



INTERNATIONAL JOURNAL OF HUMAN RIGHTS LAW REVIEW

An International Open Access Double Blind Peer Reviewed, Referred Journal

Volume 5 | Issue 4 | 2026

Art. 42

From Pixels to Personhood: AI Virtual Influencers and the Jurisprudence of Legal Personality

Rahul Bishnoi

Research Scholar,

The National Law Institute University, Bhopal, Madhya Pradesh, India

Dr. Astitwa Bhargava

Assistant Professor,

The National Law Institute University, Bhopal, Madhya Pradesh, India

Recommended Citation

Rahul Bishnoi and Dr. Astitwa Bhargava, *From Pixels to Personhood: AI Virtual Influencers and the Jurisprudence of Legal Personality*, 5 IJHRLR 614-630 (2026).
Available at www.ijhrlr.in/current-issues/.

This Article is brought to you for free and open access by the International Journal of Human Rights Law Review by an authorized Lex Assisto & Co. administrator.

For more information,
please contact humanrightlawreview@gmail.com

From Pixels to Personhood: AI Virtual Influencers and the Jurisprudence of Legal Personality

ABSTRACT

Lil Miquela has 2.6 million followers, a record deal, and a Prada campaign to her name, and she has never existed.¹ Imma models for Ikea and Amazon from a Tokyo apartment nobody has ever entered; Aitana López reportedly earns her creators more in a good month than most human models earn in Spain.² These figures occupy a genuinely strange legal position. They sign endorsement deals but cannot be sued for breaching them. They generate revenue but cannot own it. They say things that could constitute defamation or false advertising, with no settled answer as to who bears responsibility for what they say. Commentary on this problem reaches, almost by reflex, for the vocabulary of personhood - asking whether an AI influencer should be recognised as a person, or even whether she has rights. This paper argues that the reflex is a category error, and that untangling it requires taking the jurisprudence of legal personality seriously rather than treating "person" as a single switch to be flipped on or off. Working through the fiction, concession, and realist theories of personality that have shaped the concept since Savigny and Gierke, through Naffine's account of the person as a contested legal category, and through the stricter conditions that human rights law attaches to recognition before the law, the paper shows that virtual influencers satisfy neither the administrative logic that justified corporate personality nor the dignity-based logic that grounds human rights. What they need instead is a narrower, disclosure-centred legal status built around the humans and companies actually doing the work, and the paper closes by setting out what that status should look like.

KEYWORDS

Legal Personality, Virtual Influencers, Artificial Intelligence, Accountability, Disclosure

1. INTRODUCTION

In 2016, a young woman named Lil Miquela appeared on Instagram with

¹ *Virtual Humans*, 'Lil Miquela' (VirtualHumans.org)

<<https://www.virtualhumans.org/human/lil-miquela>> accessed 15 July 2026.

² 'Meet the Spanish AI Model Earning up to €10,000 a Month' (Euronews, 27 December 2024) <<https://www.euronews.com/next/2024/12/27/meet-the-first-spanish-ai-model-earning-up-to-10000-per-month>> accessed 15 July 2026.

no explanation attached to the account. She posted selfies, talked about her music, and picked a public feud with another influencer over politics. It took months before anyone confirmed what a lot of people already suspected: she was computer-generated, built by a Los Angeles startup called Brud.³ By the time the truth came out, she already had a fan base that did not much care.

Eight years on, virtual influencers are not a novelty act. Calvin Klein, Samsung, and BMW have all run campaigns built around one, and industry estimates put the virtual influencer market at roughly \$22 billion in 2024, growing faster than influencer marketing built around actual people.⁴ Some of these figures are photorealistic, designed to be mistaken for a real person at a glance. Others, like Miquela, wear their artificiality openly, sitting somewhere between mascot and celebrity.

None of this required a change in the law. A company can build a CGI character and put her on the internet without asking anyone's permission, and that is exactly what has happened. Commerce, though, has a habit of outrunning statutes written before it existed, and virtual influencers are as clean an example as any. When a human influencer promotes a product that turns out to be dangerous, or lies about having used it, there is a person to sue, a person to fine, a name that goes on the regulator's complaint. When a virtual influencer does the same thing, the "person" in question is a Blender file and a render pipeline. The company that built her sits somewhere behind the curtain, and depending on how the deal is structured, so does a separate marketing agency, a separate rights holder, and possibly a third-party AI vendor whose model wrote her dialogue in real time.

Popular commentary tends to frame the resulting gap as a question about personhood - should she have rights, is she a person, does the law need to recognise her as something. That framing is understandable, and it is almost entirely wrong. It borrows a legal category built to solve one kind of administrative problem and asks it to solve a different one. Jurisprudence has spent well over a century working out what "person" does and does not mean as a legal term of art, and none of that history supports the popular instinct to treat personhood as a dignity that an entity either possesses or lacks. This paper's central claim is that the personhood question and the accountability question have fused together in public discussion in a way jurisprudence itself never fuses them, and that separating the two is the necessary first move before any

³ n 1.

⁴ Influencer Marketing Hub, 'The State of Virtual Influencer Marketing' (Influencer Marketing Hub, 2024).

sensible policy can follow.

The paper proceeds in seven parts. Section 1 sets out a working taxonomy of what these entities actually are, since “AI virtual influencer” covers several distinct things that get discussed as one. Section 2 works through the competing jurisprudential theories of legal personality and asks what each would actually require of a virtual influencer. Section 3 turns to human rights law specifically, since the language of rights surfaces constantly in this debate and deserves harder scrutiny than it usually gets. Section 4 surveys where the existing patchwork of intellectual property, agency, and consumer protection law breaks down. Section 5 works through four illustrative cases. Section 6 proposes a regulatory framework built on the earlier sections' conclusions, and Section 7 takes the strongest objections to that framework seriously before the paper closes.

2. WHAT A VIRTUAL INFLUENCER ACTUALLY IS

Before asking what the law should do, it helps to be precise about what these entities are, because “AI virtual influencer” covers at least three distinct things that tend to get discussed as though they were one.

The first is the fully scripted CGI persona. Lil Miquela falls into this category, certainly in her earlier years: a team of writers, artists, and marketers decides what she posts, when, and why, and the “AI” label is often more aspirational than technical. Her outputs are human-authored and merely attributed to a fictional digital being – closer to a Muppet than a chatbot, a puppet with unusually good rendering.

The second is the AI-generated persona with limited autonomy. Aitana López, created by the Barcelona agency The Clueless after its founder grew frustrated with the cost and unpredictability of working with real models, uses generative tools for both her images and some of her captions, with a small creative team planning her “life” each week and reviewing content before it goes live.⁵ The line between assisted and autonomous is blurry here, and agencies have an obvious commercial incentive to talk up the autonomy for novelty value while keeping a human hand on the wheel for legal safety.

The third, smaller but growing category is the real-time interactive avatar: an influencer or brand mascot that responds live to fans in livestreams, direct messages, or voice, using a language model with minimal human review of any individual output. This is the category

⁵ 'Aitana López' (*Wikipedia*) <https://en.wikipedia.org/wiki/Aitana_L%C3%B3pez> accessed 15 July 2026; 'Meet the Spanish AI Model Earning up to €10,000 a Month' (n 2).

that raises the sharpest legal questions, precisely because it is the one where a company genuinely cannot say, word for word, what its influencer will tell a given fan on a given day.

The distinction matters because the legal analysis differs across the three. A scripted CGI puppet raises much the same questions as any brand mascot – Tony the Tiger does not need personhood, and neither does Miquela in her scripted mode. An unscripted, LLM-driven avatar is a different animal, closer to the broader question of who answers for an autonomous system's output when nobody at the company wrote or approved that particular sentence. This paper focuses mainly on the second and third categories. The first is already handled reasonably well by existing law: the character is intellectual property, and the humans who operate her are ordinary agents of the company. It is where AI genuinely enters the picture that the frameworks start to strain.

3. THE JURISPRUDENCE OF LEGAL PERSONALITY

“Person” is one of the oldest and most argued-over words in legal theory, and it is worth pausing on that argument before applying the word to anything new. The Roman *persona* originally denoted a theatrical mask, a role one stepped into rather than a fact about one's nature, and the law borrowed the word for exactly that reason.⁶ A legal person is not a claim about what an entity is; it is a statement about what the law will let that entity do: hold property, sue, be sued, enter contracts, bear duties. Four jurisprudential accounts have competed to explain when and why the law should extend that status, and each gives a different, and differently unhelpful, answer for a virtual influencer.

Friedrich Carl von Savigny's fiction theory, developed in the nineteenth century to explain corporate personality within Roman law's framework, held that only human beings are persons in any real sense; the corporation is a fiction, a convenient legal mask the state chooses to grant because doing so serves a practical end.⁷ On this account personhood is never discovered in an entity's nature – it is conferred, and it is conferred instrumentally. If Savigny is right, the question for a virtual influencer is not whether she resembles a person closely enough to qualify, but whether granting her the fiction would serve some purpose the law cannot achieve any other way. Nothing in the preceding paragraphs suggests it would: the fiction, applied to Miquela herself

⁶ The theatrical origin of *persona* and its adoption into Roman legal usage is traced in John Dewey, 'The Historic Background of Corporate Legal Personality' (1926) 35 Yale LJ 655, 656.

⁷ Friedrich Carl von Savigny, *System des heutigen römischen Rechts* (Veit 1840) vol II, discussed in Dewey (n 6) 656–57.

rather than to Brud, would grant contracting power to an entity that owns no assets against which a judgment could ever be enforced.

The realist or organic theory, most associated with Otto von Gierke, pushed the opposite way. Gierke argued that a corporate group is not merely fictional but has a genuine collective will and real social existence that the law recognises rather than invents – a *Genossenschaft*, a real association with its own life.⁸ Whatever the merits of that view for guilds, churches, and corporations, it has no purchase here. A virtual influencer has no members, no collective will, and no existence independent of the render files and marketing calendar that constitute her. She is not a group whose reality the law might belatedly recognise; she is closer to what Gierke's theory was built to distinguish itself from – a mere name attached to an asset.

Between these two poles sits the concession theory, which treats legal personality as a grant the sovereign makes to further state purposes, and historically bound up with chartering corporations for public ends such as building canals or running banks. John Salmond's classical definition captures the functional core shared by all three accounts: a legal person is “any subject-matter other than a human being to which the law attributes personality” as “a purely arbitrary creation of the law.”⁹ Arbitrary here does not mean random. It means chosen, and chosen for a reason a court or legislature could state out loud. That is the discipline missing from most popular commentary on virtual influencers, which asks whether the persona feels enough like a person to deserve the status rather than asking what work the status would do.

Ngaire Naffine's more recent taxonomy sharpens the point by distinguishing between at least four different things lawyers mean when they use the word “person”: the person as a philosophical question about moral status; the person as whatever answer positive law happens to give at a given time; the person as a bearer of rights and duties, defined by that bundle rather than by anything prior to it; and the person as a rhetorical or metaphorical device used to dramatise an entity's importance.¹⁰ Much of the discourse around AI virtual influencers slides between the first sense and the fourth without noticing the slide and treating a marketing question (“she feels like a real personality”) as though it settled a philosophical one (“she has moral status”), and then reaching for the second sense (“the law should recognise her”) as if it

⁸ Otto von Gierke, *Political Theories of the Middle Age* (Frederic William Maitland tr, Cambridge University Press 1900) xxvi–xxx; see also Dewey (n 6) 660–64.

⁹ John Salmond, *Jurisprudence* (Glanville L Williams ed, 10th edn, Sweet & Maxwell 1947) 318.

¹⁰ Ngaire Naffine, 'Who Are Law's Persons? From Cheshire Cats to Responsible Subjects' (2003) 66 MLR 346.

followed automatically. It does not. Visa Kurki's more recent "bundle theory" formalises Naffine's third sense: legal personhood just is the sum of the specific incidents i.e. the capacity to sue, to hold title, to be a defendant and that a legal system attaches to an entity, and different entities can carry entirely different bundles without any of them being more or less a "person" in some deeper sense.¹¹ A ship arrested *in rem* in admiralty proceedings carries a tiny bundle: enough personality to be sued, nothing more. A corporation carries a much larger one. Under Kurki's account the question this paper needs to answer is not "is a virtual influencer a person," which has no fixed content, but "which specific incidents, if any, should attach to her," which is a question with an actual answer.

Run the corporate case through this lens and the disanalogy becomes exact rather than merely intuitive. The United States Supreme Court's 1886 pronouncement that corporations are persons within the meaning of the Fourteenth Amendment, and the line of cases building on it since, did not rest on any claim that a corporation resembles a human being.¹² It rested on the practical fact that a corporation outlives its founders, holds assets distinct from any shareholder's assets, and needs to enter binding contracts as an ongoing enterprise rather than as a person who might die or resign tomorrow. Corporate personhood solves a continuity problem and an asset-separation problem. A virtual influencer has neither. She owns nothing - the company that made her owns everything, including her. She does not outlive anyone; if the studio dissolves, she disappears with it, exactly as a discontinued brand mascot does. The entity that actually needs continuity, contracting power, and exposure to liability is the company, and the company already has all the personality it needs, because it is already incorporated. Handing the same bundle to the character on top of that would be redundant at best.

4. HUMAN RIGHTS AND THE LIMITS OF THE PERSONHOOD ANALOGY

If jurisprudence generally treats personhood as functional rather than dignity-based, human rights law is the one corner of the field that resists the reduction, and it is worth pausing here precisely because so much of the public conversation about AI personas borrows its emotional register from rights talk without checking whether the borrowing is warranted.

Article 6 of the Universal Declaration of Human Rights states that

¹¹ Visa AJ Kurki, *A Theory of Legal Personhood* (Oxford University Press 2019) ch 1.

¹² *Santa Clara County v Southern Pacific Railroad Co* 118 US 394 (1886); see also *Citizens United v Federal Election Commission* 558 US 310 (2010).

“everyone has the right to recognition everywhere as a person before the law,” language repeated almost verbatim in Article 16 of the International Covenant on Civil and Political Rights.¹³ Both provisions were drafted against a specific historical backdrop – the mid-twentieth-century memory of states stripping categories of human beings of legal recognition altogether, and both are anchored explicitly to “everyone,” a term the surrounding instruments make clear means every human being, not every entity capable of acting in the world.¹⁴ This is not an oversight the drafters might have corrected with more imagination. Human rights instruments are built on a premise that traces back through the liberal tradition to Kant: that persons in the morally loaded sense possess a dignity grounded in rational autonomy, and that this dignity is precisely what licenses treating them as ends rather than merely as means.¹⁵ A corporation's legal personality was never meant to invoke that premise, and nobody engaged in nineteenth-century company law thought otherwise. The confusion enters only when commentators reach for corporate personhood as a precedent for AI personhood and then, in the same breath, start asking whether the AI persona “deserves” recognition – smuggling in the human-rights sense of the word while borrowing the corporate-law justification for extending it.

The clearest way to see the gap is to look at what happens when the law does extend person-like status to non-human entities outside the corporate context, because in every instance the law has been careful to avoid exactly the dignity-based framing that popular AI commentary reaches for by default. New Zealand's *Te Awa Tupua (Whanganui River Claims Settlement) Act 2017* declares the Whanganui River a legal person, but section 14 immediately vests its rights and duties in two human guardians, Te Pou Tupua, appointed jointly by the Whanganui iwi and the Crown, because a river cannot appear in court, and somebody has to.¹⁶ The point of the grant was never that the river has interests of the kind the Universal Declaration protects; it was to give a mechanism for standing to a body of water that Māori claimants had spent 140 years arguing deserved legal protection through the settlement of a Treaty claim, not through a finding that rivers think or suffer. Days after that Act passed, the Uttarakhand High Court in India tried something similar for the Ganges and Yamuna rivers, declaring them “living entities” with the rights of legal persons, and the Supreme Court of India stayed the order within months, in part because nobody could work out who would

¹³ Universal Declaration of Human Rights (adopted 10 December 1948) UNGA Res 217 A(III) art 6.

¹⁴ International Covenant on Civil and Political Rights (adopted 16 December 1966, entered into force 23 March 1976) 999 UNTS 171, art 16.

¹⁵ Immanuel Kant, *Groundwork of the Metaphysics of Morals* (Mary Gregor tr, Cambridge University Press 1997) 4:428–29.

¹⁶ *Te Awa Tupua (Whanganui River Claims Settlement) Act 2017* (NZ), s 14.

be liable when the river, exercising its new legal personality, flooded someone's farm.¹⁷ The Whanganui model survived because it paired the fiction with two accountable humans; the Ganges ruling collapsed because it did not.

The animal rights litigation pursued by the Nonhuman Rights Project makes the same point from a different direction. In its long-running effort to secure habeas corpus relief for Happy, an elephant held at the Bronx Zoo, the organisation argued that Happy's cognitive complexity entitled her to the "person" status habeas corpus requires. New York's highest court rejected the claim in 2022, holding that habeas corpus protects a category defined by legal tradition and human autonomy, and declining to extend personhood to a non-human animal however sophisticated its cognition.¹⁸ Whatever one thinks of the result, the case shows a court squarely confronting the dignity question this paper has been keeping separate from the functional one, and confirms that even a strong, evidence-backed claim to cognitive capacity has not been enough, on its own, to move a court to grant legal personhood in the human rights sense. A virtual influencer has no cognitive claim to make at all. She does not suffer, prefer, or have interests that could be frustrated. The rights-holder sense of personhood that Happy's litigation tested, and failed to secure even for a cognitively complex elephant, is simply not in play for a Instagram account run by a five-person creative studio.

What is left, once the dignity-based sense of personhood is set aside, is the narrower administrative sense this paper has been building toward: the question of whose name goes on the contract, and who answers when something goes wrong. That question does not require personhood talk at all. It requires only that an identifiable human or corporate actor stand behind the persona with clear legal exposure and which is precisely the bundle a company already carries by virtue of being incorporated, without needing to extend any of it to the character superimposed on top.

5. THE EXISTING LEGAL PATCHWORK, AND WHERE IT BREAKS

Four bodies of law currently touch virtual influencers, each incompletely.

The character herself, her design, her name, her likeness, is protected

¹⁷ *Mohd Salim v State of Uttarakhand*, Writ Petition (PIL) No 126 of 2014 (Uttarakhand High Court, 20 March 2017); stay granted in *State of Uttarakhand v Mohd Salim*, SLP (C) No 016879/2017 (Supreme Court of India, 7 July 2017).

¹⁸ *Nonhuman Rights Project Inc v Breheny* 38 NY3d 555 (NY 2022).

reasonably well under existing intellectual property law. Copyright covers the artistic work, trademark can cover the name and branding once registered, and nothing stops a company from treating Lil Miquela roughly the way Disney treats Mickey Mouse. Where the law gets interesting is the reverse question: can the outputs a virtual influencer generates on the fly (a caption an LLM writes in response to a comment, a live reply to a fan) themselves be copyrighted? The D.C. Circuit's 2025 decision in *Thaler v Perlmutter*, upheld on further appeal and left standing when the Supreme Court denied certiorari in March 2026, confirms the Copyright Office's longstanding position that a work must owe its origin to a human author to be eligible for protection at all.¹⁹ That creates an odd asymmetry. The influencer is fiercely protected intellectual property, but the words coming out of her mouth, if generated with minimal human input, may belong to no one, meaning a competitor could in principle republish an AI influencer's spontaneous captions without infringing anything, so long as the character design itself is left alone.

Agency law fares no better. The classical definition of agency in the Restatement requires a manifestation of consent by one person that another shall act on their behalf, and both parties are assumed to be legal persons capable of that manifestation.²⁰ A virtual influencer cannot consent to anything, so she cannot be an agent in the technical sense; the humans who direct her are the company's agents, and any liability the company incurs flows through ordinary vicarious liability and employment law rather than through anything resembling the influencer's own legal capacity. That is workable for the scripted category, where a human reviewed and approved everything. It gets far shakier for the live, LLM-driven category, where the output was neither reviewed nor, in any meaningful sense, authored by an identifiable human at the moment it was published which left a gap between the company's actual conduct and the harm the output caused that ordinary agency doctrine was never built to bridge.

Consumer protection law occupies the same uneasy middle ground. The Federal Trade Commission's endorsement guides require a "clear and conspicuous" disclosure of a material connection between an endorser and an advertiser, and the FTC has made clear the rule applies regardless of whether the endorser is human, but the guides remain interpretive, built for the ordinary human-influencer case, and offer no settled position on how liability should attach when the endorsing

¹⁹ *Thaler v Perlmutter* 130 F4th 1039 (DC Cir 2025), cert denied 606 US __ (2 March 2026); see also US Copyright Office, *Compendium of US Copyright Office Practices* (3rd edn) § 306.

²⁰ Restatement (Third) of Agency § 1.01 (American Law Institute 2006).

persona's specific words were never reviewed by anyone.²¹ Defamation and false advertising claims face a structurally similar problem: both require identifying who “said” the statement in a legally meaningful sense, and courts have not yet had to decide whether that is the platform, the studio, the marketing agency, or the underlying AI vendor when three or four of those parties are involved in producing a single unscripted sentence.

None of these four bodies of law is broken in the sense of failing on its own terms. Each does roughly what it was built to do. The problem is that none of them was built with an entity in mind that generates its own words in real time while owning nothing and answering to no one directly — and stitching partial coverage from intellectual property, agency, and consumer protection together does not add up to a coherent answer to the only question that actually matters, which is who is on the hook.

6. FOUR CASES

Four episodes from the short history of virtual influencers illustrate the gap in ways abstract argument cannot.

Lil Miquela's own corporate history is the cleanest illustration of the ownership-transfer problem this paper has been building toward. Brud, the Los Angeles studio that created her, was acquired by Dapper Labs in October 2021, in an all-equity deal that folded Brud's roughly thirty-person team into a new unit called Dapper Collectives.²² Miquela's content, her voice, and her commercial relationships all continued uninterrupted through the transition; her multi-million-strong following had no formal mechanism by which to learn, at the moment it happened, that the entity actually accountable for her conduct had changed. Nothing about that transition was unlawful as company acquisitions happen constantly, and nobody suggests they should not, but it shows precisely the disclosure gap a registration requirement, discussed below, is designed to close. A follower, regulator, or prospective plaintiff dealing with Miquela in November 2021 had no easy way to determine that they were now dealing with a different corporate parent than the one that had built the relationship they thought they had.

Aitana López raises a different problem: labour displacement dressed up as innovation. The Clueless built her explicitly to avoid the

²¹ Federal Trade Commission, 'Guides Concerning the Use of Endorsements and Testimonials in Advertising' 16 CFR pt 255.

²² Dapper Labs, 'Dapper Labs Acquires Brud' (Dapper Labs Newsroom, 4 October 2021) <<https://www.dapperlabs.com/newsroom/dapper-labs-acquires-brud>> accessed 15 July 2026.

cost, scheduling difficulty, and reputational risk of contracting real models, and her creators have been candid in interviews that the economics were the whole point and she can out-earn a human model at a fraction of the overhead, with no risk of the campaign disruptions a real person's life inevitably brings.²³ Nothing about that is currently illegal, and nothing in this paper's proposed framework would make it illegal; a studio's right to build a synthetic model is not in question. What the case does illustrate is a labour-market effect no part of the existing legal patchwork was designed to address, since intellectual property law, agency law, and consumer protection law all take the persona's existence as a given and ask only how it should be regulated once created, never whether its economic function raises a distinct policy question.

Shudu, created by photographer Cameron-James Wilson in 2017 and marketed as “the world's first digital supermodel,” generated a controversy the other three cases do not raise. Wilson, who is white, built a hyperrealistic Black female persona and, in some early posts, did not clearly disclose that she was computer-generated at all, leading some viewers to believe they were looking at a real Black model and drew sharp criticism once the truth came out, since a white creator profiting from a fabricated Black persona echoed a familiar pattern of appropriation with no actual person able to consent, benefit, or object.²⁴ This matters for the framework proposed below because it shows that disclosure obligations are not only about preventing consumers from being misled about a product's authenticity. They also touch questions of representation and consent that existing publicity-rights law has never had to confront, since no living person's image or likeness was technically implicated in a way current statutes could reach. A registration and disclosure regime would not fully resolve what Shudu raised, but it would at least make the originating entity and the creative decisions behind a persona visible from the outset, rather than leaving audiences to piece the story together after the fact.

Bermuda, a second Brud-created character positioned as Miquela's rival within the same fictional universe, was used to stage a public “hacking” storyline in which Bermuda claimed to have taken over Miquela's account - a scripted stunt initially presented to followers as a real security incident.²⁵ Some followers reported feeling genuinely deceived once the stunt was revealed as fabricated content. Nothing about it was illegal, since no statute requires a company to be honest

²³ 'Aitana López' (*Wikipedia*) <https://en.wikipedia.org/wiki/Aitana_L%C3%B3pez> accessed 15 July 2026; 'Meet the Spanish AI Model Earning up to €10,000 a Month' (n 2).

²⁴ 'Shudu Gram' (*Wikipedia*) <https://en.wikipedia.org/wiki/Shudu_Gram> accessed 15 July 2026.

²⁵n 1.

about whether an in-universe plot twist is real, but it illustrates a softer harm the framework below only partially reaches: audiences increasingly relate to virtual influencers as though following a continuous, real personality, and companies have already shown a willingness to exploit that parasocial trust for engagement, independent of any AI involvement at all.

7. A WORKABLE FRAMEWORK

None of the conclusions above requires granting a virtual influencer any form of personhood, in the human rights sense or the corporate sense. What they require is a narrower, functional status attached to the company standing behind her, built around four elements.

First, mandatory registration and disclosure of the controlling entity. Any commercially deployed AI or CGI persona above a defined threshold of reach i.e. a follower count or advertising spend chosen to spare hobbyists should be required to register the legal entity responsible for its content, publicly and kept current. This solves the Miquela ownership-transfer problem directly: had Dapper Labs been required to file an updated registration the moment the Brud acquisition closed, followers, regulators, and potential plaintiffs would have had a stable, public answer to who was accountable. The mechanism is not novel; it borrows from the model already used for broadcast license holders and, in the United States, the FCC's sponsorship identification rules for radio and television.

Second, a codified endorsement-liability rule specific to synthetic personas. Rather than leaving the question to interpretive FTC guidance, legislatures should state plainly that the controlling entity behind an AI or CGI influencer is strictly liable for the accuracy and legality of any endorsement or claim the persona makes, regardless of whether a human reviewed that particular output before publication. This removes the “the AI said it, not us” defence before it can be raised, mirroring the strict-liability logic already used in product liability, where a manufacturer cannot escape responsibility for a defect by noting that no individual employee personally approved it.

Third, a labelling requirement attached to the content itself, not just to the account. A one-time “this account is AI-generated” note in a bio profile is not sufficient, given how content circulates independently of its source (a screenshot, a reposted clip, or an embedded advertisement can travel far from the original disclosure). Following the general transparency logic of the EU AI Act's Article 50, but applying it specifically to commercial personas, individual pieces of sponsored content from an AI or CGI influencer should carry a persistent, machine-

readable marker identifying it as synthetic, so the disclosure travels with the content itself.²⁶

Fourth, a narrow liability shield for genuinely unscripted, real-time interaction, conditioned on real safeguards. There is a reasonable argument that a company deploying an autonomous, real-time conversational persona should not face the same strict liability for spontaneous chat responses that it faces for pre-approved marketing copy, since holding it to that standard risks pushing the industry back toward fully scripted personas and eliminating the interactivity that gives these systems much of their value. A workable compromise, similar in spirit to safe-harbour provisions used elsewhere in technology regulation, would extend limited protection for unscripted content only where the company discloses clearly, at the point of interaction, that responses are AI-generated and not pre-reviewed; implements and can demonstrate reasonable content-moderation safeguards; and retains logs sufficient to reconstruct what was said, to whom, and when, for a defined retention period. Fail any of the three conditions, and full strict liability under the second element applies.

None of the four elements matters much without a plausible enforcement path. Registration could sit with existing consumer protection agencies like the FTC in the United States, national authorities under the EU's Unfair Commercial Practices Directive framework, and their counterparts elsewhere rather than requiring an entirely new regulator, since these agencies already handle disclosure and advertising-claims enforcement for human endorsers. Platforms hosting monetised accounts above the reach threshold could be required to collect and verify registration information as a condition of enabling monetisation, shifting compliance onto entities that already run identity checks for payment processing. Penalties should scale with reach and revenue rather than sitting at a flat rate, both because a flat fine is a rounding error for a well-funded studio and business-ending for a two-person shop, and because the harm an unregistered persona can do scales with its audience in the first place.

None of the four elements would be entirely novel in isolation, which is itself an argument for feasibility rather than a mark against originality. The EU AI Act's transparency obligations already require disclosure when a person is interacting with an AI system, though the provision is broad enough to cover chatbots and voice assistants generally, leaving real interpretive uncertainty about whether a livestreaming commercial persona clearly falls within scope.²⁷ South Korea's Act on Fair Labelling

²⁶ Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence (Artificial Intelligence Act) [2024] OJ L1689, art 50.

²⁷ Regulation (EU) 2024/1689 laying down harmonised rules on artificial intelligence

and Advertising has already produced regulatory guidance suggesting virtual influencer content should be clearly marked, reflecting the country's position as one of the largest markets for the format.²⁸ China's Cyberspace Administration goes furthest, requiring clear labelling of AI-generated and deep-synthesis content under rules issued in 2022 and 2023, though enforcement appears to vary by platform.²⁹ The United States has moved least: FTC guidance remains interpretive rather than codified for synthetic personas specifically, leaving American law to rely on the same general-purpose consumer protection and defamation doctrine this paper has argued is a poor fit for the problem. Taken together, the elements of a workable regime already exist in scattered form across different systems; what is missing is a single jurisdiction assembling them into a coherent, influencer-specific regime rather than leaving the pieces distributed across broader statutes not written with this particular problem in mind.

8. OBJECTIONS AND RESPONSES

The proposal invites several objections worth taking seriously rather than waving off.

The first is that this amounts to regulatory overreach for what is mostly harmless marketing. Most virtual influencer activity is indeed harmless, and nothing in the framework restricts what a scripted, reviewed persona can say or do; it only requires that someone accountable be identifiable and that liability attach clearly when something goes wrong. The heavier obligations bite hardest on the riskiest category, unscripted real-time AI, and barely touch the low-risk scripted category at all.

The second is that registration requirements will simply be evaded through shell companies and offshore entities, the same way disclosure regimes are evaded elsewhere. That is true, and worth taking seriously; no registration scheme is evasion-proof. The same objection applies to ordinary corporate registration, and the answer is the same one used there: pair registration with meaningful penalties for platforms that host unregistered commercial personas above the reach threshold, shifting some enforcement burden onto platforms that already have far better visibility into account ownership than any regulator does.

(Artificial Intelligence Act) [2024] OJ L1689, art 50.

²⁸ Act on Fair Labelling and Advertising (Republic of Korea); see Korea Fair Trade Commission guidance on virtual influencer disclosure (2023).

²⁹ Cyberspace Administration of China, *Provisions on the Administration of Deep Synthesis Internet Information Services* (2022); Cyberspace Administration of China, *Interim Measures for the Management of Generative Artificial Intelligence Services* (2023).

The third is that strict liability for AI-generated statements will chill innovation and push smaller studios out of the market, since only well-capitalised companies can absorb the risk. That cost is real and should be weighed honestly. But the same argument was made, and largely rejected, when strict product liability was extended to manufacturers in the twentieth century; the concern that small manufacturers could not absorb liability risk proved manageable once insurance markets developed specifically to price and spread it. A codified rule known in advance is exactly what insurers can underwrite; the ambiguous, untested exposure studios currently face is arguably harder to insure against than a clear strict-liability standard would be, precisely because insurers struggle more to price uncertainty than to price a known, defined risk.

The fourth is that none of this addresses the labour-displacement concern Aitana López's case raises. That is a fair gap, and this framework does not close it, because accountability and disclosure are a different problem from labour policy and this paper has deliberately not tried to solve both at once. Whether AI personas that functionally substitute for human models should trigger some form of labour protection or licensing fee is a genuinely separate question, closer to the broader debate already underway about AI's effect on creative labour, and it deserves treatment on its own terms rather than being bolted onto a liability framework where it does not fit.

9. CONCLUSION

The instinct to ask whether AI virtual influencers deserve personhood is understandable, and it borrows a category built for a different set of problems and applies it somewhere it does not quite fit. Corporations acquired legal personhood because doing so solved administrative problems about continuity, contracting, and liability that arise from being a persistent enterprise with its own assets. Human rights instruments extend recognition before the law to “everyone” on a dignity-based premise that has nothing to do with administrative convenience and everything to do with rational autonomy, a premise even a cognitively complex elephant has so far failed to satisfy in the eyes of a court willing to take the argument seriously. Virtual influencers sit in neither category. They do not have a continuity problem, because they own nothing and outlive no one. They do not have a dignity claim, because nothing about a marketing persona suffers, prefers, or has interests capable of being frustrated.

What they have instead is a far more mundane gap: nobody is clearly, persistently, and legally required to stand behind what they say and do, and the humans who profit from that vagueness have every incentive to

keep it in place. Fixing that does not require inventing digital citizens. It requires registration, disclosure, and liability rules precise enough that when Lil Miquela's next owner takes over, or an AI-driven persona says something false or damaging on a livestream at two in the morning, there is an answer on record, in public, before anyone has to go looking for it, to the only question that actually matters: who is responsible. The technology will keep getting better at making these characters feel present, distinct, almost alive. The law's task is not to decide whether they are. It is to make sure someone real is always standing behind them, whatever they seem to be.