. It is activated primarily during moments of "breakdown" (Heidegger, 1927) or novel interference, which correspond precisely to the "Distant Worlds" in our modal geometry.

The Safety Radius of Skill: Modal Robustness and the Defense of Intellectualism

masking.

- (3) Architectural Shift (The Role of Mind): We redefine the function of propositional knowledge (K now -That). It is not a redundant driver of "smooth coping" in the center of modal space, but a Modal Extender that provides Combinatorial Flexibility. It secures safety in the Distant Periphery, distinguishing the adaptable expert from the rigid creature of habit.

By adopting this normative stance, we can explain why Alice remains an expert in the hurricane: her Modal Volume remains massive, even if her manifestation is currently masked by the cruelty of her modal location. Skill, ultimately, is the conquest of modal space.

6. Concluding Remarks

At the outset of this inquiry, we confronted the Puzzle of Degrees: On what metaphysical basis do we assert that one agent is "more skilled" than another? The prevailing answer in the current literature—offered by Representationalists like Pavese(2020) and Young (2024), and refined by Contextualists like Zhang (2025)—has been resolutely Actualist. Whether implicitly or explicitly, these theories posit that the degree of know-how supervenes on the production of success in the actual (or default) environment. They embrace what we have termed Epistemic Consequentialism, conflating the instrumental output of the system with the epistemic standing of the agent.

However, through the analysis of Environmental Masking (The Unlucky Expert) and Environmental Mimicking (The Lucky Novice), we have demonstrated that Actualist metrics are normatively deficient. By tethering skill to actual reliability, they fail to insulate competence from the distorting noise of epistemic luck. A metric that ranks a lucky novice in a laser-guided wind-tunnel above a grandmaster in a hurricane commits a category mistake: it confuses the stability of the environment with the disposition of the agent.

To resolve the rupture between performance and competence, we must abandon the flat geography of the actual world and adopt the framework of Modal Intellectualism. This view posits that skill is not a measure of what happens here and now, but of what could happen across a weighted spectrum of counterfactual scenarios. Drawing on the logic of Modal Virtue Epistemology (Beddor and Pavese 2018) and the distinction between reliability and robustness (Kikkert 2022), we finalize the definition of gradable know-how as follows:

- (M1) The Modal Gradability Thesis:* Agent S is more skilled than Agent S' at ϕ -ing if and only if the *Weighted Modal Volume* of the possible worlds in which S' succeeds ($V(Ws)$) is greater than that of S'($V(Ws')$).

Under this definition, the degree of skill corresponds to the magnitude of the agent's *Safety Radius*. The Novice relies on rigid, non-conceptual habits that secure success only in a narrow band of ideal worlds (High Reliability, Low Robustness). The Expert possesses a competence that extends into deviant, difficult, and structurally novel worlds (High Robustness). Crucially, this formulation reveals the indispensable function of Propositional Knowledge (Know -That). Contra the Anti-Intellectualist claim that thinking interferes with skilled action, we have argued that propositional knowledge functions as a Modal Extender. While muscle memory ensures "local plasticity" in the center of the modal space, the generality and combinatorial flexibility of propositional knowledge allow the agent to diagnose errors and restructure action plans in the "Distant Periphery." "Thinking" is the mechanism that insulates agency against modal fragility.

Ultimately, the difference between the Master and the Apprentice is not that the Master succeeds more often in easy conditions, nor that they possess a "higher resolution" map of a static world. It is that the Master is ontologically robust. Their agency survives the counterfactual perturbations that would destroy the novice. Skill, therefore, is nothing less than the conquest of modal space.

Declaration

Conflict of Interest: No conflict of interest to declare.

References

- Beddor, Bob, and Carlotta Pavese. 2018. "Modal Virtue Epistemology." *Philosophy and Phenomenological Research* 101 (1):61-79.
- Carter, J. Adam, and Duncan Pritchard. 2015. "Knowledge-How and Cognitive Achievement." *Philosophy and Phenomenological Research* 91 (1):181-99.
- Dreyfus H, Dreyfus S E. 1986. *Mind over Machine: The Power of Human Intuition and Expertise in the Era of the Computer*. Free Press.
- Evans, J St BT, Julie L Barston, and Paul Pollard. 1983. "On the Conflict between Logic and Belief in Syllogistic Reasoning." *Memory & cognition* 11 (3):295-306.
- Fridland, Ellen. 2014. "They've Lost Control: Reflections on Skill." *Synthese* 191 (12):2729-50.
- Johnston, Mark. 1992. "How to Speak of the Colors." *Philosophical Studies: An International Journal for Philosophy in the Analytic Tradition* 68 (3):221-63.
- Kikkert, Sophie. 2022. "Ability's Two Dimensions of Robustness." *Proceedings of the Aristotelian Society* 122 (3):348-57.

- Lewis, David. 1973. "Counterfactuals and Comparative Possibility." In *Ifs: Conditionals, Belief, Decision, Chance and Time*. Springer.
- Manley, David, and Ryan Wasserman. 2008. "On Linking Dispositions and Conditionals." *Mind* 117 (465):59-84.
- Pavese, Carlotta. 2017. "Know-How and Gradability." *Philosophical Review* 126 (3):345-83.
- Pavese, Carlotta. 2018. "Know-How, Action, and Luck." *Synthese* 198 (Suppl 7):1595-617.
- Pavese, Carlotta. 2020. "Practical Representation." In *The Routledge Handbook of Philosophy of Skill and Expertise*, Edited by Fridland, Ellen, and Carlotta Pavese. Routledge.
- Rousse, B. Scot, and Stuart E. Dreyfus. 2021. "Revisiting the Six Stages of Skill Acquisition." In *Teaching and Learning for Adult Skill Acquisition: Applying the Dreyfus & Dreyfus Model in Different Fields*, Edited by Rousse, B. Scot, and Stuart E. Dreyfus.
- Ryle, Gilbert. 2000. *The Concept of Mind*. University Of Chicago Press.
- Sosa, Ernest. 2007. *A Virtue Epistemology*. Oxford University Press.
- Stanley, Jason, and Timothy Williamson. 2001. "Knowing How." *The Journal of Philosophy* 98 (8):411-44.
- Stanley, Jason, and Timothy Williamson. 2017. "Skill." *Noûs* 51 (4):713-26.
- Young, Garry. 2024. "On the Gradability of Knowledge How, and Its Relationship to Motor Representations and Ability." *Synthese* 203 (5):139. <https://doi.org/10.1007/s11229-024-04543-w>.
- Zhang, Xiaoxing. 2025. "Gradable Know-How." *Inquiry: An Interdisciplinary Journal of Philosophy* 68 (7):2197-213.