
ALGORITHMIC CONSUMPTION AND ECOLOGICAL COST: RE-EXAMINING CONSUMER PROTECTION IN THE AGE OF AI-DRIVEN COMMERCE

Adv. Arathi Menon P R, High Court of Kerala

ABSTRACT

The proliferation of Artificial Intelligence in e-commerce, recommendation engines, and personalised advertising has fundamentally altered consumption patterns, often at significant and underexamined environmental cost. AI-driven “nudge” architectures, hyper-personalised recommendations, dynamic pricing, and algorithmically generated urgency cues such as flash sales, actively incentivise overconsumption, fuelling fast fashion cycles, packaging waste, and the rapid obsolescence of electronic goods, while the computational infrastructure underlying these systems (data centres, model training, cloud storage) carries its own substantial carbon and water footprint. This paper interrogates whether India’s existing consumer protection architecture, particularly the Consumer Protection Act, 2019 and the Consumer Protection (E-Commerce) Rules, 2020, is equipped to address these AI-specific environmental externalities. It examines the adequacy of provisions on unfair trade practices and misleading advertisements in capturing “dark pattern” design choices that drive ecologically harmful consumption, and considers whether algorithmic transparency can be read as an extension of the consumer’s right to know. The paper further explores comparative regulatory responses, including the EU’s Digital Services Act and emerging AI Act provisions on manipulative design, to assess transposable safeguards for the Indian context. It argues for a recalibrated enforcement framework under the Central Consumer Protection Authority that integrates environmental disclosure obligations, covering both the ecological footprint of AI-recommended products and the energy intensity of the recommending technology itself, into existing consumer-facing accountability mechanisms. The paper concludes by proposing a framework of “algorithmic eco-disclosure” as a necessary evolution of eco-labelling norms for the digital marketplace, situating consumer protection law as an active instrument of environmental governance rather than a passive bystander to technologically mediated consumption.

Keywords: Artificial Intelligence, Dark Patterns, E-Commerce, Consumer Protection Act 2019, Environmental Disclosure, Algorithmic Accountability

I. INTRODUCTION

The Indian e-commerce ecosystem has, within the space of a decade, migrated from static product catalogues to dynamic, machine-learned storefronts that adapt to each visitor in real time. Recommendation engines decide which products surface first; dynamic pricing engines decide what each visitor pays; and behavioural design choices, increasingly informed by predictive models of attention and willingness to pay, decide how long a visitor lingers and how much they ultimately purchase. This transformation is generally narrated in the language of efficiency and consumer convenience. What this narrative omits is the ecological ledger behind the screen: every additional purchase nudged into being by an algorithm carries embedded environmental costs in manufacturing, packaging, transportation, and end-of-life disposal, while the recommendation itself is the product of computational infrastructure that consumes electricity and water at a scale that is no longer trivial. Artificial intelligence in commerce is therefore not environmentally neutral merely because it is intangible; it externalises ecological cost twice over, once through the consumption it manufactures, and once through the computation it requires.

India's consumer protection regime has begun, tentatively, to confront the first of these externalities. The Consumer Protection Act, 2019 ("the Act") and the delegated rules and guidelines issued under it, principally the Consumer Protection (E-Commerce) Rules, 2020 and the Guidelines for Prevention and Regulation of Dark Patterns, 2023, recognise that algorithmically engineered choice architecture can amount to an unfair trade practice or a misleading advertisement. What this framework has not yet done, and arguably has not yet been asked to do, is treat the ecological consequence of that choice architecture as a cognisable harm in its own right, or extend the consumer's right to be informed to cover the energy and resource intensity of the technology doing the recommending. This paper undertakes that extension. It asks two connected questions: first, whether India's existing apparatus of unfair trade practice and misleading advertisement law is doctrinally capacious enough to capture AI-driven environmental harm; and second, assuming a degree of present inadequacy, what an enforcement framework built around "algorithmic eco-disclosure" might look like, drawing on comparative experience under the European Union's Digital Services Act and Artificial Intelligence Act, and on India's own under-utilised eco-labelling infrastructure.

The argument proceeds in six further parts. Part II maps the problem, distinguishing the behavioural channel through which AI drives ecologically harmful consumption from the infrastructural channel through which AI itself consumes ecological resources. Part III audits the adequacy of the Act, the E-Commerce Rules, and the Dark Patterns Guidelines against this dual problem. Part IV argues that algorithmic transparency is a natural extension of the consumer's existing statutory right to be informed. Part V turns comparatively to the European Union. Part VI proposes the algorithmic eco-disclosure framework and locates it within the Central Consumer Protection Authority's existing guideline-making power. Part VII concludes.

II. THE ALGORITHMIC ARCHITECTURE OF ECOLOGICALLY HARMFUL CONSUMPTION

A. Nudge Design and the Manufacture of Demand

Recommendation engines do not merely respond to demand; they actively manufacture it. Personalisation systems trained on browsing and purchase history identify the precise threshold of novelty, scarcity, and social proof required to convert a passive browser into a purchaser, and platforms deploy this capability through devices such as countdown timers, stock-scarcity indicators, and “frequently bought together” prompts that compress the time available for deliberation. The cumulative effect on categories such as apparel is well documented: empirical literature on fast fashion records a marked decline in the average number of times a garment is worn before disposal, even as overall clothing purchase volumes have risen sharply, with the United Kingdom alone recording a fifty per cent increase in clothing purchases by weight between 2012 and 2022.¹ Scholars studying digital resale platforms have observed that personalised recommendation and frictionless returns blur the behavioural distinction between primary and secondary retail, such that the apparatus of recommendation reproduces overconsumption patterns even within markets framed as sustainable.² Fashion retailers operating largely online have been reported to deploy algorithmic trend-detection to release hundreds of new styles in rapid

¹Economics Observatory, “How does fast fashion affect the environment?”, available at: <https://www.economicsobservatory.com/how-does-fast-fashion-affect-the-environment> (Visited on June 19, 2026).

²“Secondhand fashion consumers exhibit fast fashion behaviors despite sustainability narratives”, Scientific Reports (2025), available at: <https://www.nature.com/articles/s41598-025-19089-1> (Visited on June 19, 2026).

succession, a production cadence that would be commercially unsustainable without algorithmic demand prediction.³

A second consumption channel runs through electronic goods themselves. Dynamic pricing and AI-curated “upgrade” prompts compress the replacement cycle for smartphones, laptops, and small appliances, normalising functional obsolescence well before mechanical failure. Because electronic waste is among the fastest-growing waste streams globally, and Indian electronics consumption is increasingly mediated by platforms employing exactly this design logic, algorithmic upselling translates directly into downstream e-waste, a harm the consumer protection framework has, to date, treated as incidental rather than central to the regulatory inquiry.

A third channel operates through packaging and logistics. Returns, themselves frequently induced by aggressive recommendation and frictionless ordering defaults, generate disproportionate packaging waste and reverse-logistics emissions; one estimate places returns at up to half of online fashion orders in parts of Europe, a phenomenon increasingly mirrored in Indian quick-commerce models built around rapid, low-friction ordering.⁴

B. The Material Infrastructure Beneath the Interface

The second channel is less visible precisely because it is upstream of the consumer interface: the computational cost of producing the recommendation itself. Training and running the machine learning models that power recommendation, dynamic pricing, and conversational shopping assistants requires data centre capacity whose electricity demand has grown explosively. The International Energy Agency has estimated that global data centre electricity consumption stood at approximately 415 terawatt-hours in 2024 and projected this to rise to nearly 945 terawatt-hours by 2030, with artificial intelligence workloads accounting for the principal share of this growth, taking data centres’ share of global electricity demand from roughly 1.5 per cent towards 3 per cent within the decade.⁵ Independent estimates derived from comparable data suggest that

³See, generally, “Environmental and Humanitarian Impacts of Fast Fashion”, Fordham University Library Research Collection, available at: https://research.library.fordham.edu/cgi/viewcontent.cgi?article=1182&context=environ_2015 (Visited on June 19, 2026).

⁴“Recommender Systems for Social Good: The Role of Accountability and Sustainability”, arXiv:2501.05964 (2025).

⁵International Energy Agency, Energy and AI (April, 2025), data summarised in “AI Data Center Energy Consumption Statistics 2026”, available at: <https://aimultiple.com/ai-energy-consumption> (Visited on June 19, 2026).

AI systems alone could be responsible for between 32.6 and 79.7 million tonnes of carbon dioxide emissions in 2025, and for a water footprint of between 312.5 and 764.6 billion litres in the same year, water being consumed both directly for cooling and indirectly through the generation of the electricity consumed.⁶ A 2026 assessment by United Nations researchers projects that, by 2030, AI-attributable water consumption could reach 9.3 trillion litres globally, alongside an AI-attributable e-waste burden estimated at 2.5 million metric tonnes annually by the end of the decade.⁷

None of this infrastructure is geographically or legally confined to the jurisdiction in which the resulting recommendation is displayed. An Indian consumer scrolling a marketplace application on a domestic network may be served a personalised product list generated by inference running on servers in an altogether different country, drawing on an electricity grid of a different carbon intensity, and consuming water in a location entirely disconnected from the transaction. The opacity is therefore double: the consumer cannot see the manipulative design choice for what it is, and even if they could, they would have no means of knowing the energy or water cost of the underlying computation that produced it. It is this second, infrastructural opacity that existing Indian consumer protection law, oriented as it is towards the point of sale, has not yet been asked to address.

III. INDIA'S CONSUMER PROTECTION ARCHITECTURE: AN ADEQUACY ASSESSMENT

A. Unfair Trade Practice and Misleading Advertisement under the Act of 2019

The Consumer Protection Act, 2019 replaced the narrower 1986 statute with a considerably wider conception of regulable misconduct. Section 2(47) defines “unfair trade practice” expansively, to capture any unfair method or unfair or deceptive practice adopted for the purpose of promoting the sale of goods or the provision of services, including false representation as to the

⁶“The carbon and water footprints of data centers and what this could mean for artificial intelligence”, ScienceDirect (December, 2025), available at: <https://www.sciencedirect.com/science/article/pii/S2666389925002788> (Visited on June 19, 2026).

⁷“9.3 Trillion Litres of Water: UN Report Exposes Unfathomable Footprint of Data Centers as AI Booms”, Earth.org, available at: <https://earth.org/9-3-trillion-liters-of-water-un-report-exposes-unfathomable-footprint-of-data-centers-as-ai-booms/> (Visited on June 19, 2026).

standard, quality, or composition of goods, and bait advertising.⁸ Section 2(28) separately defines “misleading advertisement” to include any advertisement that falsely describes a product, gives a guarantee likely to mislead consumers, or deliberately conceals material information.⁹ Both definitions are, on their face, technology-neutral: they regulate the misleading or deceptive effect on the consumer rather than the means by which that effect is produced, and there is consequently no obvious doctrinal obstacle to applying them to AI-generated dark patterns, dynamic pricing, or manufactured urgency cues. The Central Consumer Protection Authority (“CCPA”), established under Section 10 of the Act to protect, promote, and enforce consumer rights as a class, possesses under Section 18 a broad power to issue guidelines necessary to prevent unfair trade practices, and under Section 21 the power to direct discontinuance of, and impose penalties for, false or misleading advertisements.¹⁰

What Sections 2(47) and 2(28) do not do, however, is engage with ecological harm as a distinct head of consumer injury. The “deception” that both provisions are built to capture is deception as to the nature, quality, or price of a product or service; it is not deception as to the environmental consequence of the act of purchase, nor as to the environmental cost of the recommending technology. An advertisement that omits the embedded carbon cost of a product, or a recommender system that omits its own energy footprint, commits no falsehood that either provision, read literally, currently treats as actionable, because neither requires that ecological information be disclosed in the first place. The harm a dark pattern causes by inducing an unnecessary purchase is, under the present statutory architecture, cognisable only through the lens of the consumer’s wallet and autonomy, not through the lens of the planet’s resources.

B. The Consumer Protection (E-Commerce) Rules, 2020

The E-Commerce Rules, framed under the Act, impose a layered set of disclosure duties on e-commerce entities, marketplace entities, sellers, and inventory entities. Rule 4 requires every e-commerce entity to make specified disclosures, to maintain a grievance redressal mechanism, and, significant for present purposes, to record consumer consent only through an explicit and affirmative action, expressly prohibiting automatic consent mechanisms such as pre-ticked

⁸The Consumer Protection Act, 2019 (Act 35 of 2019), s. 2(47).

⁹*Id.*, s. 2(28).

¹⁰The Consumer Protection Act, 2019, ss. 10, 18 and 21.

checkboxes.¹¹ Rule 5 imposes further liabilities on marketplace entities, including disclosure of any differential treatment afforded to particular sellers, while Rules 6 and 7 impose corresponding duties on sellers and inventory entities respectively, including the duty not to adopt unfair trade practices and to ensure that advertised features correspond to a product's actual characteristics.¹²

This disclosure architecture is a natural site onto which environmental information could be grafted, because it already establishes the principle that the law may compel platforms to display specified categories of information regardless of whether withholding that information would otherwise amount to a misrepresentation. Rule 4 does not require an e-commerce entity to disclose its legal name and grievance officer because withholding that information is independently deceptive; it requires disclosure because the rule-making authority has determined, as a matter of policy, that consumers ought to have access to that category of information before transacting. The same regulatory technique, mandated disclosure independent of any showing of falsity, is precisely the technique this paper argues should be extended to the ecological footprint of recommended products and the energy intensity of the recommending system. The Rules' present silence on environmental disclosure is accordingly a gap of policy choice, not a gap compelled by the structure of the Rules themselves.

C. The Dark Patterns Guidelines, 2023 and the June 2025 Advisory

The most direct doctrinal engagement with AI-driven choice architecture comes from the CCPA's Guidelines for Prevention and Regulation of Dark Patterns, 2023, issued under Section 18 of the Act. The Guidelines define "dark pattern" as any practice or deceptive design pattern, using user interface or user experience interactions on any platform, that is designed to mislead or trick users into actions they did not originally intend, by subverting or impairing consumer autonomy, decision-making, or choice, and that amounts to a misleading advertisement, an unfair trade practice, or a violation of consumer rights.¹³ Annexure I to the Guidelines lists thirteen specified dark patterns, including false urgency, basket sneaking, confirm shaming, forced action,

¹¹The Consumer Protection (E-Commerce) Rules, 2020, r. 4(9).

¹²*Id.*, rr. 5, 6 and 7.

¹³Guidelines for Prevention and Regulation of Dark Patterns, 2023, cl. 2(e).

the subscription trap, interface interference, bait and switch, drip pricing, disguised advertisement, nagging, trick questions, SaaS billing, and rogue malware, with illustrations appended for each.¹⁴ The Guidelines have generated real enforcement activity. The CCPA found, by order dated 19 June, 2024, that InterGlobe Aviation Limited (IndiGo) had engaged in “confirm shaming” by displaying messages discouraging passengers who declined add-on insurance, and in opaque seat-assignment design that obscured the option to skip paid seat selection; the airline subsequently amended its interface.¹⁵ BookMyShow was issued a notice after the platform was found to be adding a charitable contribution to every ticket through a pre-ticked checkbox, a practice falling squarely within “basket sneaking” under Annexure I, and likewise amended its interface following the CCPA’s intervention.¹⁶ More recently, the Authority has moved from corrective notices to monetary penalties: in 2026 it fined McAfee Software India Private Limited and PhysicsWallah Limited for subscription-renewal and consent interfaces that offered users only an option to “accept risk” without a genuine option to decline, distinguishing these cases from IndiGo’s on the basis that no penalty had followed corrective action in that earlier matter.¹⁷ The State Consumer Disputes Redressal Commission, Chandigarh, has separately treated misleading product-authenticity advertising combined with deceptive design as an actionable unfair trade practice in proceedings against Flipkart.¹⁸ Recognising continuing non-compliance, an independent survey found that an overwhelming majority of a large sample of Indian online platforms continued to deploy at least one dark pattern as of late 2025, the CCPA issued a further Advisory on 5 June, 2025 directing all e-commerce and quick-commerce platforms to conduct a self-audit within three months and to publish a self-declaration of compliance.¹⁹

¹⁴Id., Annexure I.

¹⁵Trilegal, “Guidelines for Prevention and Regulation of Dark Patterns, 2023”, available at: <https://trilegal.com/wp-content/uploads/2023/12/Guidelines-for-Prevention-and-Regulation-of-Dark-Patterns-2023.pdf> (Visited on June 19, 2026).

¹⁶Chambers and Partners, “Dark Patterns and the Limits of Self-Policing: A Critical Look at the CCPA’s June 2025 Advisory”, available at: <https://chambers.com/articles/dark-patterns-and-the-limits-of-self-policing-a-critical-look-at-the-ccpa-s-june-2025-advisory> (Visited on June 19, 2026).

¹⁷Medianama, “Consumer Protection Authority Fines PhysicsWallah and McAfee for Dark Patterns”, available at: <https://www.medianama.com/2026/06/223-consumer-protection-authority-fines-physicswallah-and-mcafee-for-dark-patterns-what-counts-as-consent-now/> (Visited on June 19, 2026).

¹⁸Ashwini Chawla v. Flipkart, State Consumer Disputes Redressal Commission, Chandigarh, discussed in Neeti Niyaman, “Dark Patterns in Digital Platforms: India’s Rules & Guidelines”, available at: <https://neetiniyaman.com/dark-patterns-digital-platforms/> (Visited on June 19, 2026).

¹⁹Medianama, “CCPA Publishes Dark Pattern Self-Audits of 18 E-Commerce and Quick-Commerce Platforms”, available at: <https://www.medianama.com/2025/11/223-ccpa-dark-pattern-self-audit-declarations-18-e-commerce-quick-commerce-platforms/> (Visited on June 19, 2026).

This body of guidance and enforcement establishes, beyond doubt, that India's regulator is willing to treat manipulative algorithmic design as a cognisable consumer harm. It establishes, with equal clarity, that the cognisable harm remains confined to autonomy and economic injury. None of the thirteen specified dark patterns, and none of the enforcement actions taken to date, engages with the ecological consequence of the manipulation, the additional unit of fast fashion produced, the additional device manufactured ahead of necessity, or the computational footprint of the manipulating system itself. The Guidelines are, moreover, of uncertain legal force: Annexure I is expressly illustrative rather than binding, and the 2025 Advisory relies on voluntary self-audit rather than mandatory, independently verified compliance, a structure critic have characterised as reducing enforcement to a box-ticking exercise.²⁰

D. Structural Gaps: Why “Dark” Need Not Mean “Green”

Three structural gaps emerge from this survey. First, a definitional gap: the operative legal concepts, unfair trade practice, misleading advertisement, and dark pattern, are defined by reference to the deceptive or manipulative effect on the individual consumer's decision, not by reference to the aggregate or downstream environmental effect of the decision so induced. A perfectly accurate, non-deceptive recommendation can still be ecologically catastrophic in volume; current law has no doctrinal hook with which to reach it. Second, an institutional gap: the CCPA's enforcement to date has targeted point-of-sale autonomy violations rather than systemic ecological externality, and none of its guidance requires platforms to measure, let alone disclose, the carbon or water cost of the AI systems generating the scrutinised recommendations. Third, an evidentiary gap: even a motivated authority would lack any statutory disclosure obligation compelling platforms to produce the energy-intensity data necessary to evaluate an environmental harm claim, since none of the Act, the E-Commerce Rules, or the Dark Patterns Guidelines impose a duty to measure or disclose such data in the first place.

²⁰Supra note 18.

IV. ALGORITHMIC TRANSPARENCY AS AN EXTENSION OF THE RIGHT TO KNOW

The Act recognises, among the substantive rights of a consumer, the right to be informed about the quality, quantity, potency, purity, standard, and price of goods, products, or services, so as to protect the consumer against unfair trade practices.²¹ This right to know has historically been understood, and litigated, in relation to the physical attributes of a product, ingredients, safety warnings, pricing breakdowns. There is no principled reason, however, why its scope should stop at the product and exclude the process by which the product was surfaced to the consumer in the first place. If a consumer has a right to know the quantity and quality of what they are buying precisely because that information shapes a more autonomous purchasing decision, a consumer has an equally strong claim to know that the very interface presenting the choice has been engineered, by an algorithm trained to maximise engagement or conversion, to compress the deliberation that genuine autonomy requires. Algorithmic transparency, in other words, is not a novel right requiring fresh constitutional or statutory invention; it is the logical extension of the existing right to be informed into the procedural dimension of the transaction, supplementing the right's traditional focus on the product with a corresponding focus on the process of choice architecture itself.

This reframing matters because it supplies the missing doctrinal hook identified in Part III. Once algorithmic transparency is understood as falling within the existing right to be informed, two further extensions follow naturally. The first is disclosure of the ecological footprint of the AI-recommended product itself, not merely its price and quality, but its embedded environmental cost, a category of information already partially addressed, if patchily, by India's Ecomark scheme discussed in Part VI. The second, more novel, extension is disclosure of the energy intensity of the recommending technology, an item of information that has no real analogue in traditional consumer law because traditional commerce had no comparable computational intermediary standing between the consumer and the product. It is to the comparative law that has begun to grapple with this intermediary that the paper now turns.

²¹The Consumer Protection Act, 2019, s. 2(9).

V. COMPARATIVE REGULATORY RESPONSES: LEARNING FROM THE EUROPEAN UNION

A. Article 25 of the Digital Services Act

The European Union's Digital Services Act, which became directly applicable across the Union in February 2024, was the first instrument of Union law to engage explicitly with manipulative interface design, even though the term "dark pattern" itself appears only in the accompanying recital rather than the operative article.²² Article 25(1) prohibits providers of online platforms from designing, organising, or operating their online interfaces in a manner that deceives or manipulates recipients of their service, or that otherwise materially distorts or impairs the ability of recipients to make free and informed decisions.²³ Recital 67 clarifies that this prohibition is intended to capture dark patterns, defined as structures, designs, or functionalities of online interfaces used to deceive, manipulate, or otherwise materially distort or impair a recipient's ability to make autonomous, informed choices. Article 25(3) further identifies specific illustrative practices, including repeated solicitation after a recipient has already declined, commonly termed "nagging."²⁴

Two structural features distinguish Article 25 from India's Dark Patterns Guidelines. First, Article 25 sits within a binding Regulation directly applicable across the Union, rather than within an executive guideline of contested legal force. Second, Article 25(1) is deliberately open-textured, addressed to any structure, design, or functionality that distorts free and informed decision-making, rather than to a closed or even illustrative list of named practices, which gives regulators interpretive room to reach novel manipulative techniques, including those generated by AI personalisation, without amending the underlying legal text each time a new pattern emerges. India's Annexure I, by contrast, enumerates thirteen named patterns and expressly disclaims any binding interpretive force, a structure better suited to public communication than to adjudicative certainty.

²²Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October, 2022 on a Single Market For Digital Services (Digital Services Act), recital 67.

²³*Id.*, art. 25(1).

²⁴*Id.*, art. 25(3)(b).

B. Article 5 of the Artificial Intelligence Act

The Artificial Intelligence Act goes further still by regulating the technology rather than only the interface it produces. Article 5(1)(a) prohibits the placing on the market, putting into service, or use of an AI system that deploys subliminal techniques beyond a person's consciousness, or purposefully manipulative or deceptive techniques, with the objective or effect of materially distorting a person's behaviour by appreciably impairing their ability to make an informed decision, in a manner that causes or is reasonably likely to cause significant harm.²⁵ Article 5(1)(b) separately prohibits AI systems that exploit the vulnerabilities of a person or specific group by reason of age, disability, or social or economic situation, again where the result distorts behaviour so as to cause significant harm.²⁶ The European Commission's guidelines on these provisions clarify that AI-personalised advertising is not, without more, manipulative; what matters is whether the system deploys subliminal, purposefully manipulative, or deceptive techniques that subvert autonomy or exploit vulnerability, with compliance with data protection law treated as a relevant, if not conclusive, mitigating factor.²⁷

Article 5 is significant for the present argument because it locates the regulatory target one step upstream of the interface, at the level of the AI system itself, which is precisely the locus at which environmental footprint is generated and could, in principle, be measured and disclosed. A regulatory architecture willing to ask whether a recommendation system manipulates consciousness is institutionally well placed to ask, in the same breath, what that system costs the grid and the water table, yet, notably, even the AI Act's prohibited-practices regime does not currently impose any environmental disclosure obligation on commercial recommender or pricing systems, leaving this as much a gap in the comparator jurisdiction as in India.

C. The Proposed Digital Fairness Act

The European Commission is, at the time of writing, preparing a Digital Fairness Act, expected in the course of 2026 as a successor to Article 25 of the Digital Services Act, intended to

²⁵Regulation (EU) 2024/1689 of the European Parliament and of the Council of 13 June, 2024 laying down harmonised rules on Artificial Intelligence (Artificial Intelligence Act), art. 5(1)(a).

²⁶*Id.*, art. 5(1)(b).

²⁷Future of Privacy Forum, "Red Lines under the EU AI Act: Understanding Manipulative Techniques and the Exploitation of Vulnerabilities", available at: <https://fpf.org/blog/red-lines-under-the-eu-ai-act-understanding-manipulative-techniques-and-the-exploitation-of-vulnerabilities/> (Visited on June 19, 2026).

address addictive service design and AI-driven deceptive marketing more comprehensively than the DSA's interface-focused prohibition currently permits.²⁸ Enforcement sweeps by international consumer protection networks have found extensive non-compliance even with the existing framework: a 2024 joint sweep of over six hundred traders found that more than three-quarters deployed at least one dark pattern, with sneaking practices, such as the inability to disable subscription auto-renewal, among the most common.²⁹ The proposed Act signals that even the comparatively mature European framework regards interface-level prohibition as necessary but insufficient, and is moving towards regulating the underlying business model and algorithmic incentive structure, direction with direct relevance to the environmental dimension this paper foregrounds, since that business model is precisely the locus of the overconsumption incentive.

D. Transposability and Its Limits

Three features of the EU framework appear transposable to the Indian context without requiring wholesale legislative reform. First, the open-textured, effects-based prohibition of Article 25(1), addressed to any design that materially distorts free and informed decision-making, could usefully supplement, rather than replace, Annexure I's closed list, giving the CCPA the same interpretive latitude over novel AI-driven techniques that EU regulators currently enjoy. Second, the AI Act's technique of regulating the AI system itself, and not merely the interface it produces, supplies a doctrinal template for extending disclosure obligations upstream to the recommender or pricing engine rather than confining them to the resulting webpage. Third, the EU's reliance on systematic, independently conducted compliance sweeps, rather than platform self-audit alone, offers a corrective to the evidentiary weaknesses of the CCPA's June 2025 Advisory.

Transposition is not, however, unconstrained. The European Union's framework operates against a backdrop of comprehensive data protection regulation under the General Data Protection Regulation, treated by the AI Act's own guidelines as a mitigating factor in assessing manipulation; India's younger Digital Personal Data Protection Act, 2023 occupies similar terrain but with a shorter enforcement record, so any transplanted standard would lean on a thinner surrounding infrastructure. Equally, the EU's institutional capacity, spanning data protection authorities, competition regulators, and the Commission's own enforcement directorate, exceeds

²⁸Digital Fairness Act Updates, available at: <https://www.digital-fairness-act.com/> (Visited on June 19, 2026).

²⁹Ibid.

what a single body such as the CCPA, already stretched across unfair trade practice, misleading advertisement, and product safety jurisdiction, can replicate without additional resourcing. Any Indian adoption of these features must therefore be calibrated to existing institutional capacity rather than imported wholesale.

VI. TOWARDS A RECALIBRATED ENFORCEMENT FRAMEWORK: ALGORITHMIC ECO-DISCLOSURE

A. Building on Existing Eco-Labelling Infrastructure

India is not starting from a blank slate on environmental disclosure. The Ecomark scheme, instituted in 1991 and recently overhauled through rules notified in 2024 that align it with the government's Lifestyle for Environment Mission, already establishes a cradle-to-grave methodology for certifying products as environmentally friendly, administered jointly by the Central Pollution Control Board and the Bureau of Indian Standards.³⁰ The scheme has suffered from low uptake and weak consumer recognition since inception, regularly attributed to its voluntary character and the absence of any link to mandatory point-of-sale disclosure of the kind the E-Commerce Rules already impose for matters such as country of origin and grievance officer details.³¹ The proposal advanced here is not to create a parallel eco-labelling bureaucracy but to graft Ecomark-equivalent disclosure onto the existing, already mandatory, E-Commerce Rules architecture, converting a voluntary certification regime into a disclosure obligation enforceable through the same Rule 4 mechanism that already compels disclosure of grievance officers and return policies.

B. A Two-Tier Disclosure Architecture

The framework proposed here, termed algorithmic eco-disclosure, has two tiers, corresponding to the two channels of ecological harm identified in Part II. The first tier addresses the product: where a good or service is surfaced to a consumer through an AI-driven recommendation, dynamic pricing, or urgency-cue mechanism, the platform should be required to

³⁰Ministry of Environment, Forest and Climate Change, Government of India, Eco-Mark (Quality Control) Rules, 2024.

³¹The Lex-Warrier, "The Indian Eco-labelling Scheme: Ecomar", available at: <https://lex-warrier.in/archives/indian-eco-labelling-scheme-ecomark.html> (Visited on June 19, 2026).

display, alongside the recommendation, a standardised ecological footprint indicator for that product, drawing on Ecomark-equivalent criteria, extended where necessary to electronics and apparel categories most implicated in algorithmically accelerated obsolescence and fast-fashion cycles. The second tier addresses the technology: platforms deploying recommender or dynamic pricing systems above a threshold of scale, for instance, measured by transaction volume or active user base, mirroring the “very large online platform” threshold the Digital Services Act uses to calibrate obligations, should be required to disclose, in aggregate and standardised form, the energy intensity of the AI system generating personalised commercial content, expressed as a simple comparative metric analogous to the star-rating disclosures long familiar to Indian consumers from energy-labelled appliances, rather than as raw and unintelligible technical data.

Both tiers should be triggered specifically by the deployment of an AI-driven choice-architecture mechanism, rather than attaching to e-commerce transactions generally, so as to target the precise locus of ecological harm identified in this paper, the algorithmic acceleration of consumption, without imposing a disproportionate compliance burden on routine, non-personalised online retail.

C. Institutional Locus: The Central Consumer Protection Authority

The CCPA already possesses, under Section 18 of the Act, the power to issue guidelines necessary to prevent unfair trade practices and protect consumer interest, the very power under which the Dark Patterns Guidelines were issued.³² No new statutory authority is strictly necessary to operationalise algorithmic eco-disclosure: the CCPA could issue an amendment to, or a supplementary annexure within, the existing Guidelines, explicitly recognising that the omission of ecological footprint disclosure from an AI-driven recommendation itself amounts to a “deceptive design pattern” within the existing definition of dark pattern, given that an interface concealing the ecological consequence of an algorithmically induced purchase decision impairs the consumer’s autonomy and decision-making within the existing statutory language just as surely as an interface concealing a hidden fee.³³ This interpretive move would allow the proposed disclosure obligation to attach to the existing, judicially and administratively tested machinery of unfair trade practice and misleading advertisement enforcement under Sections 18 and 21 of the

³²The Consumer Protection Act, 2019, s. 18.

³³Supra note 15.

Act, rather than requiring the creation of a freestanding environmental regulator with overlapping jurisdiction.

Enforcement should, however, depart from the voluntary self-audit model adopted in the CCPA's June 2025 Advisory, whose limitations, most starkly illustrated by independent findings that a large majority of platforms surveyed continued to deploy unaudited dark patterns notwithstanding the Advisory, counsel against repeating that model for environmental disclosure.³⁴ A periodic, CCPA-commissioned third-party verification mechanism, modelled loosely on the compliance-sweep methodology used by international consumer protection networks, would supply the independent verification that self-declaration alone cannot.

D. Anticipated Objections

Three objections merit brief consideration. The first concerns measurability: the energy intensity of a given recommendation is difficult to isolate, since inference workloads are typically shared across millions of simultaneous queries on shared infrastructure. This objection, while valid, is not fatal; aggregate, platform-level disclosure, analogous to a corporate sustainability report rather than a per-transaction receipt, sidesteps the attribution problem while still supplying the comparative information a consumer or regulator needs. The second concerns disproportionate burden on smaller players: a threshold calibrated to scale, as proposed above, exempts platforms below a defined size from the technology-tier obligation while retaining the lighter product-tier obligation. The third, and most serious, objection is that a disclosure-only regime, without independent verification, risks replicating the soft-law fate of the Dark Patterns Guidelines and the 2025 Advisory, becoming a compliance ritual rather than a substantive constraint on design; this is precisely why the proposal insists on third-party verification as a non-negotiable feature, rather than an optional enhancement, of the framework.

³⁴Supra note 18; see also Internet Freedom Foundation, "Dark Patterns and weak remedies: Why the CCPA Advisory falls short", available at: <https://internetfreedom.in/dark-patterns-and-weak-remedies-why-the-ccpa-advisory-falls-short/> (Visited on June 19, 2026).

VII. CONCLUSION

The conventional account of artificial intelligence in commerce treats environmental harm, if it treats it at all, as an unfortunate but incidental by-product of an otherwise welfare-enhancing technology. This paper has argued for a different starting point: that algorithmically engineered choice architecture and the computational infrastructure that powers it are environmental actors in their own right, and that consumer protection law, built, in India as elsewhere, around the regulation of deception and manipulation at the point of sale, is institutionally well placed to regulate them, provided its existing doctrinal categories are read with appropriate ambition. The Consumer Protection Act, 2019, the E-Commerce Rules, 2020, and the Dark Patterns Guidelines, 2023 already supply the CCPA with the statutory hooks needed to treat the concealment of ecological footprint and algorithmic energy intensity as a deceptive design pattern; the comparative experience of the European Union's Digital Services Act and Artificial Intelligence Act demonstrates both the promise and the limits of interface-level and system-level prohibition respectively; and India's own, presently under-used, Ecomark infrastructure supplies a ready template for the product-tier disclosure this paper proposes.

What remains is the institutional will to read the consumer's right to be informed as a right that travels with the transaction into its algorithmic and ecological hinterland, to insist, in short, that the digital marketplace account not only for what the consumer is told, but for what the planet is made to bear in the telling. Consumer protection law, on this reading, ceases to be a passive bystander to technologically mediated consumption and becomes what this paper has called it throughout: an active instrument of environmental governance for the algorithmic age.