Scientific Research Publishing

Beijing Law Review, 2026, 17(4), 1285-1316
https://www.scirp.org/journal/blr
ISSN Online: 2159-4635
ISSN Print: 2159-4627

The Dual-Track Evolution of Circumstantial Evidence at the International Court of Justice: Expanding Access, Rising Thresholds

Peiran Lu

School of Law, Beijing Foreign Studies University, Beijing, China
Email: lupr2003@163.com

How to cite this paper: Lu, P. (2026). The Dual-Track Evolution of Circumstantial Evidence at the International Court of Justice: Expanding Access, Rising Thresholds. *Beijing Law Review, 17,* 1285-1316.
https://doi.org/10.4236/blr.2026.174062

Received: June 11, 2026
Accepted: November 24, 2026
Published: November 27, 2026

Copyright © 2026 by author(s) and Scientific Research Publishing Inc.
This work is licensed under the Creative Commons Attribution-NonCommercial International License (CC BY-NC 4.0).
http://creativecommons.org/licenses/by-nc/4.0/

Open Access

Abstract

The International Court of Justice has, since the *Corfu Channel* case of 1949, recognised that a State victim of an internationally wrongful act may have "a more liberal recourse to inferences of fact and circumstantial evidence" where the respondent State enjoys exclusive territorial control over the relevant facts. This rule is conventionally read as a single doctrine that lowers the evidentiary burden on applicants confronting an opponent's monopoly over proof. This article challenges that reading and the assumption of a single doctrinal lineage. It first disentangles three distinct lines of jurisprudence that are commonly conflated: the Court's own standard of proof for charges of exceptional gravity, which descends from *Corfu Channel*; the "only reasonable inference" test, which entered international law from the jurisprudence of the ad hoc criminal tribunals as the inferential counterpart of proof beyond reasonable doubt; and the plausibility requirement governing provisional measures, which is functionally separate from both. On the basis of this corrected genealogy, the article argues that the circumstantial-evidence rule has evolved along two analytically distinct axes: a horizontal axis, along which the *scope* of inferential proof has widened from physical facts to State policy and subjective intent, and a vertical axis, along which the *threshold* of probative force has risen. The central claim is that these movements are functionally connected—and that the connection is visible in the Court's response to the expansion. Confronted with the newly accessible domain of subjective intent, the Court did not elaborate its own "single conclusion" formula but imported a threshold calibrated to criminal conviction. That borrowing is the compensating constraint. The framework offers a fresh vantage point on the debate over whether the standard for inferring genocidal intent has become unattainably demanding, and reframes it as a question about the fit between a threshold and the institutional setting into which it has been transposed.

Keywords

International Court of Justice, Circumstantial Evidence, Inference of Fact, Standard of Proof, Genocidal Intent

1. Introduction

The rule on circumstantial evidence articulated by the International Court of Justice (ICJ, or "the Court") in the *Corfu Channel* case occupies a peculiar position in the law of international adjudication. It is, on the one hand, routinely invoked as a doctrine of accommodation: a recognition that an applicant State which cannot penetrate the evidentiary monopoly of its opponent should not, for that reason alone, be denied a remedy[1]. The Court there held that exclusive territorial control may impede direct proof of facts giving rise to responsibility, and that the victim State "should be allowed a more liberal recourse to inferences of fact and circumstantial evidence"[2]. On this reading the rule is protective of applicants, a procedural counterweight to the structural asymmetry of inter-State litigation.

On the other hand, the same body of jurisprudence has more recently attracted sustained criticism for setting the bar of proof so high as to render certain claims practically incapable of being established. The requirement, crystallised in *Croatia v. Serbia*, that genocidal intent be "the only inference that could reasonably be drawn" from a pattern of conduct has been described by scholars and even by intervening States as threatening to make findings of genocide "near-impossible"[3]. On this reading the rule is, if anything, hostile to applicants.

How can a single doctrine be at once a tool of accommodation and an instrument of near-prohibitive stringency? The conventional literature has tended to address the rule's protective and restrictive faces separately, or to treat the latter as a recent and regrettable departure from the former. This article advances a different account. It argues that the two faces are not contradictory but are the visible manifestations of two analytically distinct movements within the jurisprudence, movements that have unfolded in tandem and, more importantly, are functionally linked.

The argument proceeds from a deceptively simple analytical move: the disaggregation of the *Corfu Channel* rule into two separable dimensions. The first

[1] *Corfu Channel* (*United Kingdom v Albania*) (*Merits*) [1949] ICJ Rep 4, 18.

[2] ibid.

[3] See *Application of the Convention on the Prevention and Punishment of the Crime of Genocide* (*Croatia v Serbia*) (*Judgment*) [2015] ICJ Rep 3, para 148; and Joint Declaration of Intervention of Canada, the Kingdom of Denmark, the French Republic, the Federal Republic of Germany, the Kingdom of the Netherlands, and the United Kingdom of Great Britain and Northern Ireland, filed pursuant to Article 63 of the Statute in *Application of the Convention on the Prevention and Punishment of the Crime of Genocide* (*The Gambia v Myanmar*) (15 November 2023), paras 51-52, available at the Court's website (https://icj-cij.org/node/203299), cautioning against an interpretation of the standard for inferring genocidal intent so demanding as to make findings of genocide near-impossible. The substance of that concern is discussed in Part 4 below.

concerns the *conditions of access* to inferential proof—the circumstances in which an applicant may resort to circumstantial evidence at all. The second concerns the *threshold of probative force*—how compelling that circumstantial evidence must be before the Court will act upon it. Once these dimensions are held apart, a pattern emerges across seven decades of jurisprudence. Along what this article terms the *horizontal axis*, the conditions of access have steadily expanded: from the physical facts of mine-laying within territorial waters in *Corfu Channel*, through the attribution of clandestine support to armed groups in *Nicaragua*, to the proof of State policy and the subjective intent of officials in the *Genocide Convention* cases. Along the *vertical axis*, the threshold of probative force has correspondingly risen: from facts "leading logically to a single conclusion", through the demand for "fully conclusive" evidence in *Bosnia v. Serbia*, to the "only reasonable inference" formulation of *Croatia v. Serbia*.

The central thesis of this article is that these two axes are not merely co-present but functionally connected. The expansion of access along the horizontal axis—the application of inferential proof to ever more consequential and politically charged matters, culminating in the ascription of genocidal intent to a State—has generated a corresponding pressure to raise the threshold along the vertical axis. The elevation of the threshold operates, in effect, as a compensating constraint: as the domain in which circumstantial evidence may be deployed widens, the Court has demanded more of that evidence in order to guard against the over-extension of inferential reasoning into matters where the stakes of error are gravest. This is what the article calls the *dual-track evolution* of the rule.

Establishing that thesis requires, however, a preliminary corrective that occupies a substantial part of the analysis. The standards at issue are commonly discussed as though they formed a single doctrinal sequence descending from *Corfu Channel*. They do not. Three distinct lines of jurisprudence, with different origins and different functions, converge in the modern doctrine: the Court's own standard of proof for charges of exceptional gravity, which genuinely descends from *Corfu Channel*; the "only reasonable inference" test, which entered international law from the jurisprudence of the ad hoc criminal tribunals, where it developed as the inferential counterpart of proof beyond reasonable doubt; and the plausibility requirement governing provisional measures, which is functionally separate from both. Part 4.1 disentangles these lines. The correction matters because it changes what has to be explained. The vertical elevation is not the internal ripening of a 1949 formula; it is, in its decisive component, an *importation*. And that importation is itself the strongest evidence for the functional claim: confronted with the newly accessible domain of subjective intent, the Court did not elaborate the formula it had used for half a century but reached for a more demanding test from a neighbouring jurisprudence—a choice intelligible as a response to the pressure that the expansion of access had created.

This reframing carries a payoff for a live and unresolved debate. The question whether the "only reasonable inference" standard is unattainably demanding has

divided commentators, with some maintaining that the standard, properly understood, is met where the pattern of conduct cannot be explained other than by genocidal intent, and others warning that an unduly strict application could make findings of genocide practically unattainable[4]. This article suggests that part of the reason the debate has reached an impasse is that it has treated the vertical axis in isolation—asking only whether the threshold is too high in the abstract—while neglecting the horizontal expansion that conditions and, on the account offered here, partially explains it. Seen against the widening scope of the rule, the elevation of the threshold appears less as an arbitrary act of judicial caution than as a structurally intelligible response to the rule's own success.

A clarification of method, case selection, and scope is in order at the outset. The analysis is doctrinal: it proceeds through a close reading of the Court's own reasoning in a defined line of authority. The cases on which it principally relies—*Corfu Channel*, *Nicaragua*, *Bosnia v. Serbia*, and *Croatia v. Serbia*—are chosen because they constitute the backbone of the Court's articulation of the circumstantial-evidence rule: *Corfu Channel* is its origin and remains the invariable point of departure; *Nicaragua* is the principal instance of the rule's application beyond the physical facts that occasioned it; and the two *Genocide Convention* cases are the decisions in which the Court has most fully elaborated both the reach of inferential proof and the standard governing it. These are, in other words, the authorities in which the two dimensions this article isolates are most clearly visible and most consequentially in tension. Other decisions bearing on proof and evidence before the Court—for instance *Oil Platforms*, *Pulp Mills*, or *Armed Activities on the Territory of the Congo*—are not treated as separate anchors, not because they are irrelevant, but because they refine matters (the treatment of expert evidence, the weight of documentary and testimonial material) that lie alongside rather than at the core of the specific rule on inference from circumstantial evidence with which this article is concerned; they are noted where they bear on the argument. The provisional-measures phase of *South Africa v. Israel* is included not as authority on the merits but precisely for what it reveals about the calibration of the evidentiary standard across procedural phases, as Part 4 explains.

This article is, further, a study in the law of evidence; it is not a commentary on the merits of any pending or decided dispute. The *Genocide Convention* cases, including the provisional-measures phase of *South Africa v. Israel* and the January 2026 oral proceedings in *The Gambia v. Myanmar*, are examined solely as jurisprudential material bearing on the evidentiary standard, and no position is taken, or should be inferred, on the underlying factual or political questions those cases raise. The analysis seeks to identify a pattern in the Court's reasoning and to offer an explanatory account of it, not to adjudicate the controversies in which that reasoning has been deployed.

[4]The contours of this debate are set out in Part 4. For the interpretation of reasonable alternatives, see Barigye et al. (2025). For the concern about an unduly strict application, see the Joint Declaration of Intervention (n 3) para 51.

The article proceeds in six further parts. Part 2 returns to *Corfu Channel* and disaggregates the rule into its access and threshold dimensions, distinguishing both from the classical categories of burden and standard of proof. Part 3 traces the horizontal expansion of access across the jurisprudence. Part 4 turns to the threshold: it opens, in Part 4.1, by disentangling the three lines of jurisprudence described above and establishing the true genealogy of each standard, before tracing the vertical elevation itself and situating the recent practice within it. Part 5 develops the article's central claim concerning the functional link between the two axes, draws out its consequences for the contemporary debate, and takes up the normative question that the corrected genealogy raises—whether a threshold calibrated to criminal conviction is appropriately transposed to proceedings concerning State responsibility. Part 6 considers the implications for the practice of international dispute settlement more broadly. Part 7 concludes.

2. The Foundational Rule: *Corfu Channel* and Its Internal Structure

The rule on circumstantial evidence was born of a genuine evidentiary predicament. The *Corfu Channel* case, the first contentious matter decided by the Court, arose from the explosion of mines in Albanian territorial waters on 22 October 1946, which struck two British destroyers passing through the Corfu Strait and caused the deaths of forty-four crew members[5]. The United Kingdom contended that Albania was internationally responsible for the consequences of the mine-laying; Albania denied any knowledge of it. The dispute thus turned on a question of proof that Albania was uniquely positioned to frustrate: could the United Kingdom establish that Albania knew of the laying of mines within its own territorial waters, when all the relevant facts—the identity of those who laid the mines, the operations of the Albanian coastal defences, the records of any decision to acquiesce—lay within Albania's exclusive control[6]?

The Court's response to this predicament is the foundation of the entire doctrine. It reasoned that the exclusive control exercised by a State within its frontiers can make it difficult to furnish direct proof of facts giving rise to responsibility, and that the injured State should therefore have "a more liberal recourse to inferences of fact and circumstantial evidence"[7]. The Court added that such indirect evidence "is admitted in all systems of law, and its use is recognized by international decisions", and that it "must be regarded as of special weight when it is based on a series of facts linked together and leading logically to a single conclusion"[8]. Applying this approach, the Court drew upon the close surveillance of the strait by the Albanian coastal authorities, the geographical configuration of the area, and the feasibility of observing the mine-laying operation from the shore, in

[5] *Corfu Channel* (n 1) 10, 12-13.
[6] ibid 18.
[7] ibid.
[8] ibid.

order to conclude that the laying of the minefield could not have escaped the knowledge of the Albanian Government[9].

For the purposes of the present analysis, what matters is not only the result but the internal architecture of the rule the Court announced. That architecture can be disaggregated into two analytically distinct components, the separation of which is the methodological foundation for everything that follows.

The first component may be termed the *access dimension*. It specifies the condition under which the more liberal recourse to inference becomes available: namely, that the respondent enjoys "exclusive control" over the relevant facts, such that direct proof is effectively foreclosed to the applicant. This is a threshold of *entitlement*—it identifies the circumstances that trigger the rule. It says nothing, in itself, about how much evidence is ultimately required; it speaks only to whether the applicant may have resort to circumstantial evidence at all. In *Corfu Channel* the access condition was satisfied by the physical location of the relevant facts within Albanian territorial waters and under Albanian governmental control.

The second component may be termed the *threshold dimension*. It specifies how compelling the circumstantial evidence must be before the Court will treat the inferred fact as established: the evidence is "of special weight" when it rests on "a series of facts linked together and leading logically to a single conclusion". This is a standard of *probative sufficiency*—it governs the evaluation of the evidence once resort to it has been permitted. The phrase "leading logically to a single conclusion" is the operative formula, and its precise content will become decisive in Part 4.

Because the argument that follows turns on fine distinctions between related evidentiary concepts, it is worth pausing to fix terminology. Four notions must be kept apart. The *burden of proof* answers the question *who must prove a given fact*: in principle, the party that asserts a fact bears the burden of establishing it (*actori incumbit probatio*). The *standard of proof* answers the question *how convinced the Court must be* before it treats a fact as established—whether, for instance, the evidence must be "fully conclusive" or merely persuasive on the balance of probabilities. These two are the classical categories of the law of evidence, and they are logically independent of one another: a rule may allocate the burden to one party while setting the standard at any level of stringency. The two dimensions this article isolates are not substitutes for these classical categories but analytical refinements that cut across them, specific to the operation of *circumstantial* proof. The *access dimension* concerns a prior question that neither burden nor standard directly addresses: *in what circumstances may a party discharge its burden by inference at all*, rather than by direct evidence—a question that becomes pressing precisely where direct proof is foreclosed by an opponent's exclusive control. The *threshold dimension* then concerns how the standard of proof operates once in-

[9]ibid 18-22.

ference is permitted: *how strong the inferential link must be*—how completely the circumstantial evidence must exclude competing explanations—before the applicable standard is met. In the vocabulary of the classical categories, the access dimension conditions the *mode* by which the burden may be discharged, and the threshold dimension specifies what the *standard* demands of inferential, as opposed to direct, evidence. The burden of proof, by contrast, remains throughout on the party asserting the fact; nothing in this article's account shifts it. Keeping these four notions distinct is essential to what follows, because the central claim—that access has widened while the threshold has risen—is precisely a claim that two of them have moved in different directions while the burden has stayed where it was.

The distinction between these two dimensions is easily overlooked, because in *Corfu Channel* itself they were articulated in a single passage and applied to a single set of facts. Yet the distinction is real and consequential. The access dimension and the threshold dimension answer different questions—*when may inference be used?* and *how strong must the inference be*?—and there is no logical necessity that they evolve in parallel, or even in the same direction. A rule could, in principle, widen access while holding the threshold constant; or raise the threshold while holding access constant; or move both in the same direction, or in opposite directions. Which of these patterns the jurisprudence has in fact exhibited is an empirical question about the case law, and it is to that question that the next two Parts are addressed.

It bears emphasis that the rule, as formulated in *Corfu Channel*, was not presented by the Court as an innovation peculiar to international law. The observation that circumstantial evidence "is admitted in all systems of law" situates the rule within the broader tradition of inferential proof common to domestic legal orders, and aligns with the understanding, developed in the foundational scholarship on general principles, that international courts draw upon principles of evidence and procedure recognised across national systems (Cheng, 1953; Riddell and Plant, 2009)[10]. The rule is thus best understood not as a sui generis indulgence but as the transposition into the inter-State setting of a mode of proof that is, in its essentials, universal. This pedigree matters for the analysis that follows, because it suggests that the constraints which attend inferential proof in domestic systems—above all, the concern that inference not be permitted to substitute for proof where the consequences of error are grave—are equally available to the international judge as resources for shaping the rule's evolution.

The Analytical Payoff of Disaggregation

The disaggregation of the rule into access and threshold dimensions is not a mere

[10]The classic study remains B Cheng, *General Principles of Law as Applied by International Courts and Tribunals* (Cheng, 1953), which examines the application by international tribunals of principles of proof and procedure drawn from domestic legal systems. For modern treatments of evidence and fact-finding before the Court, see Riddell and Plant (2009) and Devaney (2016), ch 1.

taxonomic convenience; it is an analytical instrument that makes visible a structure the unitary reading conceals. The dominant treatment of the circumstantial-evidence rule, in the commentary and in the practice alike, runs the two dimensions together, speaking of "the rule on circumstantial evidence" as though it were a single parameter that is either more or less favourable to applicants. On that approach, the protective face of the rule (its accommodation of applicants facing exclusive control) and its restrictive face (its demand that inferences be compelling) appear as competing characterisations of one thing, and the analyst is left to ask which characterisation is the more accurate—whether the rule is, on balance, a friend or a foe to applicants.

The disaggregated approach dissolves this false choice. Once access and threshold are recognised as distinct parameters, it becomes apparent that the rule can be, and is, protective along one dimension and restrictive along the other: it is protective in *whom* and *what* it admits to inferential proof (the access dimension, which has expanded), and restrictive in *how much* it demands of that proof (the threshold dimension, which has risen). The protective and restrictive faces are not competing descriptions of a single parameter but accurate descriptions of two different parameters moving in two different directions. The apparent paradox with which this article opened—that the rule is at once accommodating and near-prohibitive—is thus revealed as no paradox at all, but the predictable appearance of a two-dimensional structure viewed as though it were one-dimensional. The disaggregation is, in this sense, the key that unlocks the entire analysis: it is only because the two dimensions are held apart that their independent trajectories, and ultimately their functional relationship, can be brought into view.

3. Track One: The Horizontal Expansion of Access

The first of the two tracks concerns the access dimension—the range of circumstances in which the Court has been willing to permit resort to inferential proof. Traced across the principal cases, the jurisprudence reveals a steady widening of that range. The rule has migrated from a narrow application to physical events within a respondent's territory to an expansive application embracing the policies of States and the mental states of their officials. This Part documents that migration in three stages.

3.1. From Physical Facts to Clandestine Attribution: *Nicaragua*

The first significant extension came in *Military and Paramilitary Activities in and against Nicaragua*. The case concerned, among much else, whether the United States bore responsibility for its support of the contra forces operating against the Government of Nicaragua. The core facts—the decisions taken within the United States Government, the operations of its intelligence agencies, the scale and character of the assistance furnished to the contras—were, like the facts in *Corfu Channel*, substantially within the exclusive control of the respondent. Nicaragua's evidentiary position was, in this respect, structurally comparable to that of the

United Kingdom in *Corfu Channel*, although the Court did not frame its approach in those terms[11].

The Court cited *Corfu Channel* on the weight of unconfirmed testimony and, elsewhere in the judgment, for the proposition that a State's knowledge cannot be inferred from the mere fact of its control over its territory—a proposition the Court applied to the applicant's own territory, in relation to the alleged flow of arms to the armed opposition in El Salvador, rather than as a ground for relaxing proof against the respondent[12]. Read alongside *Corfu Channel*, its treatment of the evidence illustrates a wider use of inferential proof. The facts at issue in *Corfu Channel* were physical and localised: the presence of a minefield, the observability of its laying. The facts at issue in *Nicaragua* were of a different order: they concerned the *attribution* to a State of clandestine conduct effected through intermediaries, conduct whose very design was to remain deniable. In addressing such facts the Court had recourse to inferential and circumstantial proof, while cautioning that evidence drawn from secondary sources had to be treated with care and weighed for its reliability[13].

The significance of *Nicaragua* for the access dimension is twofold. First, it confirmed that the rule was not confined to the species of facts before the Court in *Corfu Channel*—physical events within territorial waters—but extended to the proof of clandestine State conduct more generally. The rule thereby acquired a generality it had not previously been shown to possess: it was revealed as a tool applicable wherever exclusive control over facts obstructs direct proof, whatever the nature of those facts. Second, *Nicaragua* introduced an important qualification concerning the *evaluation* of inferential material, treating press articles and books with caution and regarding them not as evidence capable of proving facts but as material capable, at most, of corroborating facts established by other means[14]. It should be stressed that this qualification is a rule about the *admissibility and weight of particular categories of source material*—about which items of evidence may be relied upon, and for what purpose—and not a heightening of the *standard of probative force* required of the inference ultimately drawn. To insist that a newspaper report cannot by itself prove a fact is to say something about the evidentiary value of newspaper reports; it is not to say that, once a body of admissible evidence has been assembled, the inference it supports must meet a more exacting standard than the "single conclusion" formula demanded. The distinction matters for the argument of this article, because it shows that *Nicaragua*'s evidentiary caution operated on the *access* side of the analysis—refining what materials may feed an inference—without disturbing the *threshold*. It is worth adding, to forestall a

[11] *Military and Paramilitary Activities in and against Nicaragua* (*Nicaragua v United States of America*) (*Merits*) [1986] ICJ Rep 14.

[12] See *Nicaragua* (n 11) paras 68 and 155, citing *Corfu Channel* on the weight of unconfirmed hearsay and the limits of inferring knowledge from territorial control, respectively.

[13] See *Nicaragua* (n 11) paras 57, 62-63 and 68.

[14] *Nicaragua* (n 11) paras 60 and 62-63. Paragraph 63 distinguishes the use of press material to establish that a fact was publicly known.

misreading, that this caution was not the seed from which the later threshold grew. As Part 4.1 shows, the alternative-excluding test that came to govern inferences of intent was not an internal elaboration of the Court's own practice but was drawn from a different jurisprudence altogether. What *Nicaragua* does show is a disposition to scrutinise the materials of inference—a disposition consistent with, but genealogically distinct from, the threshold that the Court would later import.

It is worth observing what *Nicaragua* did *not* do. Although it widened the access dimension, it did not disturb the threshold formula inherited from *Corfu Channel*. In establishing the facts imputable to the United States, the Court proceeded on the footing that it enjoyed freedom in estimating the value of the various elements of evidence, and its reasoning can be read as consistent with the "single conclusion" approach to the weight of circumstantial evidence; it did not announce, or apply, any heightened standard of probative force of the kind later articulated in *Bosnia v. Serbia*[15]. In the terms of this article, *Nicaragua* moved the horizontal axis without moving the vertical one. This observation is important, because it demonstrates that the two axes are genuinely separable in the practice of the Court, and not merely in analytical theory: here is a case that advanced one while leaving the other in place.

3.2. From Conduct to Policy and Intent: The *Genocide Convention* Cases

The most consequential expansion of access occurred in the two *Genocide Convention* cases between Bosnia and Herzegovina and Serbia, and between Croatia and Serbia. Here the rule was applied not merely to the attribution of clandestine conduct but to the proof of a State's genocidal *intent*—a matter that is at once subjective, collective, and of the utmost gravity.

In *Bosnia v. Serbia*, the Court was required to determine whether acts of genocide had been committed and, if so, whether they were attributable to the respondent[16]. The proof of genocidal intent—the *dolus specialis* of the crime—presented an evidentiary challenge of a kind not previously confronted under the *Corfu Channel* rule. Intent is not a physical fact that can be observed from a shoreline, nor even a clandestine operation that can be reconstructed from documents and reports. It is a mental state, and where it is the intent of a State rather than of an identifiable individual, it must be discerned from the pattern of conduct of many actors. The Court accepted, consistently with the logic of *Corfu Channel*, that in-

[15]See *Nicaragua* (n 11) paras 57, 60 and 75-125, in which the Court established the facts imputable to the respondent on the basis of the evidence before it. The Court's reasoning in this section proceeds without articulating or applying any standard of probative force more exacting than that stated in *Corfu Channel*; the heightened "fully conclusive" standard for charges of exceptional gravity was expressly restated two decades later in *Bosnia v Serbia* (n 16) para 209.

[16]*Application of the Convention on the Prevention and Punishment of the Crime of Genocide* (*Bosnia and Herzegovina v Serbia and Montenegro*) (*Judgment*) [2007] ICJ Rep 43. The Court's determination whether genocide had been committed and, if so, whether it was attributable to the respondent occupies the greater part of the Judgment; the standard of proof is stated at para 209 and the treatment of inference of intent from a pattern of conduct at para 373.

tent could be established inferentially, from a pattern of conduct; the access dimension thus expanded once more, now to embrace the proof of subjective intent[17].

A clarification is required here, because the inference of a *State's* genocidal intent involves a conceptual step that the inference of physical facts does not, and the coherence of the article's account depends on identifying it. A State has no mind of its own; its intent, in the sense relevant to responsibility, is a construct derived from the intent of the natural persons whose conduct is attributable to it. The proof of State genocidal intent by inference therefore presupposes, and operates alongside, the ordinary rules of attribution. The Court's analysis in the *Genocide Convention* cases in fact proceeds in two logically distinct stages that it is important not to conflate: first, whether the underlying acts constitute genocide, which requires that those acts be accompanied by the *dolus specialis*; and second, whether those acts, and the intent with which they were committed, are attributable to the respondent State under the law of State responsibility—whether they were the acts of its organs, or of persons or groups acting on its instructions or under its direction or control[18]. The inference of intent from a pattern of conduct operates within the first stage: it is the intent accompanying the relevant acts that is inferred from the pattern those acts compose. But the intent so inferred becomes the *State's* intent, for the purposes of responsibility, only in so far as the acts it accompanies are attributable to the State under the second stage. The two inquiries are analytically independent—a State may bear responsibility for failing to prevent genocide committed by actors whose conduct is not attributable to it, as *Bosnia v. Serbia* held—but they are practically intertwined, because the pattern of conduct from which intent is inferred is, characteristically, the conduct of the very organs and agents whose acts the applicant seeks to attribute to the State. For the purposes of this article, the significant point is that the expansion of the access dimension to embrace subjective intent does not dispense with attribution; it adds, to the already demanding task of attributing conduct to a State, the further task of establishing inferentially the intent with which that conduct was carried out. This compounding of evidentiary burdens is part of what makes the threshold question, addressed in Part 4, so consequential in this domain.

Croatia v. Serbia confirmed and consolidated this expansion. The Court reiter-

[17]ibid para 373, at 196-197 (holding that "for a pattern of conduct to be accepted as evidence of [the] existence [of *dolus specialis*], it would have to be such that it could only point to the existence of such intent"); see also *Croatia v Serbia* (n 3) paras 145 and 148.

[18]On the distinction between the commission of genocide and its attribution to a State, and the Court's application of the rules of attribution reflected in Articles 4 and 8 of the International Law Commission's Articles on Responsibility of States for Internationally Wrongful Acts, see *Bosnia v Serbia* (n 16): the test of responsibility at paras 379-384; attribution on the basis of the conduct of State organs (Article 4) at paras 385-395; and attribution on the basis of direction or control (Article 8, applying the "effective control" test) at paras 396-412. On the decoupling of the obligation to prevent genocide from responsibility for its commission—such that a State may be responsible for a failure to prevent even where the underlying acts are not attributable to it—see ibid paras 428-438.

ated that the *dolus specialis* of genocide could be inferred from a pattern of conduct, and proceeded to examine an extensive evidentiary record concerning the conduct of the parties during the conflict, with a view to determining whether such an inference was warranted[19]. By the time of *Croatia v. Serbia*, then, the access dimension of the rule had travelled a considerable distance from its origin. What had begun in *Corfu Channel* as a means of proving a State's knowledge of physical events within its own waters had become a means of proving a State's genocidal intent from the aggregate pattern of conduct attributable to it.

3.3. The Trajectory of Track One

Viewed as a whole, the jurisprudence on the access dimension describes a clear trajectory of expansion. The categories of fact to which inferential proof has been applied have widened at each stage: physical and localised events (*Corfu Channel*); the attribution of clandestine conduct to a State (*Nicaragua*); the existence of a State policy and the subjective intent underlying a pattern of conduct (*Bosnia* and *Croatia*). At no point has the Court retreated from a category of access it had previously recognised; the movement has been consistently outward.

The consequence of this expansion is that the rule on circumstantial evidence has been transformed in its practical significance. In *Corfu Channel* it was a specific response to a specific kind of evidentiary obstruction. By the time of the *Genocide Convention* cases it had become a general instrument for the proof of facts within a respondent's exclusive control, applicable across the full range of matters that arise in inter-State litigation—including the most grave and the most resistant to direct proof. This generality is, in one sense, a measure of the rule's success: it has proved adaptable to evidentiary predicaments far removed from the one that occasioned it. But the very breadth of the rule's application raises a question that the access dimension cannot answer on its own. If inferential proof may now be deployed to establish matters as consequential as a State's genocidal intent, what is to prevent the over-extension of inference—the drawing of grave conclusions from patterns of conduct that are in truth equivocal? The answer, this article argues, lies along the second track, to which the analysis now turns.

4. Track Two: The Vertical Elevation of the Threshold

Part 3 has documented the outward movement of the access dimension; this Part turns to a movement in the opposite direction. If the access dimension has expanded horizontally, the threshold dimension has risen vertically—the Court has demanded progressively more of the circumstantial evidence it will act upon, even as it has permitted that evidence to be applied to progressively graver matters. Yet to trace the vertical elevation accurately, one must first understand what has risen and where it came from. The modern doctrine is commonly presented as though

[19] *Croatia v Serbia* (n 3) para 148; and see paras 143-147 on the inference of *dolus specialis* from a pattern of conduct.

the "only reasonable inference" formula were the latest stage of a single standard descending from *Corfu Channel*. That presentation is, as Part 4.1 will show, wrong in a way that matters. The formulations in play belong to three genealogically distinct lines of jurisprudence—lines that converged at a specific historical moment and in a specific doctrinal domain, and whose convergence is itself part of the phenomenon this article seeks to explain.

4.1. The Genealogy of the Standards: Three Lines of Jurisprudence

The formulations that the Court has deployed in this area—"conclusive evidence", "fully conclusive", "could only point to", "the only inference that could reasonably be drawn", "plausible"—are commonly treated as variations on a single theme, as though each were a reformulation of its predecessor in a continuous doctrinal descent from *Corfu Channel*. That treatment is natural enough: the formulations do appear in a single line of authority, they are invoked alongside one another in the same judgments, and they share the surface feature of addressing the strength of evidence required. But the appearance of continuity is misleading. Beneath it lie three genealogically distinct lines of jurisprudence, with different points of origin, different institutional functions, and—crucially—different normative calibrations. The failure to disentangle them has produced two errors that run through the existing commentary: a misattribution of the "only reasonable inference" test to the Court's own evidentiary tradition, and a consequent mischaracterisation of what has to be explained. What has to be explained is not why the Court raised its own standard—for in its decisive component it did not raise its own standard but adopted someone else's. This section sets the record straight.

The first line: the standard of proof for grave charges. This line is internal to the Court and runs continuously from its first contentious case. In Corfu *Channel* the Court, in the passages preceding its treatment of circumstantial evidence, employed several formulations to describe what it required before it would hold a State responsible on a charge of exceptional gravity: it spoke of decisive legal proof, of conclusive evidence, and of a degree of certainty (Del Mar, 2012)[20]. Nearly six decades later, in *Bosnia v. Serbia*, the Court expressly invoked *Corfu Channel* for what it described as the long-recognised position that "claims against a State involving charges of exceptional gravity must be proved by evidence that is fully conclusive", requiring that it be "fully convinced" that the alleged acts had been "clearly established"[21]. *Croatia v. Serbia* reiterated the same standard[22]. A precision is required here that the literature does not always observe: *Bosnia v. Serbia* applied this exacting standard to the allegations of the commission of gen-

[20] *Corfu Channel* (n 1) 17. On the several formulations there employed, and on the Court's standards of proof more generally, see Del Mar (2012).
[21] *Bosnia v Serbia* (n 16) para 209.
[22] *Croatia v Serbia* (n 3) para 178.

ocide and to their attribution, but applied a distinct and less demanding standard—"proof at a high level of certainty appropriate to the seriousness of the allegation"—to the separate claims that the respondent had failed to prevent and to punish genocide[23]. The "fully conclusive" standard is therefore neither a novelty of 2007 nor a uniform requirement across all claims; it is a differentiated standard of proof with a genuine *Corfu Channel* pedigree.

The second line: the "only reasonable inference" rule. This line is not internal to the Court at all. The notion of an "only reasonable inference" entered international law through the jurisprudence of the modern ad hoc criminal tribunals, where it developed as a function of the criminal standard of proof beyond reasonable doubt (Milanovic, 2025)[24]. Two features of its origin deserve emphasis, because both bear directly on the argument of this article. First, in its native habitat the rule is not specific to genocidal intent, nor indeed to mental states at all: the ad hoc tribunals applied it across a wide range of cases, to facts relevant to the *actus reus* as well as to *mens rea*. It is, in origin, a general rule about how circumstantial evidence should be evaluated—which is precisely the function this article ascribes to the threshold dimension. Second, the rule is internally calibrated to the demands of criminal conviction: it is the inferential counterpart of proof beyond reasonable doubt, and its stringency is a function of that standard, which exists to protect individuals against wrongful conviction. Alongside this general rule, the ad hoc tribunals also developed the specific practice of inferring genocidal intent from the surrounding facts and circumstances where direct evidence of a plan or policy is absent (Barigye et al., 2025)[25]. On the account offered here, it was this body of jurisprudence, rather than any internal development of the *Corfu Channel* formula, that best explains the language the Court adopted in *Bosnia v. Serbia*—that a pattern of conduct "could only point to" the existence of specific intent—and crystallised in *Croatia v. Serbia* as the requirement that the applicant's inference be "the only inference that could reasonably be drawn". The Court itself did not attribute the 2007 formulation to the tribunals; in *Croatia v. Serbia* it set the Tribunal's formulation alongside its own and treated the two as substantively the same. The connection is not merely one of vocabulary: in both cases the Court relied extensively on the findings of the International Criminal Tribunal for the

[23] *Bosnia v Serbia* (n 16) paras 209-210. The distinction between the two standards—"fully conclusive" for the commission and attribution of genocide, and "proof at a high level of certainty appropriate to the seriousness of the allegation" for the obligations to prevent and to punish—is frequently elided in the commentary but is material to any account of the Court's evidentiary practice.

[24] See Milanovic (2025), tracing the "only reasonable inference" formula to the jurisprudence of the ad hoc criminal tribunals, where it functions as the inferential counterpart of the standard of proof beyond reasonable doubt and is applied to circumstantial evidence generally rather than to the proof of specific intent alone. Milanovic does not, however, suggest that the Court took the formula from the tribunals; that account is advanced in the present article. On the relationship between the formula and the criminal standard of proof, see also Ambos (2016). For a separate analysis of the Karadžić trial judgment, see Sterio (2017).

[25] See *Prosecutor v Akayesu* (Judgment) ICTR-96-4-T (2 September 1998) para 523; *Prosecutor v Jelisić* (Appeal Judgment) IT-95-10-A (5 July 2001) para 47. For discussion of this practice and its reception by the Court, see Barigye et al. (2025).

former Yugoslavia, findings themselves reached under the criminal standard of proof[26].

The third line: *plausibility at the provisional-measures stage*. *This line is again internal to the Court but entirely separate in function from the first two*. *Since its 2009 order in* Questions relating to the Obligation to Prosecute or Extradite (Belgium v. Senegal), *the Court has required*, *as a condition of indicating provisional measures*, *that the rights asserted by the applicant be* "*at least plausible*"; *the requirement has been applied consistently in the subsequent practice*, *including in* The Gambia v. Myanmar *and* South Africa v. Israel[27]. Plausibility is not a standard of proof at all in the sense in which the first two lines are concerned with proof: it is a threshold of arguability governing access to interim protection, and it operates on the asserted rights rather than on the evidence for the underlying facts.

The relationship between the lines. *Once these genealogies are set out*, *the structure of the modern doctrine becomes clearer*, *and the account offered in this article requires a corresponding refinement*. *The horizontal expansion traced in Part 3 is a development within the* Corfu Channel *line*: *it is that line's access rule that has been extended from physical facts to subjective intent*. *The vertical elevation traced in what follows is*, *by contrast*, *of* hybrid *genealogy*. *Its standard-of-proof component* ("*fully conclusive*") *is a* Corfu Channel *inheritance*, *applied with increasing explicitness to grave charges*. *But its inferential component—the alternative-excluding structure of the* "*only reasonable inference*" *test*, *which does the decisive work in the cases—was* borrowed from the criminal line. The elevation of the threshold is therefore not simply the internal ripening of a 1949 formula; it is the importation, into the law of State responsibility, of an evidentiary technique developed for and calibrated to criminal conviction.

This refinement does not displace the article's central claim; it sharpens it. The claim is not, and was never, that the *Corfu Channel* rule spontaneously became more demanding as it aged. It is that, when the access dimension expanded into the domain of subjective intent—a domain in which the ad hoc tribunals had already developed a body of technique calibrated to the most demanding standard of proof known to law—the Court, confronted with the risk of over-extension that this expansion had created, did not extend its own "single conclusion" formula to meet it but reached across to a neighbouring jurisprudence and took from it a ready-made and markedly more stringent test. The borrowing *is* the compensat-

[26]On the Court's reliance on the findings of the International Criminal Tribunal for the former Yugoslavia, reached under the criminal standard of proof, see *Bosnia v Serbia* (n 16) paras 212-223; and *Croatia v Serbia* (n 3) paras 434-436 and 440, where the Court drew on the Tribunal's findings in assessing the objective of forced displacement. On the Court's treatment of the Tribunal's formulation as substantively the same as its own 2007 criterion, see *Croatia v Serbia* (n 3) paras 146-148.

[27] *Questions relating to the Obligation to Prosecute or Extradite* (*Belgium v Senegal*) (*Provisional Measures*, *Order of 28 May 2009*) [2009] ICJ Rep 139, para 57. For the subsequent development of the requirement, see, among others, *Application of the Convention on the Prevention and Punishment of the Crime of Genocide* (*The Gambia v Myanmar*) (*Provisional Measures*, *Order of 23 January 2020*) [2020] ICJ Rep 3, paras 43-44; and *Application of the Convention on the Prevention and Punishment of the Crime of Genocide in the Gaza Strip* (*South Africa v Israel*) (Provisional Measures, Order of 26 January 2024) [2024] ICJ Rep 3, para 54.

ing move. And that reframing opens a question which Part 5 takes up and which the existing debate has left largely implicit: whether a threshold whose stringency is calibrated to the protection of individuals against wrongful criminal conviction is well suited to proceedings concerning the responsibility of States.

4.2. The Stratification of the Standard of Proof: *Bosnia v. Serbia*

The first explicit elevation of the threshold came in *Bosnia v. Serbia*, not in the form of a new formula for the weight of circumstantial evidence as such, but through the articulation of a heightened standard of proof for grave allegations. As the preceding section has shown, the Court did not present this as an innovation: it invoked *Corfu Channel* for the long-recognised position that claims against a State involving charges of exceptional gravity must be "proved by evidence that is fully conclusive", and required that it be "fully convinced" that the allegations had been "clearly established". Allegations of the commission of genocide, the Court made clear, fell within this category—as did their attribution to the respondent. To the distinct claims that the respondent had failed to prevent and to punish genocide, the Court applied a materially less demanding standard, requiring "proof at a high level of certainty appropriate to the seriousness of the allegation".

What is doctrinally significant here is not the novelty of the "fully conclusive" standard, which as we have seen has a *Corfu Channel* pedigree, but the explicit *stratification* of the standard of proof according to the gravity and character of the allegation. The Court did not apply a single, uniform standard across all questions of fact; it differentiated, demanding "fully conclusive" evidence for the gravest charges while applying a less exacting standard to the distinct claims concerning prevention and punishment (Del Mar, 2012)[28]. This stratification is significant for the threshold dimension of the circumstantial-evidence rule because, where a fact of exceptional gravity is to be established by inference, the heightened standard of proof governs the evaluation of that inference. The circumstantial evidence must, in such cases, be strong enough to meet the "fully conclusive" benchmark—a benchmark notably more demanding than the "special weight" attaching to facts "leading logically to a single conclusion" in *Corfu Channel*.

The relationship between the gravity of the allegation and the stringency of proof was thus made explicit: the graver the matter to be established, the more compelling the inferential evidence must be, and the more rigorously the links in the chain of inference must be scrutinised. This is the conceptual hinge on which the subsequent elevation in *Croatia v. Serbia* would turn.

4.3. From "Single Conclusion" to "Only Reasonable Inference": *Croatia v. Serbia*

The decisive elevation of the threshold occurred in *Croatia v. Serbia*. Addressing

[28] *Bosnia v Serbia* (n 16) paras 209-210. On the Court's approach to standards of proof more generally, see Del Mar (2012).

the proof of *dolus specialis* by inference from a pattern of conduct, the Court held that, "in order to infer the existence of *dolus specialis* from a pattern of conduct, it is necessary and sufficient that this is the only inference that could reasonably be drawn from the acts in question"[29]. This formulation—"the only inference that could reasonably be drawn"—is the high-water mark of the threshold's elevation, and its relationship to the original *Corfu Channel* formula requires careful examination.

At first sight the two formulae might appear to say much the same thing. *Corfu Channel* spoke of facts "leading logically to a single conclusion"; *Croatia v. Serbia* speaks of an inference being "the only [one] that could reasonably be drawn". Both seem to require that the circumstantial evidence point to one conclusion rather than several. Yet on closer inspection the two formulae allocate the burden of inferential proof differently, and the difference is one of substance.

The "single conclusion" formula is, in its orientation, *positive*: it asks whether the series of facts, taken together, points to the conclusion the applicant urges. The inquiry is directed at the affirmative coherence of the applicant's case—whether the facts, linked together, lead to the proposed conclusion. The "only reasonable inference" formula is, by contrast, *negative* in a way the earlier formula is not: it asks not merely whether the facts point to the applicant's conclusion, but whether that conclusion is the *only* one that can reasonably be drawn—which is to say, whether every competing reasonable explanation has been excluded. Under the former, the applicant succeeds by establishing that the facts support its conclusion; under the latter, the applicant succeeds only by additionally establishing that the facts *cannot reasonably support any alternative*.

This shift—from pointing to a conclusion to excluding all reasonable alternatives—effects a material transfer of the burden of inferential proof. It is no longer enough for the applicant to construct a chain of facts that coheres around its desired conclusion; the applicant must also negate the field of competing explanations. Correspondingly, the respondent's task is eased: the respondent need not displace the applicant's chain in its entirety, but need only advance a single alternative explanation that is itself reasonable, for the existence of such an alternative defeats the proposition that the applicant's inference is the *only* reasonable one. The "only reasonable inference" standard thus arms the respondent with a powerful defensive instrument—the reasonable alternative—that the "single conclusion" formula did not so clearly supply.

The structure of the difference repays closer analysis, for it is not merely a matter of emphasis. Consider a body of circumstantial evidence that admits of two reasonable interpretations, one of which corresponds to the conclusion the applicant urges. Under a formula oriented to whether the facts "lead logically to a single conclusion", the applicant's position is not hopeless: the inquiry is directed at whether the facts, taken as a coherent series, point to the proffered conclusion,

[29] *Croatia v Serbia* (n 3) para 148.

and the existence of some other reading does not necessarily defeat that coherence so long as the applicant's reading is the one to which the series most naturally leads. Under a formula requiring that the applicant's inference be "the only [one] that could reasonably be drawn", the same body of evidence yields a different result: the existence of a second reasonable interpretation is, by definition, fatal, for it establishes that the applicant's inference is not the only reasonable one. The two formulae thus diverge precisely in the case of genuinely equivocal evidence—evidence consistent with more than one reasonable account—and they diverge in a manner systematically unfavourable to the applicant under the later formula. Where evidence is equivocal, the "single conclusion" approach may permit a finding for the applicant if its reading is the more natural; the "only reasonable inference" approach forbids it so long as the competing reading remains reasonable.

This analysis also clarifies what the "only reasonable inference" standard does *not* require, a point on which the contemporary commentary has rightly insisted. The standard does not demand the exclusion of every *conceivable* alternative, however fanciful; it demands the exclusion only of *reasonable* alternatives. An applicant is not defeated by the bare logical possibility of an innocent explanation, but only by an alternative explanation that a reasonable tribunal would regard as genuinely open on the evidence. The qualifier "reasonably" thus performs essential work: it prevents the standard from collapsing into an impossible demand for the negation of all logical possibilities, while still requiring that the field of *plausible* competing explanations be foreclosed. The burden the standard imposes is therefore demanding but not, in principle, impossible: it is the burden of showing that, among the explanations a reasonable tribunal would entertain, the applicant's is the only one the evidence will bear.

A schematic example, deliberately abstracted from any actual dispute, makes the divergence concrete. Suppose an applicant seeks to establish that a respondent State forcibly deported a protected group with the intent to destroy it in part, and adduces a record showing the following: that the group was systematically expelled from a defined region; that the expulsions were executed with considerable violence; that a proportion of the group died in the course of the expulsions; and that officials of the respondent made statements expressing hostility toward the group. Suppose further that the respondent adduces, or the record itself discloses, that the expelled population was permitted to cross an international frontier and settle beyond it, that transport was in some instances provided, and that the region was one the respondent sought to control for reasons connected to the armed conflict. Now consider how this single record fares under each formula. Under the "single conclusion" approach, the applicant may succeed: the assembled facts—expulsion, violence, deaths, hostile statements—form a coherent series that points to an intent to destroy, and a tribunal persuaded that this is the reading to which the series most naturally leads could find the intent established, notwithstanding the presence of some contrary indications. Under the "only reasonable inference" approach, the same record fails: the facts concerning the permitted crossing of the

frontier, the provision of transport, and the strategic value of the territory sustain a competing and reasonable explanation—that the respondent's intent was to remove the group from the region (forcible displacement) rather than to destroy it as such—and the existence of that reasonable alternative is, by the terms of the standard, fatal to the inference of genocidal intent, however well the applicant's own reading coheres. The record has not changed; what has changed is the question asked of it. The first formula asks whether the applicant's inference is *supported*; the second asks whether it is *uniquely* supported, to the exclusion of reasonable competitors. The same evidence that satisfies the former can fail the latter, and it is precisely in the space between these two questions that the elevation of the threshold does its work.

The application of the standard in *Croatia v. Serbia* illustrates the point in practice, and along strikingly similar lines to the schematic example just given. The Court accepted that the evidentiary record disclosed a pattern of conduct of considerable breadth, and that this pattern could support the conclusion that grave violations of international humanitarian law had occurred. But it held that the inference of genocidal intent was not the only one that could reasonably be drawn: the same pattern of conduct was, in the Court's assessment, consistent with an intent to compel the displacement of a population rather than to destroy it as such[30]. Because a competing reasonable explanation existed, the inference of *dolus specialis* could not be sustained, and the claim failed. It was the *availability of a reasonable alternative*, not any deficiency in the affirmative coherence of the applicant's chain of facts, that proved decisive. This is the operative signature of the elevated threshold.

4.4. The Contemporary Debate over Attainability

The elevation of the threshold has not gone unremarked. A substantial body of contemporary commentary has questioned whether the "only reasonable inference" standard, as applied by the Court, sets the bar of proof so high that findings of genocidal intent become practically unattainable. The concern has been voiced not only by academic commentators but, strikingly, by States: in their joint declaration of intervention in *The Gambia v. Myanmar*, a group of States cautioned against an interpretation of the standard that would render the threshold for inferring genocidal intent so demanding as to make findings of genocide near-impossible[31]. Others have defended the standard, arguing that it is satisfied where a pattern of conduct cannot be explained other than by genocidal intent, and that properly understood it does not require the exclusion of fanciful or merely conceivable alternatives but only of reasonable ones[32].

The debate has tended to crystallise around a single question: whether the

[30] *Croatia v Serbia* (n 3) paras 434-440.

[31] See the Joint Declaration of Intervention referred to at n 3 above. For scholarly analysis of the standard for demonstrating genocidal intent before the Court, see Van Santen (2025).

[32] Barigye et al. (2025).

standard, considered in itself, is pitched at an appropriate level. Some commentators read the Court's "fully conclusive" requirement as approximating the criminal standard of proof beyond reasonable doubt, and question whether such a standard is appropriate to proceedings concerning State responsibility; others read it as admitting of gradations and resist the assimilation to the criminal standard (Del Mar, 2012; Van Santen, 2025)[33]. What unites most contributions to the debate, however, is that they engage the threshold *in isolation*—they ask whether the vertical axis has been set too high, taken on its own terms.

It is precisely this isolation that Part 5 will question. For present purposes, the point is to register that the elevation of the threshold documented above is not a matter of merely doctrinal interest but is at the centre of an active and consequential controversy, and that the terms in which that controversy has been conducted have largely abstracted the threshold from the conditions of access that accompany it.

4.5. Stratification across Procedural Phases: *South Africa v. Israel*

A further dimension of the threshold's operation has come into view in the most recent practice, concerning the standard applicable at the provisional-measures phase as distinct from the merits. In its order on provisional measures in *South Africa v. Israel*, the Court applied the "plausibility" standard that governs the indication of provisional measures, asking whether the rights asserted by the applicant were at least plausible, rather than whether the applicant's allegations had been established to the standard applicable at the merits[34]. The Court did not, at the provisional-measures stage, apply the "only reasonable inference" standard; that standard is, by its nature, reserved for the merits, where the question is whether the relevant intent has been proved, not whether the asserted rights are plausible[35].

The precise evidentiary significance of the order must be stated carefully, lest it be overread. *South Africa v. Israel* is invoked in this article for one proposition only: that the Court calibrates the evidentiary demand it makes according to the function of the procedural phase, applying a threshold of mere plausibility at the interlocutory stage. The order does *not* pronounce upon, alter, or supply support for the standard governing the proof of genocidal intent at the merits. A finding that asserted rights are "plausible" for the purpose of indicating provisional measures is not a finding, even provisionally, that genocidal intent has been or can be established; the plausibility of a right is a threshold for protective interloc-

[33]The question whether the "fully conclusive" standard is equivalent to a single criminal-style standard of proof beyond reasonable doubt, or instead admits of internal gradation, is canvassed in Van Santen (2025: pp. 175-182), which sets out the competing lines of reasoning. See also Del Mar (2012) on the Court's approach to standards of proof.

[34]*South Africa v Israel* (n 27) paras 54, 59-62 and 74. The plausibility standard governs only the indication of provisional measures and is not a standard for the proof of the merits.

[35]ibid.

utory relief, not a measure of the probative sufficiency of evidence of intent. It would therefore be a serious misreading—and one this article expressly disclaims—to treat the order as evidence that the merits threshold is, or should be, lower than the "only reasonable inference" standard articulated in *Croatia v. Serbia*, or indeed to draw from it any inference about the merits standard at all. The order is relevant here not as authority on the *height* of the merits threshold but as an illustration of the *phenomenon of phase-specific calibration*—the very phenomenon that, on the account offered in this article, also operates as between different categories of fact at the merits stage itself. Its evidentiary value to the argument is structural and analogical, not doctrinal: it confirms that the Court varies what it demands of evidence according to the function of the determination in hand, and it does so without the article relying on it to establish anything about the merits standard for genocidal intent, which rests entirely on the merits authorities discussed in Parts 4.2 and 4.3.

This stratification across procedural phases is consistent with, and indeed reinforces, the account of the threshold offered here. It demonstrates that the Court calibrates the demands it makes of an applicant's evidence according to the function of the procedural stage in question: a relatively undemanding standard of plausibility suffices to justify the protective, provisional, and reversible measures appropriate to an interlocutory phase, whereas the exacting "only reasonable inference" standard governs the definitive determination of responsibility at the merits. The differentiation across phases is, in effect, a further instance of the principle articulated in *Bosnia v. Serbia*: the stringency of proof tracks the gravity and finality of the determination to be made.

The oral proceedings on the merits in *The Gambia v. Myanmar* took place in January 2026. The Court delivered its judgment on the merits on 19 November 2026[36]. The proceedings underscore that the threshold dimension of the circumstantial-evidence rule remains a matter of active doctrinal development, and that the relationship between the breadth of the rule's application and the height of its threshold—the relationship this article seeks to illuminate—is of immediate practical importance.

4.6. The Trajectory of Track Two

The threshold dimension, like the access dimension, describes a clear trajectory across the jurisprudence—but a trajectory of elevation rather than expansion. Measured by what an applicant must now show, the demand has risen: the "single conclusion" formula of *Corfu Channel*, the stratified "fully conclusive" standard of *Bosnia v. Serbia*, and the "only reasonable inference" standard of *Croatia v. Serbia* mark, in that order, successively greater demands upon the circumstantial evidence on which grave findings may be based. But it bears repeating, in light of Part 4.1, that this is a trajectory in the *level of demand* rather than the internal

[36] *The Gambia* v *Myanmar* (Judgment of 19 November 2026), General List No 178, paras 1, 17 and 58 (institution in 2019, preliminary objections judgment in 2022 and merits hearings in January 2026).

development of a single doctrine. The rise is real; its mechanism is convergence rather than growth. The Court's own standard of proof became more explicitly stratified, while the operative inferential test was drawn from the jurisprudence of the ad hoc tribunals. What unites these developments is not a shared doctrinal lineage but a shared direction of effect. The shift from a positive orientation (does the evidence point to the applicant's conclusion?) to a negative one (is the applicant's conclusion the only reasonable one?) has materially increased the burden borne by applicants and correspondingly strengthened the defensive position of respondents. The recent practice confirms that this elevated threshold is now firmly established at the merits stage, even as the demands made at interlocutory stages are calibrated differently.

Two trajectories have thus been established: an outward expansion along the horizontal axis of access, and an upward elevation along the vertical axis of threshold. The question that remains—and it is the central question of this article—is whether these two trajectories are merely coincident, or whether they are connected. It is to that question that Part 5 is devoted.

5. The Functional Link: Why the Two Tracks Move Together

The preceding two Parts have established two independent trajectories in the jurisprudence: the horizontal expansion of access and the vertical elevation of the threshold. This Part advances the article's central claim—that these trajectories are not merely coincident but functionally connected—and draws out the consequences of that claim for the contemporary debate canvassed in Part 4.

5.1. The Claim Stated

The claim may be stated in a sentence: *the widening of access is the cause of which the elevation of the threshold is the compensating effect*. As the Court extended inferential proof to ever more consequential matters—from physical events, through clandestine attribution, to subjective intent—it enlarged the risk that equivocal patterns of conduct would be made to bear conclusions their ambiguity could not support. That risk was at its most acute precisely where the access dimension reached furthest: the ascription, by inference, of genocidal intent to a State. The elevation of the threshold is the Court's answer to that risk. The corrected genealogy set out in Part 4.1 makes the mechanism visible: confronted with the need for a more exacting filter, the Court did not elaborate its own formula but imported one from the criminal tribunals, a test whose stringency was already calibrated to the gravest category of adjudicative determination. The importation itself is the compensating move.

The two tracks, on this view, are two aspects of a single development. The rule could not expand along the access dimension—could not be applied to the proof of genocidal intent, in particular—without raising the question of how to prevent its misuse in that gravest of applications. The elevation of the threshold is the answer the Court has given. The breadth of the rule and the height of its threshold

are, in this sense, mutually conditioning: the wider the rule reaches, the more it must demand.

5.2. The Argument for the Claim

Four considerations support reading the relationship between the two tracks as functional rather than coincidental.

The first is the *sequence and correlation* of the developments. The elevation of the threshold did not precede the expansion of access, nor did it occur independently of it. The "single conclusion" formula remained undisturbed through *Nicaragua*, which expanded access to clandestine attribution but did not reach subjective intent. It was only when the rule was applied to the proof of genocidal intent—the furthest reach of the access dimension—that the threshold was explicitly elevated, first through the stratified standard of *Bosnia v. Serbia* and then through the "only reasonable inference" formula of *Croatia v. Serbia*. A refinement is needed here, and it strengthens rather than weakens the point. The "only reasonable inference" test did not come into existence when the Court adopted it; as Part 4.1 has shown, it had been available in the jurisprudence of the ad hoc tribunals since the late 1990s. What is significant is therefore not the moment of the test's *emergence* but the moment of its *adoption*. The Court had the test available to it for years without recourse to it, and turned to it precisely when—and in the domain where—the access dimension reached its most consequential application. A correlation between a pre-existing instrument's availability and its eventual use is weak evidence of anything; a correlation between the moment of *taking up* that instrument and the moment at which a specific structural pressure arose is considerably stronger. That is the correlation the record discloses, and it is difficult to explain on the hypothesis that the two developments are independent.

The second consideration is the *internal logic* of inferential proof, which the Court itself acknowledged in grounding the rule in principles common to all systems of law. In domestic legal traditions, the rigour demanded of circumstantial evidence characteristically increases with the gravity of the matter to be proved; the law is most exacting of inference where the stakes are highest. The Court's insistence in *Bosnia v. Serbia* that charges of exceptional gravity be established by "fully conclusive" evidence is an explicit articulation of exactly this logic[37]. If the stringency of inferential proof properly tracks the gravity of the matter, then the extension of the rule to the gravest matters (genocidal intent) necessarily entails a heightening of the stringency demanded. The elevation of the threshold is, on this view, not a contingent policy choice but an entailment of the principled logic of inferential proof once that logic is applied to the expanded domain of access.

The third consideration—and the one that the genealogy set out in Part 4.1 brings most sharply into view—is the *character of the Court's response* to the expansion. Faced with the need for a threshold adequate to the newly accessible do-

[37] *Bosnia v Serbia* (n 16) para 209.

main of subjective intent, the Court did not elaborate its own inherited formula. It did not ask what "leading logically to a single conclusion" required in the context of intent; it adopted, instead, a test developed elsewhere, in the jurisprudence of the ad hoc criminal tribunals, and calibrated there to the standard of proof for criminal conviction. That choice is difficult to explain except as a response to felt need: the *Corfu Channel* formula was available and had served for half a century, and its abandonment in favour of an imported and markedly more demanding test, precisely at the point where the rule reached its most consequential application, is strong evidence that the expansion of access was generating pressure for a constraint that the existing formula could not supply. The two developments are linked not merely because they coincide in time, but because the second is intelligible as a solution to a problem created by the first.

The fourth consideration is *functional necessity*. Suppose, counterfactually, that the Court had expanded access to the proof of genocidal intent while leaving the threshold at the original "single conclusion" level. The result would have been a rule under which a State's genocidal intent might be established by a pattern of conduct that merely *pointed to* such intent, without excluding the reasonable possibility that the conduct was attributable to some other intent—the intent to displace a population, for instance, rather than to destroy it. Such a rule would license the gravest of findings on equivocal evidence. The elevation of the threshold forecloses this possibility. It is, in this sense, functionally necessary to the responsible operation of an expanded rule: without it, the expansion of access would have outrun the evidentiary discipline required to keep inference within reliable bounds.

A word is in order about the strength of the causal claim these considerations support. It is not contended that the Court consciously elevated the threshold *in order to* compensate for the expansion of access, as though the connection were a matter of deliberate institutional design recorded in the reasoning of the judgments. The judgments do not announce any such strategy, and the argument advanced here does not depend on their having done so. The claim is rather structural: that the elevation of the threshold is intelligible as a response to the risk generated by the expansion of access, that it occurred where and when that risk became acute, and that it conforms to the internal logic of inferential proof which the Court has itself endorsed. The relationship is functional in the sense that the two developments fit together as risk and constraint, whether or not the fit was the product of conscious calculation. This is a familiar form of explanation in the analysis of judicial doctrine, where the coherence of a body of case law is frequently best understood not as the execution of a preconceived plan but as the cumulative result of decisions each responsive to the pressures of its moment, which together exhibit a structure their authors need not have explicitly intended. The dual-track account is an explanation of this kind: it identifies a structure in the jurisprudence and offers a functional rationale for it, without attributing to the Court a design it never declared.

5.3. An Objection Considered

It may be objected that the elevation of the threshold is sufficiently explained by the gravity of genocide alone, without reference to the expansion of access. On this objection, the Court raised the threshold simply because genocide is the "crime of crimes", and the heightened standard reflects the seriousness of the charge rather than any structural relationship to the breadth of the rule. The expansion of access, the objection runs, is incidental; what drove the threshold up was the subject-matter, not the scope.

This objection does not defeat the account offered here, for two reasons. First, the gravity of the subject-matter and the expansion of access are not competing explanations but complementary aspects of the same development. It is precisely *because* the access dimension expanded to embrace the proof of genocidal intent that the gravity of the subject-matter came to bear on the circumstantial-evidence rule at all. Had the rule never been extended beyond physical facts and clandestine conduct, the gravity of genocide would have had no occasion to influence the threshold of inferential proof, for inferential proof would never have been applied to it. The gravity of the subject-matter operates *through* the expansion of access, not independently of it. The two explanations thus converge: the expansion of access is the channel through which the gravity of the matter exerts its pressure on the threshold.

Second, and more decisively, the gravity-based explanation does not account for the *form* the Court's response took. If the elevation were simply a reflex of the subject-matter's seriousness, one would expect the Court to have raised its own standard—to have said, in effect, that where the charge is genocide the facts must lead to a single conclusion with particular clarity. That is not what happened. The Court adopted a test from a different jurisprudence, one developed for criminal adjudication and calibrated to the standard for conviction. Gravity alone cannot explain that choice, because gravity is a property of the allegation, not a reason to prefer an imported instrument over an available domestic one. What does explain it is the structural situation: the Court had extended inferential proof into a domain—the subjective intent of collective actors—in which its own formula offered little guidance and in which another body of tribunals had already developed a worked-out technique. The borrowing is a response to that situation, and the situation was created by the expansion of access. Gravity supplies the motive for demanding more; the expansion of access supplies both the occasion and the explanation for the particular form the additional demand took.

5.4. Consequences for the Contemporary Debate

The functional account has a direct bearing on the debate over the attainability of the "only reasonable inference" standard. That debate, as Part 4 observed, has largely engaged the threshold in isolation, asking whether the vertical axis has been pitched too high considered on its own. The dual-track framework suggests that this framing is incomplete, and that its incompleteness helps to explain why the debate has proved difficult to resolve.

If the elevation of the threshold is functionally linked to the expansion of access, then the height of the threshold cannot be evaluated in abstraction from the breadth of the rule it governs. A threshold that might appear excessively demanding if the rule were confined to its original, narrow domain may appear appropriate—even necessary—once it is recognised that the rule now reaches the proof of a State's genocidal intent. The critics who contend that the standard is unattainably high, and the defenders who contend that it is appropriately calibrated, may thus be addressing different aspects of a single structure: the critics attending to the absolute height of the vertical axis, the defenders responding, if often implicitly, to the constraint that the breadth of the horizontal axis requires. The dual-track framework does not, by itself, resolve the normative question of whether the Court has struck the right balance. But it relocates that question onto more adequate terrain. The proper question is not whether the threshold is high in the abstract, but whether the height of the threshold is *proportionate to the breadth of access*—whether, in other words, the compensating constraint is calibrated to the risk that the expansion of the rule has generated. That is a question about the relationship between the two axes, and it cannot be addressed so long as the threshold is examined in isolation from the access dimension that conditions it.

This reframing yields a more tractable way of advancing the debate. Rather than asking, in the abstract, whether "only reasonable inference" is too demanding a standard, the inquiry becomes: given that the rule now extends to the proof of genocidal intent, is the negative, alternative-excluding structure of the standard a proportionate response to the risk of over-extension that such proof entails? Posed in this form, the question admits of principled argument on both sides—argument directed at the *fit* between threshold and access—in a way that the abstract question of the threshold's height, considered alone, does not.

5.5. Dimensions of the Proportionality Inquiry

To say that the proper question concerns the proportionality of threshold to access is not to leave the matter at the level of slogan. The dual-track framework suggests at least three dimensions along which that proportionality might be assessed, each of which converts the abstract worry about the threshold's height into a more specific and answerable question.

The first dimension concerns the *gravity of the inferred fact*. The risk against which the elevated threshold guards is the risk of grave findings on equivocal evidence, and the gravity of the finding is therefore directly relevant to the proportionality of the constraint. Where the fact to be inferred is one of the utmost seriousness—a State's genocidal intent—a demanding threshold is more readily justified than where the inferred fact is comparatively anodyne. This suggests that the "only reasonable inference" standard may be well-fitted to the proof of genocidal intent specifically, without its following that the same standard should govern every application of the circumstantial-evidence rule. The proportionality inquiry thus invites a differentiation that the abstract debate about the threshold's

height tends to obscure: the standard appropriate to the gravest inferences need not be the standard appropriate to all.

The second dimension concerns the *availability of alternative explanations as a function of the subject-matter*. The alternative-excluding structure of the elevated threshold bears more heavily on applicants in some domains than in others, because reasonable alternative explanations are more readily available in some domains than in others. The proof of subjective intent is the paradigm of a domain in which alternatives abound: a given pattern of conduct will frequently be consistent with more than one intent, as *Croatia v. Serbia* itself illustrates, where conduct consistent with an intent to destroy was also consistent with an intent to displace. In such a domain, the requirement to exclude reasonable alternatives is especially exacting. Whether that exacting demand is proportionate depends on a judgment about whether, in the domain in question, the multiplicity of available explanations reflects genuine evidential equivocality (in which case the demand is appropriate) or merely the ingenuity of respondents in manufacturing alternatives (in which case it may not be). This is a question that can be investigated domain by domain, rather than resolved once and for all in the abstract.

The third dimension concerns the *interaction between the threshold and the other resources available to the Court for the management of evidence*. The stringency of the threshold does not operate in isolation; it operates alongside the Court's powers to seek evidence, to draw inferences from a party's conduct in the proceedings, and to attribute weight to the findings of other bodies. A demanding threshold may be more or less proportionate depending on the extent to which these other resources are available to mitigate the applicant's evidentiary burden. Where the Court makes vigorous use of its capacity to assist in the elucidation of facts, a high threshold may be tolerable; where it does not, the same threshold may operate to defeat meritorious claims. The proportionality of the threshold is, in this sense, relative to the broader evidentiary practice within which it is embedded—a practice whose other elements lie beyond the scope of this article but whose relevance the dual-track framework brings into view.

These three dimensions do not yield a determinate answer to the question whether the Court has struck the right balance. They are offered not as a formula but as a structured agenda for the inquiry that the dual-track framework recommends: an inquiry into the fit between the height of the threshold and the breadth of the access it governs, conducted with attention to the gravity of the inferred fact, the availability of alternatives in the relevant domain, and the other evidentiary resources at the Court's disposal. The contribution of the framework is to make that agenda visible, and thereby to move the debate beyond the impasse that the isolated consideration of the threshold has produced.

5.6. The Transposition Question: A Criminal Threshold in a State-Responsibility Setting

The genealogy set out in Part 4.1 raises a question that the dual-track framework

brings into focus and that the existing debate has largely left implicit. The "only reasonable inference" test was developed in the ad hoc criminal tribunals as the inferential counterpart of proof beyond reasonable doubt. That standard is not an arbitrary level of stringency: it is calibrated to the specific moral and institutional situation of the criminal trial, in which the State's coercive power is brought to bear on an individual and in which the asymmetry between the risk of convicting the innocent and the risk of acquitting the guilty is deliberately resolved in the defendant's favour. The question is whether a test so calibrated is appropriately transposed to proceedings between States concerning international responsibility.

There are respectable arguments on each side, and it is no part of this article's purpose to resolve them. In favour of transposition, it may be said that the gravity of a finding that a State has committed genocide is of a wholly different order from that of an ordinary finding of internationally wrongful conduct; that such a finding carries consequences—for the respondent's standing, for the historical record, and potentially for the individuals whose conduct is thereby characterised—that are in some respects comparable to those of a criminal conviction; and that the Court's own stratification of the standard of proof, distinguishing charges of exceptional gravity from ordinary claims, already embodies the principle that stringency should track gravity. On this view the borrowing is not a category mistake but a recognition that certain civil determinations approach the criminal in their seriousness.

Against transposition, it may be said that the values that justify the criminal standard do not straightforwardly obtain between States. A State is not an individual facing loss of liberty; the presumption of innocence, whose procedural expression the criminal standard is, protects a person against the coercive power of the very authority adjudicating the case, a structure with no analogue in inter-State litigation. It may further be objected that the transposition produces an incoherence within the system as a whole: if the standard for establishing State responsibility for genocide is at least as demanding as that for convicting an individual, then a State may escape responsibility on evidence that would suffice to convict its officials—a result difficult to reconcile with the Genocide Convention's structure, which addresses both individual criminal responsibility and State obligations[38]. And there is a systemic consideration: because the Court has relied heavily on findings of the ad hoc tribunals, findings themselves reached under the criminal standard and confined to the individuals and incidents charged, the effect may be to make the proof of State responsibility parasitic on the contingencies of criminal prosecution.

The dual-track framework does not adjudicate this question, but it locates it correctly. On the account offered here, the transposition was not an abstract choice about the appropriate standard for State responsibility; it was a response to

[38]For the relationship between the standards of proof in proceedings on State responsibility and in international criminal trials, see, for background, Milanovic (2025) and Van Santen (2025: pp. 175-182); the argument in the text is the present author's.

a specific structural pressure—the expansion of access into a domain where the risk of over-extended inference was gravest and where a ready-made technique was available in a neighbouring jurisprudence. Recognising that origin does not settle whether the borrowing was well judged. But it does suggest that the question is better framed as one about *fit*—whether a threshold calibrated to one institutional setting serves the purposes of another—than as the abstract question, which has dominated the literature, of whether the standard is simply "too high".

6. Implications for International Dispute Settlement

The dual-track account, though developed through the jurisprudence of the ICJ, carries implications for the conduct of international dispute settlement more broadly. This Part identifies three.

The first concerns the litigation strategy of parties before the Court. The dual-track structure entails that an applicant relying on circumstantial evidence in a grave matter bears a burden that is, in a precise sense, double. It must satisfy the access dimension by establishing that the respondent's exclusive control over the relevant facts warrants resort to inference; and it must satisfy the elevated threshold by demonstrating not merely that the facts point to its conclusion but that no reasonable alternative explanation survives. The implications for the framing of an applicant's case are considerable. The construction of a chain of facts that coheres around the desired conclusion, while necessary, is no longer sufficient; the applicant must anticipate and foreclose the reasonable alternatives that a respondent will advance. Correlatively, a respondent's most economical defensive strategy, where the elevated threshold applies, is not to dismantle the applicant's chain of facts in its entirety but to identify and substantiate a single reasonable alternative explanation, the existence of which defeats the proposition that the applicant's inference is the only reasonable one. The reasonable alternative is, under the elevated threshold, the respondent's instrument of choice.

The second implication concerns the stratification of standards across procedural phases. The differentiation, evident in the recent practice, between the plausibility standard governing provisional measures and the "only reasonable inference" standard governing the merits has a strategic dimension. An applicant seeking interlocutory protection need not, at that stage, marshal evidence sufficient to exclude all reasonable alternatives; it must show that the rights it asserts are plausible, among the conditions for provisional measures. The evidentiary effort appropriate to the provisional-measures phase is therefore different in kind, and not merely in degree, from that required at the merits. A party that calibrates its evidentiary strategy to the function of the procedural stage in question—modest at the interlocutory phase, exacting at the merits—will deploy its resources to greater effect than one that treats the standard of proof as uniform across the life of the case.

The third implication concerns the transferability of the analysis beyond the ICJ. The rule on circumstantial evidence is not unique to the Court; the underlying

problem to which it responds—the asymmetry of access to proof where one party controls the relevant facts—arises across the institutions of international adjudication, including the dispute settlement system of the World Trade Organization, the International Tribunal for the Law of the Sea, and investor-State arbitration. The dual-track framework offers a lens through which the evidentiary practice of these other fora might be examined: in each, one may ask whether the conditions of access to inferential proof and the threshold of probative force have evolved together, and whether the relationship between them exhibits the compensating structure identified here in the jurisprudence of the ICJ. Whether the dual-track pattern is peculiar to the Court or is a more general feature of international evidentiary law is a question this article cannot answer, but it is one the present framework is designed to make askable.

These implications are, it must be stressed, methodological rather than substantive. They concern how parties should reason about evidence and how analysts should reason about evidentiary doctrine; they take no position on the outcome of any dispute. The value of the dual-track framework, in this domain as in the doctrinal analysis that precedes it, lies in the clarity it brings to the structure of inferential proof, not in any particular result it dictates.

7. Conclusion

The rule on circumstantial evidence announced in *Corfu Channel* has long been read as a single doctrine with a single function: the accommodation of applicants who cannot penetrate a respondent's monopoly over proof. This article has argued that the rule is better understood as having evolved along two analytically distinct axes. Along the horizontal axis of access, the rule has expanded outward, from the physical facts of *Corfu Channel*, through the clandestine attribution of *Nicaragua*, to the proof of State policy and genocidal intent in the *Genocide Convention* cases. Along the vertical axis of threshold, the rule has risen upward, from the "single conclusion" formula of *Corfu Channel*, through the stratified "fully conclusive" standard of *Bosnia v. Serbia*, to the "only reasonable inference" standard of *Croatia v. Serbia*. The recent practice, including the stratification of standards across procedural phases in *South Africa v. Israel*, confirms that the elevated threshold is now firmly established at the merits stage.

The central claim of the article is that these two trajectories are functionally connected. The expansion of access along the horizontal axis generated the risk of over-extension—the danger that grave conclusions would be drawn from equivocal patterns of conduct—and the elevation of the threshold along the vertical axis is the Court's compensating response to that risk. The breadth of the rule and the height of its threshold are mutually conditioning: the wider the rule reaches, the more it must demand. That claim rests, however, on a corrected account of where the standards come from. The elevation was not the internal ripening of the 1949 formula. Its standard-of-proof component descends from *Corfu Channel*, but its decisive inferential component—the requirement that the applicant's inference be

the only reasonable one—was imported from the jurisprudence of the ad hoc criminal tribunals, where it had developed as the counterpart of proof beyond reasonable doubt. That the Court, at the very moment when access reached its most consequential application, declined to elaborate its own formula and reached instead for a more demanding test from a neighbouring jurisprudence is the strongest evidence that the expansion was generating pressure for a constraint the existing formula could not supply. The borrowing is the compensating move.

The dual-track framework offers a fresh vantage point on the contemporary debate over the attainability of the "only reasonable inference" standard. That debate has largely examined the threshold in isolation, asking whether it is too high in the abstract. The framework developed here suggests that the height of the threshold cannot properly be assessed apart from the breadth of access it governs, and that the more adequate question is whether the threshold is proportionate to that breadth—whether the compensating constraint is calibrated to the risk the expansion of the rule has created. So reframed, the debate is relocated onto terrain on which principled argument can proceed. The corrected genealogy adds a further and sharper question: whether a test whose stringency derives from the need to protect individuals against wrongful criminal conviction is well fitted to proceedings in which one State accuses another of the gravest of internationally wrongful acts—a question about institutional fit that the abstract debate over the threshold's height cannot reach.

It remains to identify the limits of the present inquiry and the directions it leaves open. The analysis has been confined to the jurisprudence of the ICJ; whether the dual-track pattern is replicated in the evidentiary practice of other international courts and tribunals is a question for further study. The article has, moreover, deliberately abstained from the normative judgment of whether the Court has in fact struck the right balance between access and threshold; it has sought to clarify the structure within which that judgment must be made, not to make it. And the relationship between the circumstantial-evidence rule and adjacent features of the Court's evidentiary practice—the treatment of expert evidence, the effect of a party's non-appearance, the evidentiary weight of the reports of international bodies—invites the extension of the dual-track analysis to a wider field. These are tasks for another occasion. What this article has sought to establish is narrower but, it is hoped, secure: that the evidentiary doctrine before the International Court of Justice is the product not of one movement but of the convergence of several, and that understanding their relationship—the expansion that created the pressure, the borrowing that answered it, and the questions of institutional fit that the borrowing has left open—is the key to understanding the rule as it stands today.

Funding

This work was supported by the Fundamental Research Funds for the Central Universities (Grant No. 2026JX024).

Declaration

The author bears sole responsibility for the views expressed and for any errors that remain.

Conflicts of Interest

The author declares no conflicts of interest regarding the publication of this paper.

References

Ambos, K. (2016, April 1). *Karadzic's Genocidal Intent as the "Only Reasonable Inference"?* EJIL: Talk! https://www.ejiltalk.org/karadzics-genocidal-intent-as-the-only-reasonable-inference/

Barigye, N. Y., Hendrickse, M., & Todeschini, V. (2025, May 26). *Genocidal Intent in Armed Conflict: Unpacking the ICJ's "Only Reasonable Inference" Standard*. Opinio Juris. https://opiniojuris.org/2025/05/26/genocidal-intent-in-armed-conflict-unpacking-the-icjs-only-reasonable-inference-standard/

Cheng, B. (1953). *General Principles of Law as Applied by International Courts and Tribunals*. Stevens & Sons.

Del Mar, K. (2012). The International Court of Justice and Standards of Proof. In K. Bannelier, T. Christakis, & S. Heathcote (Eds.), *The ICJ and the Evolution of International Law: The Enduring Impact of the Corfu Channel Case* (pp. 98-123). Routledge.

Devaney, J. G. (2016). *Fact-Finding before the International Court of Justice*. Cambridge University Press. https://doi.org/10.1017/cbo9781316498972

Milanovic, M. (2025, September 18). *Proving Genocide*. EJIL: Talk! https://www.ejiltalk.org/proving-genocide/

Riddell, A., & Plant, B. (2009). *Evidence before the International Court of Justice*. British Institute of International and Comparative Law.

Sterio, M. (2017). The Karadžić Genocide Conviction: Inferences, Intent, and the Necessity to Redefine Genocide. *Emory International Law Review, 31,* 271-298.

Van Santen, B. (2025). Searching for a Plan: Demonstrating Genocidal Intent before the ICJ. *New York University Journal of International Law and Politics, 57,* 155-194.