



SEPTEMBER 10, 2023 / JÉRÔME DE COOMAN

THE ECONOMIC CONSEQUENCES OF GENERATIVE AI ON THE ART MARKET

When an art buyer cannot establish with certainty the origin of the work, namely when there is a risk that the purchase transaction concerns a forgery and not a genuine artwork, the buyer adjusts her willingness-to-pay based on the average probability that the artwork is authentic. Given that the artistic outputs of generative artificial intelligence system are indistinguishable from human-made ones, this blogpost argues there is a risk that the latter will be devalued. This would be especially true given that the market value of machine-enabled artwork is expected to be lower than the human-made one. Art authentication could solve the problem, but law- and policymakers should be aware that such rule of origin cannot be dyadic given that some machine-enabled artworks owe their existence more to the human user of GenAI than others.

With the advent of generative artificial intelligence (GenAI) come tremendous economic consequences. This blogpost does not discuss whether and how GenAI will disrupt the working world. Rather, it focuses on the market for artworks and argues that GenAI leads to George Akerlof's lemons problem (1)—albeit slightly modified to encompass the idiosyncrasies of the art market (2)—because human-made and machine-enabled artworks are indistinguishable (3). The blogpost ends by recommending bottom-up and top-down rules of authentication (4).

1. The original lemons problem

Imagine that on the market for used cars there are only two different types of products: the good-quality used cars (in American slang, the peaches) and the bad-quality ones (the lemons). Being on the market of *used* cars, sellers know whether the car being sold is a lemon or not. Buyers do not. This asymmetric information incentivises dishonest sellers to sell a lemon at the price of a peach. Buyers, in turn, will value a particular car based on the average probability that that car is a peach (assuming they are risk neutral) and adapt their willingness-to-pay accordingly. Therefore, the market price equilibrium will systematically be lower than what a peach is worth, but higher than what a lemon is worth. Sellers of peaches will, consequently, either exit the market or voluntarily decrease the quality of their products (Arrowian moral hazard).

This is a market inefficiency par excellence: asymmetric information forces buyers to misallocate their resources, resulting in a systemic inefficiency that harms both buyers' and other sellers' welfare.

2. Asymmetric information on the art market

First and foremost, discussing asymmetric information on the art market requires a distinction between aesthetic and market values. The fluctuating valuation of Leonardo da Vinci's masterpiece Salvator Mundi provides an interesting background. In the 1950s, it was sold as a copy of the original, completed by one of da Vinci's students for approximately USD 60 (approximately USD 755 today). After its authentication as a

genuine Da Vinci, it was auctioned for USD 450.3 million in 2017 (approximately USD 557.3 million today). However, experts expressed doubts afterwards and the painting has not been exhibited since. Authentication did not change the aesthetic value of *Salvator Mundi*, but did multiply its market value by more than seven hundred thousand—before disappearing from the museum's walls. Therefore, aesthetic value is not the prime factor of market value. Authentication is. Thus, a lemon is not necessarily a low aesthetically value artwork, but one that is not genuine (*i.e.*, faked or forged).

Without authentication, the buyer cannot know whether she acquires an original artwork or a forgery, whereas the seller usually already possesses this knowledge—assuming she is the forger or knows from the forger that the artwork is forged. Therefore, the buyer will revise her willingness-to-pay to take into account the average probability that the artwork is forged. The upshot? Sellers of genuine artworks cannot receive the true value of their product, but forgers can dishonestly sell forgeries at a price that largely exceeds their worth. Again, asymmetric information leads to inefficient market allocation and eventually to the potential bankruptcy of art dealers.

3. GenAI and the lemons problem

This blogpost does not investigate whether GenAI can be malevolently used to flood the art market with forgeries. Rather, it questions whether users have an incentive to sell a machine-enabled artwork as a human-made one. It seems like this is the case.

On the one hand, they are indistinguishable. The case of Generative Adversarial Networks (GANs) is particularly enlightening in this regard. In a nutshell, a GAN is composed of two competing sub-models that are simultaneously trained. A generative model G generates new data (output) based on simple random variables (input). A discriminative model D then estimates the probability that the output is real (*i.e.*, belongs from the training dataset) or is generated by G . These models are called adversarial because they play a zero-sum game: either the discriminator correctly classifies the examples as real or generated, or alternatively fails and, *ergo*, the generator successfully fools the discriminator. Eventually, the discriminator is unable to classify the examples it has to label and reaches

a fifty percent success rate (*i.e.*, the same accuracy as predicting heads or tails when flipping a coin). Indistinguishability cannot be better defined.

On the other hand, the market value of human-made and machine-enabled artworks under symmetric information might well be different for a threefold reason. First, it has been demonstrated that beholders tend to rate machine-enabled artwork (disclosed as such) at a lower aesthetic level than human-made ones (perception bias). Second, auctions of machine-enabled artwork at prices that exceed expert evaluation seem to be drying up. Although the hammer price of *The Portrait of Edmond de Belamy* was more than 43 times its upper estimate, subsequent auctions were far less sensational: at best, they were auctioned at the upper estimate; at worse, not auctioned at all. One reason might be that the announcement effect is over. Third, machine-enabled artworks might not be copyrightable. This has an economic impact. Copyright is an exclusive property right that grants its owner (the author or a third party to whom the latter would have transferred the right) a temporary monopoly. Without copyright, the remuneration of the author is reduced to a trickle. Actually, this was at the core of the *Painer* ruling by the CJEU. Freelance photographer Ms. Painer neither conferred any right over portrait photographs she took, nor consented to their publication. When they were nevertheless published, she claimed a copyright and requested “payment of appropriate remuneration and damages for her loss” (§38). The uncopyrightability of machine-enabled artwork therefore means that the remuneration of GenAI users will be limited.

4. Conclusion and recommendations

Given that GenAI introduces new indistinguishable products on a market that was already confronted with a lemons problem, the solution is intuitive. It suffices to expand the authentication rule to machine-enabled artworks. So far, authentication distinguished genuine from forged artworks. From now on, it can similarly be used to draw a line between human-made and machine-enabled artworks.

However, authentication is currently binary: either the work is genuine or it is not. Transposed to GenAI, this suggests a distinction between human-made (GenAI-free) and machine-enabled artworks, as soon as GenAI is used in the creative process. This is likely inappropriate. Such a dyadic

authentication would lead to an equivalent consideration of the output of, e.g., a text-to-image GenAI when prompted with one single sentence and a work of art that required hundreds, if not thousands of detailed instructions. The latter piece owes its existence more to the human user than the former. Should they still be both classified as machine-enabled?

Geographical rules of origin that establish the provenance of goods and services for tax purposes proposes a solution. Pursuant to the Union Customs Code, "goods the production of which involves more than one country or territory shall be deemed to originate in the country or territory where they underwent their last, substantial, economically-justified processing or working". The CJEU emphasised that an operation is substantial if "the product resulting therefrom has its own properties and a composition of its own, which it did not possess before" (§6). Transposed mutatis mutandis to machine-enabled artwork, this substantial transformation test suggests that (a) an artwork may owe its existence to both human and machine, (b) the fact that a human edited the artwork before being made available to the public does not mean it is human-made, and (c) that the artwork will only be labelled as human-made if the human input is substantial vis-à-vis the role of the machine.

In this regard, it is worth emphasizing that the still-under-discussion AI Act contains an embryonic top-down rule of origin limited, however, to deep fake (art. 52(3)). But a rule of origin could also be achieved bottom-up by the art community itself. In economic jargon, this is known as signalling: artists literally signal to potential buyers the quality (or origin) of the product. This can be done by artists filming themselves while creating the artwork. Actually, photography was used to prove the authorship of the Tiara of Saïtapharnès in 1896. Unfortunately, records of book writers hankered over their keyboards while typing would prove absolutely nothing. Against that background, both OpenAI, academic publishers and a professional organisation for published writers required that the use of GenAI should be acknowledged in an ad hoc section. That said, given that indistinguishability incentivises dishonest dealings, such voluntary reporting scheme are not a *credible* signal if based solely on the authors' goodwill. However, a technical solution exists to make such disclosure trustworthy. It is perfectly conceivable to design generative AI systems in such a way that they watermark the generated output as such.

Either top-down or bottom-up, a rule of origin has at least one decisive benefit. It allows those who value human-made artworks more than machine-enabled ones to buy the product at the correct price. Yet, it has

absolutely no impact on buyers who do not care about the artwork's provenance as long as it is of high aesthetic quality. A rule of origin does nothing but introduce greater transparency to ensure symmetrical information between buyers and sellers.

Published under licence CC BY-NC-ND.

HI, I'M JÉRÔME DE COOMAN

Jérôme De Cooman is Ph.D. candidate and teaching assistant at University of Liege since 2018 in EU Competition Law and EU law, (Big) Data and AI Applications courses. Since 2021, he is Administrative Manager at the Brussels School of Competition, Junior Editor for the Yearbook of Antitrust and Regulatory Studies (YARS) at the Centre for Antitrust and Regulatory Studies (CARS, University of Warsaw, Poland), and Junior Member of the Academic Society for Competition Law (ASCOLA).



**THIS BLOGPOST WAS WRITTEN BY
AUTHOR**



JÉRÔME DE COOMAN

JÉRÔME DE COOMAN IS PH.D. CANDIDATE AND TEACHING ASSISTANT AT UNIVERSITY OF LIEGE SINCE 2018 IN EU COMPETITION LAW AND EU LAW, (BIG) DATA AND AI APPLICATIONS COURSES. SINCE 2021, HE IS ADMINISTRATIVE MANAGER AT THE BRUSSELS SCHOOL OF COMPETITION, JUNIOR EDITOR FOR THE YEARBOOK OF ANTITRUST AND REGULATORY STUDIES (YARS) AT THE CENTRE FOR ANTITRUST AND REGULATORY STUDIES (CARS, UNIVERSITY OF WARSAW, POLAND), AND JUNIOR MEMBER OF THE ACADEMIC SOCIETY FOR COMPETITION LAW (ASCOLA).

[VIEW ALL POSTS](#)

PREVIOUS POST

KANN DER WAS? EIN
KRITISCHER BLICK AUF DIE
LETZTEN
ÄNDERUNGSVORSCHLÄGE FÜR
DEN AI-ACT

NEXT POST

AI ACT AND GDPR: ON THE
PATH TOWARDS OVERLAP OF
THE ENFORCEMENT
STRUCTURES

CATEGORIES

BLOG-BEITRÄGE

RAILS - BLOG

Robotics and AI Law Society

[ALL BLOGPOSTS](#)
[RAILS WEBSITE](#)
[ABOUT](#)
[LEGAL DISCLOSURE & PRIVACY POLICY](#)



AUTHORS LIST

Default Search ▾

Search

Agnieszka Jabłonowska, LL.M. (EUI), Amir Cahane, Ann-Kristin Mayrhofer, Anna Urbanek, Anton Kamke

1 **2** **3** ... **18** **NEXT »**

RAILS - ROBOTICS AND AI LAW SOCIETY, 2020