THE EVIDENTIARY AND EPISTEMIC BREAKDOWN OF PATENTABILITY IN AUTONOMOUS MOLECULAR DESIGN: RECONCILING SECTION 3(D) AND THE SKILLED-PERSON STANDARD WITH BLACK-BOX IN-SILICO DRUG DISCOVERY

Siddharth Shankar, LLM, Vishwakarma University

ABSTRACT

Autonomous molecular design yields candidate compounds without articulable reasons, supported by efficacy evidence that is computed rather than observed. This article asks whether India's patentability architecture can absorb that shift: the skilled-person standard behind section 2(1)(ja) of the Patents Act, 1970, the therapeutic-efficacy barrier in section 3(d), and the procedural rules in section 104A and in section 63(4) of the Bharatiya Sakshya Adhiniyam, 2023. It surveys the doctrine from Bishwanath Prasad and Roche v. Cipla through Novartis and the Delhi High Court's 2025–2026 decisions, and identifies the lacunae. It argues that the breakdown is real but mislocated. The skilled person does not collapse, because computation makes candidate generation cheap without making efficacy predictable; inventive contribution migrates from structure to evidence. Section 3(d) holds, but generative optimisation targets properties that it discounts. The procedural layer fails: section 104A presupposes concealed physical manufacture, and section 63(4) authenticates files, not models. The article proposes clarifications to sections 2(1)(ja), 3(d) and 6, an adverse-inference supplement to section 104A, and a short, conditional, access-safeguarded provisional exclusivity tied to auditable provenance records.

Keywords: autonomous molecular design; section 3(d); person skilled in the art; therapeutic efficacy; section 104A; Bharatiya Sakshya Adhiniyam; provisional exclusivity

1. Introduction

Patent law must decide what counts as evidence that an invention is an invention. Autonomous molecular design unsettles the answer in two ways at once. The route to a candidate compound is a trajectory through a learned latent space that neither applicant nor examiner can reconstruct as a chain of reasons. And the evidence of the candidate's merit is typically computed, not observed: docking scores, predicted affinities and in-silico ADMET profiles. Indian pharmaceutical patent law is unusually exposed to both shifts, because section 3(d) of the Patents Act, 1970, as read by the Supreme Court, demands proof of therapeutic efficacy and treats weaker proxies with suspicion.[1]

This article asks whether three pillars of that regime survive black-box design: the skilled-person fiction through which inventive step is judged, the efficacy barrier of section 3(d), and the procedural machinery of section 104A of the Patents Act and section 63 of the Bharatiya Sakshya Adhiniyam, 2023 (BSA). It argues that the breakdown is real but mislocated. The skilled person does not collapse, because computation makes candidate generation cheap without making efficacy predictable; inventive contribution migrates from structure to evidence. Section 3(d) holds, but generative optimisation targets properties that section 3(d) discounts. The procedural layer is where the system genuinely fails, because section 104A presupposes concealed physical manufacture and section 63(4) authenticates files rather than validating models. Part 7 proposes a measured response: clarifications to sections 2(1)(ja), 3(d) and 6, and a short, conditional, access-safeguarded provisional exclusivity tied to auditable provenance records.

Purely computational method claims, which engage section 3(k), are outside the scope of this article. Technical terms are explained in Part 2.1.[2]

2. Technical Foundations

2.1 Definitions

Autonomous molecular design means software that invents possible drug molecules on its own.

[1] *Novartis AG v. Union of India*, (2013) 6 S.C.C. 1 (India) [hereinafter *Novartis*].

[2] Patents Act, 1970, No. 39 of 1970, § 3(k) (India) [hereinafter Patents Act]. *See also* S.S. Rana & Co., *India's Draft Biotechnology Patent Guidelines (2026): What Pharma & Biotech Brand Owners Need to Know*, S.S. RANA & CO. (Sept. 17, 2026), https://ssrana.in/articles/indias-draft-biotechnology-patent-guidelines-2026-what-pharma-biotech-brand-owners-need-to-know/ (reporting that the draft treats purely computational methods lacking technical effect as falling within § 3(k)).

Scientists state what they want, such as a molecule that fits a disease-related protein, and the software proposes, scores and ranks structures, repeating the cycle many times, with no human choosing each one. It learns patterns from large collections of known molecules and composes new ones that fit them (the common techniques are called diffusion models and equivariant neural networks). What it has learned is stored as numbers in a *latent space*, an internal map of chemical possibilities. The software can produce a molecule from that map but cannot give a step-by-step reason for choosing it; that is the 'black box' discussed in this article. Its instructions, or *conditioning inputs*, typically include the shape of the target protein's *binding site* (the 'lock' the molecule must fit), a *seed scaffold* (a starting framework, often taken from a known drug) and the properties wanted.

A *candidate compound* is a molecule that the software proposes for further testing as a possible medicine, at a stage when nobody knows whether it will work in patients. This article uses the term for the software's output at any stage, from an early laboratory hit to a molecule chosen for human trials.

Scaffold hopping redesigns a molecule's core framework while keeping the features that let it act on the same target, like cutting a differently shaped key that opens the same lock. *Bioisosteric replacement* swaps a small part of a molecule for a similarly behaving part, such as one atom for another of similar size. *Deuteration* swaps hydrogen atoms for a heavier, chemically similar form of hydrogen, usually so that the body breaks the drug down more slowly.

Evidence about a molecule is either computed or measured. Work by computer simulation is *in silico*; work in laboratory dishes is *in vitro*; work in living animals or people is *in vivo*. Laboratory testing is called *wet-lab validation*. *Docking* simulates on screen how well a molecule fits a target's binding site, and the *docking score* is the simulation's numerical estimate of that fit: a ranking aid, not a measurement. *Binding affinity* is the experimentally measured strength with which a molecule attaches to its target. *ADMET* stands for absorption, distribution, metabolism, excretion and toxicity: how a drug enters the body, where it goes, how the body alters and removes it, and whether it causes harm. A *computational ADMET profile* is software's prediction of those properties, not a laboratory result. *Bioavailability*, the share of a dose that reaches the bloodstream, is one of them. Such properties show whether a molecule can behave as a usable medicine, not that it treats the disease. *Therapeutic efficacy*,

the ability to produce the intended clinical benefit in patients, is what section 3(d) tests; *potency*, the amount of substance needed to produce an effect in a laboratory test, is different. In clinical trials, Phase I mainly tests safety and dosing, and Phase II is where effectiveness is first tested in patients.

A model's reliability depends on how it was tested. A *held-out set* is data kept back from training and used as an examination; the examination is fairer if it uses molecules disclosed after training ended or molecules unlike those the model learned from. The *applicability domain* is the range of molecules for which predictions can be trusted; outside it, they are guesses. A model is *calibrated* when its stated confidence matches how often it is right, and *data leakage* occurs when examination answers have crept into the training data. Finally, *model weights* are the numbers that make up a trained model, a *training-data manifest* is an inventory of the data it learned from, and a *hash value* is a short digital fingerprint of a file that changes if any part of the file changes.

2.2 Three evidential properties

Three properties of autonomous design matter evidentially. The first is opacity: the software maps its instructions (a target, a starting structure, desired properties) to candidate molecules without a human-readable rationale, and after-the-fact explanations describe how the model behaves, not why this molecule was chosen. The second is surrogacy: the objectives optimised are computable proxies, such as docking scores, predicted binding strength and predicted clearance or toxicity, for the clinical result that the law cares about. The third is equifinality: independently built systems can arrive at the same molecule, so an identical product does not imply an identical process.

The proxies are imperfect in ways that courts will meet. A benchmark of deep-learning docking methods, including diffusion-based ones, found that they frequently predict physically implausible ways in which a molecule sits in its target and generalise poorly to new targets, with none outperforming classical docking tools.[3] The clinical data show a similar divide. Molecules from AI-native developers cleared Phase I at 80–90 percent against historical rates of 40–65 percent, but Phase II success, where efficacy is first tested, was about 40 percent,

[3] Martin Buttenschoen, Garrett M. Morris & Charlotte M. Deane, *PoseBusters: AI-Based Docking Methods Fail to Generate Physically Valid Poses or Generalise to Novel Sequences*, 15 CHEM. SCI. 3130 (2024).

close to conventional discovery, on small numbers.[4] Computation has compressed design; it has not compressed uncertainty about efficacy. That asymmetry organises what follows.

3. The Skilled Person

3.1 A fiction with elastic edges

Section 2(1)(ja) defines inventive step by reference to a technical advance or economic significance that makes the invention not obvious to a person skilled in the art, but the Act nowhere defines that person.[5] The figure is judge-made. *Bishwanath Prasad* requires more than a workshop improvement,[6] and the Division Bench in *Roche v. Cipla* prescribed five steps: identify the ordinary skilled person; identify the inventive concept; impute to a skilled but unimaginative person the common general knowledge at the priority date; identify the differences; and decide, excluding hindsight, whether those differences would have been obvious.[7] The same judgment tied the fiction to disciplinary composition by holding that Cipla's expert, who was not a medicinal chemist, could not speak as the skilled person.[8] The Controller General's 2026 draft guidelines fold this test together with the Delhi High Court's four-step formulation while cautioning that neither is a rigid formula.[9]

The fiction is also already elastic. In genus-species disputes between common inventors, the Delhi High Court has assessed obviousness through a 'person in the know', who knows what the applicant knew. The *Risdiplam* Division Bench applied the idea at the interim stage in October 2025, although the earlier dapagliflozin judgments that introduced it were set aside as precedent by the Supreme Court.[10] Commentators ask whether the external, hypothetical

[4] Madura K.P. Jayatunga et al., *How Successful Are AI-Discovered Drugs in Clinical Trials? A First Analysis and Emerging Lessons*, 29 DRUG DISCOVERY TODAY 104009 (2024).
[5] Patents Act, *supra* note 2, § 2(1)(ja).
[6] *Bishwanath Prasad Radhey Shyam v. Hindustan Metal Indus.*, (1979) 2 S.C.C. 511 (India).
[7] *F. Hoffmann-La Roche Ltd. v. Cipla Ltd.*, RFA(OS) 92/2012 & RFA(OS) 103/2012, ¶ 151 (Delhi H.C. Nov. 27, 2015) (India) (Div. Bench) [hereinafter *Roche v. Cipla*].
[8] *Id.* ¶ 168.
[9] Office of the Controller Gen. of Patents, Designs & Trade Marks, *Draft Guidelines for Examination of Pharmaceutical Patent Applications* (Sept. 4, 2026) (India) [hereinafter Draft Guidelines]; *see* Gaurav Mishra, *Draft Pharma Patent Examination Guidelines 2026*, INTELLEPEDIA (Sept. 7, 2026), https://www.bananaip.com/intellepedia/draft-pharma-patent-examination-guidelines-2026/.
[10] *F. Hoffmann-La Roche AG v. Natco Pharma Ltd.*, FAO(OS)(COMM) 43/2025 (Delhi H.C. Oct. 9, 2025) (India) (Div. Bench) [hereinafter *Risdiplam*]; Victor Vaibhav Tandon, *Genus-Specie Disputes: Thoughts on the Risdiplam Division Bench Judgment (Part I)*, SPICYIP (Oct. 2025), https://spicyip.com/2025/10/part-i-genus-specie-disputes-thoughts-on-the-risdiplam-division-bench-judgment.html (noting the Supreme Court's order of Aug. 2, 2024).

construct is being converted into a real-inventor yardstick.[11] In *Intra-Cellular*, which treated coverage by a genus patent as disclosure, one commentator notes that the court never named the standard it was in substance applying.[12]

3.2 The augmented skilled person

Abbott argues that once inventive machines become standard in a field, the skilled person should be understood as one who uses such a machine, or as the machine itself, so that the bar to patentability rises until inventive activity becomes obvious.[13] Dornis foresees a systemic meltdown of patent doctrine absent redesign.[14] The Indian formulation sharpens the worry. The skilled person is defined by the absence of imagination, and running a generally available generative model is mechanical, not imaginative. On that logic, every compound that a standard tool would output lies within the unimaginative person's reach.

3.3 Reassessment: try-ability and expectation

The argument proves less than it appears to. Obviousness in the chemical arts has two components: whether the claimed compound would have been envisaged or tried, and whether the skilled person would have had reason to expect it to succeed for the claimed purpose. Tools transform the first. Enumeration and scoring costs approach zero, so the set of obvious-to-try structures expands. They leave the second largely untouched as to efficacy. The Division Bench in *Roche v. Cipla* held that structural similarity alone does not show obviousness; a reason to make the particular change in order to achieve the sought property must be shown.[15] The Phase II data in Part 2 are the empirical counterpart: models do not give the skilled person a reasonable expectation of therapeutic efficacy. The tool-augmented skilled person can generate and rank candidates but cannot expect efficacy. Inventive contribution therefore migrates from

[11] Victor Vaibhav Tandon, *Genus-Specie Disputes: Thoughts on the Risdiplam Division Bench Judgment (Part II)*, SPICYIP (Oct. 2025), https://spicyip.com/2025/10/part-ii-genus-specie-disputes-thoughts-on-the-risdiplam-division-bench-judgment.html.

[12] *Intra-Cellular Therapies, Inc. v. Controller of Patents*, C.A.(COMM.IPD-PAT) 24/2023, 2026:DHC:5394 (Delhi H.C. July 6, 2026) (India) [hereinafter *Intra-Cellular*]; Arnav Kaman, *One Atomic Substitution, Two Giant Hurdles of Disclosure and Efficacy*, SPICYIP (Aug. 17, 2026), https://spicyip.com/2026/08/one-atomic-substitution-two-giant-hurdles-of-disclosure-and-efficacy.html.

[13] Ryan Abbott, *Everything Is Obvious*, 66 UCLA L. REV. 2 (2019).

[14] Tim W. Dornis, *Artificial Intelligence and Innovation: The End of Patent Law as We Know It*, 23 YALE J.L. & TECH. 97 (2020).

[15] *Roche v. Cipla*, *supra* note 7; *see* Rupali Samuel, *Cipla's Erlocip Infringes Roche's Patent, Finds DHC*, SPICYIP (Dec. 1, 2015), https://spicyip.com/2015/12/ciplas-erlocip-infringes-roches-patent-finds-dhc.html.

the compound to its validation and to selection under uncertainty.

Three rules follow. First, common general knowledge should include tools in general use at the priority date, proved by evidence of adoption rather than mere availability. Second, the expectation-of-success inquiry should be endpoint-specific: a model validated for solubility raises obviousness for solubility only. Third, the person in the know should not be extended to a platform operator's private models. Treating unpublished latent space as constructive prior art would reward secrecy and deter disclosure, and it is unnecessary, because an applicant's own seed scaffold is already available as an admission in the section 3(d) inquiry, as Part 4 explains.

4. Section 3(d) and the Efficacy Barrier

4.1 Doctrine

Section 3(d) excludes the mere discovery of a new form of a known substance that does not enhance known efficacy, and deems salts, esters, polymorphs, isomers and 'other derivatives' to be the same substance unless they differ significantly in properties with regard to efficacy.[16] *Novartis* held that, for medicines, efficacy means therapeutic efficacy, and that increased bioavailability does not establish it unless research data show that it produces an enhancement.[17] Four later developments matter. The *Roche v. Cipla* Division Bench treated section 3 as a threshold of eligibility rather than an exception to section 2(1)(j), and described a new chemical entity that is structurally dissimilar but functionally similar to an existing one as merely a substance under section 3(d), to be assessed for inventive step only if enhanced efficacy is shown.[18] *DS Biopharma* and *Taiho* require the Controller to identify a specific known substance, even where the prior art is a Markush claim.[19] *Intra-Cellular* held that deuterated analogues of a prior-art compound were new forms of a known substance, and that pharmacokinetic improvement without research data linking it to therapeutic benefit did not suffice.[20] And the 2026 draft guidelines describe the inquiry as data-driven, with roughly

[16] Patents Act, *supra* note 2, § 3(d) & Explanation.
[17] *Novartis*, *supra* note 1.
[18] *Roche v. Cipla*, *supra* note 7, ¶ 61; *see* Rupali Samuel, *3(d) a Patent Eligibility Provision: DHC in Roche v. Cipla*, SPICYIP (Dec. 1, 2015), https://spicyip.com/2015/12/3d-a-patent-eligibility-provision-dhc-in-roche-v-cipla.html.
[19] *DS Biopharma Ltd. v. Controller of Patents & Designs*, C.A.(COMM.IPD-PAT) 6/2021 (Delhi H.C. Aug. 30, 2022) (India); *Taiho Pharm. Co. v. Controller of Patents*, C.A.(COMM.IPD-PAT) 6/2022 (Delhi H.C. May 15, 2025) (India).
[20] *Intra-Cellular*, *supra* note 12.

seventeen worked section 3(d) examples, deuterated analogues included.[21]

4.2 Scaffold hopping and bioisosteric generation

On the *Roche v. Cipla* reading, scaffold hopping, which by design changes the core framework while preserving the features that let the molecule act on its target, produces precisely the structurally dissimilar but functionally similar entity that section 3(d) treats as a substance, not an invention. The question whether AI-driven optimisation is an unpatentable new form or a qualifying contribution is therefore mis-posed. Section 3(d) is a threshold that precedes inventive step, so structural novelty, however algorithmically produced, cannot bypass it. The compound must show significantly enhanced therapeutic efficacy over an identified known substance.

Two frictions follow. The first is a mismatch of objectives. Generative design optimises several goals at once: acting on the intended target and not others, resisting rapid breakdown in the body, solubility, low predicted toxicity and ease of manufacture. *Novartis* and *Intra-Cellular* treat physico-chemical and pharmacokinetic gains as something other than therapeutic efficacy, and whether reduced toxicity counts is unresolved.[22] The comparative advantage of computational design therefore lies where section 3(d) is least receptive. The second friction is the identification of the comparator. *DS Biopharma* requires a specific known substance, but for a generated analogue the nearest known compound by computed structural similarity need not be the scaffold from which the model started. The applicant's provenance record names the seed and so supplies the comparator, which makes honest disclosure an admission against interest. That incentive to withhold is the very opacity that the system needs to dissolve.

4.3 Evidentiary sufficiency

Can simulated binding affinities, docking scores or computational ADMET profiles establish a prima facie case of enhanced therapeutic efficacy without wet-lab validation at filing? On present law, no. *Novartis* required research data even for bioavailability, a measured pharmacokinetic quantity, and *Intra-Cellular* held measured exposure insufficient without a demonstrated therapeutic link. A predicted score sits further along an inferential chain: a

[21] Draft Guidelines, *supra* note 9; Mishra, *supra* note 9.
[22] *Novartis*, *supra* note 1 (rejecting improved physico-chemical properties and bioavailability as such); *Intra-Cellular*, *supra* note 12.

predicted fit, then actual binding, then action on the target in a living body, then benefit to patients. Each link is empirical, and the benchmark and clinical data in Part 2 show that the links are weak.

That answer is 'not alone', not 'never'. A computed value from a model that was validated in advance for the specific property, on molecules disclosed after its training or unlike its training data, and that states reliable confidence levels, is relevant circumstantial evidence that an expert can explain, and it can support a plausible efficacy hypothesis at filing. Whether later wet-lab data may then complete the showing is a timing question that the Act does not answer. The European Patent Office's treatment of post-published evidence offers one model,[23] and Indian law would benefit from express rule-making on post-filing efficacy data in the section 3(d) context.

5. Procedural Burden and the Black Box

5.1 Section 104A

Section 104A allows a court, in an infringement suit on a process patent, to direct the defendant to prove that its process differs where the patent is for obtaining a new product, or where an identical product is substantially likely to have been made by the patented process and the patentee could not, by reasonable efforts, find out the process used. The patentee must first prove that the product is identical to the product directly obtained by the patented process, and the defendant need not disclose manufacturing or commercial secrets where that would be unreasonable.[24] The Delhi High Court has treated the provision as a calibrated departure from the ordinary burden, held that product identity, not similarity, is a sine qua non, and held that the provision prevails over general discovery rules.[25]

Three features of autonomous design make section 104A a poor fit. First, the product directly obtained by a computational pipeline is information, a ranked list of structures, while the physical compound is made by ordinary synthesis. Whether data is a product for the purposes of section 48(b) is doubtful, and if the claimed compound is itself patentable the patentee has

[23] Enlarged Bd. of Appeal, Eur. Patent Off., Decision G 2/21 (Mar. 23, 2023).
[24] Patents Act, *supra* note 2, § 104A.
[25] *F. Hoffmann-La Roche AG v. Zydus Lifesciences Ltd.*, I.A. 5827/2024 in CS(COMM) 159/2024 (Delhi H.C. 2025) (India) [hereinafter *Roche v. Zydus*]; *see* Rohin Pujari, *India – Delhi High Court Denies Roche Access to Zydus' Confidential Information*, CONVENTUS LAW (Aug. 20, 2025), https://conventuslaw.com/report/india-delhi-high-court-denies-roche-access-to-zydus-confidential-information/.

no need of a process claim.[26] Second, equifinality breaks the inference on which both limbs rest. A new compound is not obtainable only by the patented pipeline, and identity of product says little about identity of process when independent models converge. Third, the defendant's rebuttal cannot be mechanistic. Because the model is non-interpretable, difference can be shown only through provenance: architecture, training-data manifests, weight hashes and logs. That is exactly what section 104A(2) protects as a secret. Applied mechanically, the provision will either impose a burden that cannot be discharged or never trigger.

5.2 Electronic records under the BSA

Section 63 of the BSA replaced section 65B of the Indian Evidence Act, 1872. The certificate under section 63(4) has two parts: Part A discloses the hash value of the record, and Part B is certified by an expert.[27] In *Pune Bar Association* (May 2026) the Supreme Court upheld the requirement and described the hash value as an electronic fingerprint. It indicated that, besides a notified Examiner of Electronic Evidence, a person with special skill in computer science and cyber forensics may sign Part B if the court is satisfied of that expertise, while declining to rule conclusively on the point.[28]

For black-box evidence the regime does less than it appears to. It addresses integrity: the file is the file that was hashed. It does not address validity, whether the model's outputs are reliable for the endpoint, or completeness, whether the logs record the process that produced the compound. Those questions fall to expert opinion under section 39, and Indian courts expect expert opinion to be reasoned.[29] A neural network cannot supply reasons, but a human expert can supply reasons about validation methodology, applicability domain and data leakage. Courts will need protocols to receive weights and logs in confidence, of the kind used for the confidentiality club in *Roche v. Zydus*, and to separate certification of authenticity from evaluation of reliability.

5.3 Provenance as the common denominator

Both failures point to the same missing object: an auditable, certified provenance record. Without it neither burden-shifting nor authentication can do epistemic work. With it, the

[26] Patents Act, *supra* note 2, § 48(b).
[27] Bharatiya Sakshya Adhiniyam, 2023, No. 47 of 2023, § 63(4) & sched. (India) [hereinafter BSA].
[28] *Pune Bar Ass'n v. Union of India*, 2026 SCC OnLine SC 1297 (India) (decided May 22, 2026).
[29] BSA, *supra* note 27, § 39; *State of Himachal Pradesh v. Jai Lal*, (1999) 7 S.C.C. 280 (India).

comparator in section 3(d), the rebuttal in section 104A and the certificate in section 63 all acquire content. Provenance is also double-edged. It names the known substance and exposes weights and data, and so cannot responsibly be demanded without protection for secrets and a corresponding benefit to the discloser.

6. Lacunae

The survey exposes six gaps. No Indian authority addresses whether the skilled person is taken to use generally available computational tools. No authority addresses whether computed values may count as efficacy evidence under section 3(d), or whether post-filing data may complete the showing. The comparator problem for generated analogues is unaddressed. Section 104A has not been tested on computational processes, and its premises fail for them. The scope of the expert who may certify Part B of the BSA certificate remains open after *Pune Bar Association*. And there is no framework for provenance records, so that the evidence on which every other inquiry depends has no standard form, custodian or confidentiality regime.

7. Doctrinal Alternatives and Proposal

7.1 Structural amendments

Three targeted amendments would address most of the doctrinal friction. To section 2(1)(ja), an explanation should provide that the person skilled in the art is taken to have access to computational tools in general use in the field at the priority date, but not to undisclosed models or their outputs. To section 3(d), an explanation should provide that efficacy shall be established by experimental or clinical data; that predictions from a computational model may corroborate, but not alone establish, efficacy where the model is shown to be validated for the endpoint; and, codifying *DS Biopharma*, that the Controller shall identify the known substance. As to section 6, no amendment is needed to the requirement that the applicant be the true and first inventor. A rule requiring a declaration of computational assistance and of the natural person's act of selection or validation would align Indian practice with the United States Patent and Trademark Office's November 2025 guidance, under which only natural persons can be inventors and conception remains the test.[30]

[30] Patents Act, *supra* note 2, §§ 2(1)(y), 6(1)(a); Revised Inventorship Guidance for AI-Assisted Inventions, 90 Fed. Reg. 54,636 (Nov. 28, 2025).

These amendments cannot do everything. If section 3(d) requires efficacy data at filing, applicants who have computed but not yet validated a candidate must choose between filing prematurely and withholding the compound as a trade secret, which defeats disclosure. Loosening section 3(d) to accommodate computed evidence would reopen the evergreening that the provision exists to prevent, and the access stakes are concrete: in the *Risdiplam* litigation a generic launch was permitted at about ₹15,900 per bottle against a reported originator price of ₹6 lakh.[31]

7.2 A provisional computational exclusivity

A sui generis instrument can absorb the timing problem without touching section 3(d). I propose a provisional computational exclusivity for compounds designed with material use of an autonomous design system that cannot yet meet section 3(d) for want of efficacy data. Registration would be conditional on depositing a provenance record, certified under section 63(4) of the BSA and held in confidence, comprising conditioning inputs, model identifiers and weight hashes, training-data manifest hashes, generation logs and a validation statement. The term would be short and non-renewable, and would lapse unless within it the registrant files a patent application supported by efficacy data that satisfies section 3(d), in which case the exclusivity would merge into, and not extend, the patent term. Remedies would be limited to a reasonable royalty, with no injunction against products already on the market, and the instrument would be expressly subject to sections 84 and 92 and to the public-health flexibilities confirmed at Doha.[32] A model provision follows:

> (1) A person who has designed a compound with material use of an autonomous molecular design system may apply to register a provisional computational exclusivity on depositing a provenance record certified under section 63(4) of the Bharatiya Sakshya Adhiniyam, 2023.
>
> (2) The exclusivity shall subsist for three years from registration, shall not be renewed or extended, and shall lapse unless within that period the registrant files a patent application supported by efficacy data satisfying section 3(d) of this

[31] *Delhi High Court's Ruling in F. Hoffmann-La Roche AG & Anr. v. Natco Pharma*, INDIA CORPORATE LAW (Nov. 2025), https://corporate.cyrilamarchandblogs.com/2025/11/delhi-high-courts-ruling-in-f-hoffmann-la-roche-ag-anr-v-natco-pharma/; *Risdiplam*, *supra* note 10.

[32] World Trade Org., *Declaration on the TRIPS Agreement and Public Health*, WT/MIN(01)/DEC/2 (Nov. 14, 2001); Patents Act, *supra* note 2, §§ 84, 92.

Act.

(3) The registrant's remedy for infringement shall be limited to a reasonable royalty. No injunction shall issue against a product marketed before registration or where the public interest so requires.

(4) Nothing in this section limits sections 84 and 92 of this Act.

(5) Where the provenance record is in issue, the court shall proceed in confidence and may draw an adverse inference from a failure to produce it.

7.3 Section 104A and objections

For computational processes, section 104A should be supplemented by a rule that, once the patentee makes a prima facie case that the defendant practises the patented process, a defendant who fails to produce a certified provenance record may be taken not to have shown that its process differs, with the protection of sub-section (2) preserved through confidential proceedings. That locates the burden with the one party able to meet it without forcing disclosure of weights to a rival.

The exclusivity may be objected to as evergreening by another name. Its lapse condition, royalty-only remedy and exclusion of products already marketed answer that charge: it protects a candidate during the interval in which efficacy is being established, and gives nothing once efficacy fails. It may also be objected that the provenance deposit burdens applicants. That burden is the price of the exclusivity, and an applicant who prefers secrecy remains free to rely on trade secret protection and forgo the instrument.

8. Conclusion

Autonomous molecular design does not abolish the skilled person or defeat section 3(d). It displaces the locus of invention from structure to evidence, and it exposes a procedural layer built for a different kind of concealment. The skilled person remains serviceable if made tool-aware. Section 3(d) remains the right gate, but its evidentiary demands and its timing must be made explicit. Sections 104A and 63(4) cannot, without a shared provenance record, do the epistemic work asked of them. The proposals here are deliberately modest: three clarifications, an adverse-inference supplement, and a provisional exclusivity that lapses unless efficacy is

proved. They keep faith with *Novartis*, whose logic is that monopoly should follow demonstrated therapeutic benefit, while recognising that demonstration now arrives in two stages, computed and observed. The analysis rests on first-instance and interim decisions, draft guidelines and a small clinical evidence base, and should be revisited as these settle.