



UNIVERSITY
OF WARSAW



FACULTY OF
ECONOMIC SCIENCES

WORKING PAPERS

No. 16/2026 (510)

TWO BLIND WALLS: A TRANSFERABLE PEDAGOGICAL DESIGN FOR ART-MARKET LITERACY ACROSS ART AND ECONOMICS EDUCATION

TOMASZ KOPCZEWSKI
JUSTYNA LASKOWSKA-LISICKA
JAN LISICKI
KOSTIANTYN OKHRIMENKO
TOMASZ POTOCKI

WARSAW 2026

ISSN 2957-0506



Two Blind Walls: A Transferable Pedagogical Design for Art-Market Literacy Across Art and Economics Education

Tomasz Kopczewski^{1}, Justyna Laskowska-Lisicka², Jan Lisicki¹, Kostiantyn Okhrimenko¹, Tomasz Potocki³*

¹ *University of Warsaw, Faculty of Economic Sciences*

² *Adam Mickiewicz Institute*

³ *University of Rzeszów*

* *Corresponding author: tkopczewski@wne.uw.edu.pl*

Abstract: This article presents an ongoing pedagogical project at the intersection of economics and art education. It starts from a double blind wall: economics rarely uses the art market as a teaching laboratory, while art education often protects aesthetic judgement from market reductionism. We propose a transferable protocol combining diagnostic questions, valuation tasks, anchoring and signalling treatments, public-art perception probes, and dashboard feedback. Its first implementation uses Wojciech Fangor in Warsaw. The aim is not representative measurement, but a replicable design through which students observe how attention, information, narrative, and market signals shape the encounter with art.

Keywords: art valuation, economics education, art education, anchoring, signalling, asymmetric information, pedagogical narrative laboratory, Know Thyself, Wojciech Fangor, replication

JEL codes: A22, A29, C93, D83, D91, I23, Z11

Acknowledgements: We thank Robert Borowski for the technical implementation of the survey on the ProfiTest platform; Ewa Weychert for consultation in the early phase of project design; participants in the open open seminar "Know Yourself Through Economic Theories" held at the Faculty of Economic Sciences, University of Warsaw, on 24 January 2026, for their feedback during the design phase; and the students -participants of the 2026 surveys, who agreed to make their intuitions about value visible. Errors of fact, interpretation, and judgment remain our own.

1. Introduction

In central Warsaw, three sites lie within a few minutes' walk of one another. At the Uniwersytet station of the second metro line, students enter and leave the city center through walls designed by Wojciech Fangor. A short walk south, at the Warszawa Śródmieście railway station, opened in 1963, a series of mosaics he completed for the building's public hall is slowly degrading. Between the two, alongside the Museum of Modern Art, a recently named street - *ulica Wojciecha Fangora* - was unveiled in October 2024. Fangor (1922–2015) was the only Polish artist to have had a solo exhibition at the Solomon R. Guggenheim Museum, and his paintings now sell at Polish auctions for sums above seven million złoty. His name, in this part of the city, is everywhere. His work is mostly unseen.

This article begins from that paradox, but it is not primarily an article about Fangor. Fangor is the Warsaw implementation of a broader pedagogical design: a transferable sequence of valuation tasks, signalling treatments, attention probes, and feedback devices through which students can examine how artworks are seen, ignored, priced, narrated, and institutionalised. The design was built for use on both sides of a disciplinary divide. For economics educators, it offers a classroom laboratory for information asymmetry, anchoring, signalling, reputation, and attention. For art and design educators, it offers a way to address the market-mediated life of artworks without reducing artistic value to price. The same protocol can be re-run with another artist, in another city, and within another educational tradition.

The need for such a design comes from a double pedagogical omission. Economics education has powerful tools for analysing uncertainty, market signals, asymmetric information, and bounded attention, but it rarely uses the art market as a central teaching case. Art education, by contrast, has historically maintained a principled distance from market valuation, often for good reasons: to protect the autonomy of aesthetic judgement, to resist commodification, and to keep open forms of attention not organised by price. Yet this distance also has a cost. Students may learn to look at artworks without learning how reputational signals, institutional validation, auction records, scarcity, provenance, and visibility shape the conditions under which artworks circulate. Economics often fails to see art as a rich pedagogical field for its own theories; art education often avoids the economic grammar through which artworks acquire public and market life. These are the two blind walls of our title.

We do not treat these blind walls as simple failures. They are structured by the histories and values of the disciplines themselves. The reluctance of art education to speak in the language of markets is not a mistake; it is partly a defence against reductionism.

The reluctance of economics education to teach through art is also understandable; the art market resists the clean inferential moves that economists often prefer. But the educational consequence is similar on both sides. Students are left without a shared vocabulary for describing objects that are at once visual, symbolic, institutional, attentional, and economic.

To make this gap visible, we implemented the design in Warsaw through a set of experiments based on the work of Wojciech Fangor. In the main study, more than a hundred participants - students from different faculties of the University of Warsaw and members of the public attending the open seminar "Know Yourself Through Economic Theories" - completed an online questionnaire containing diagnostic questions and experimental valuation tasks. The instrument is in continuous use: each new cohort enters the same Shiny dashboard under a session-token. For the analyses reported here, the dataset was frozen at $N = 131$ completed responses. Later cohorts enter the dashboard for teaching and future comparison, but are not included in the present results. We treat the findings not as population-level estimates, but as diagnostic signals for educational design.

Three findings are especially important. First, when the same respondent was asked how much they would pay for a contemporary painting they liked, and then how much they would pay for the same hypothetical painting if treated purely as a financial investment, the investment figure was on average nearly seven times higher than the aesthetic figure. For roughly one in seven respondents, the ratio was ten times or more. Second, when respondents were shown a painting and given a numerical price suggestion before estimating its value, their valuations anchored strongly on that number, even when the anchor had no informational relationship to the artwork. Third, when asked to name young promising Polish artists, sixty-six percent answered that they could not name any (95% CI [57%, 74%]). Many of those who did answer supplied either deceased masters or internationally recognisable brands rather than young contemporary Polish practitioners.

Taken together, these findings describe not simply a lack of knowledge, but a particular cognitive infrastructure. Respondents were able to produce prices, but often lacked names, categories, historical context, and stable points of reference. They could respond to signals, but not always distinguish between aesthetic judgement, market reputation, institutional validation, and speculative expectation. This is why the design matters pedagogically. It does not tell students what art is worth. It lets them observe how their own valuations move when information, price anchors, biographical signals, curatorial language, and public visibility are introduced or withheld.

The design consists of several linked elements. A diagnostic battery asks respondents about their familiarity with art movements, including a fictitious movement used to detect declarative inflation. A dual willingness-to-pay question separates aesthetic value from investment value. An anchoring experiment tests the effect of arbitrary numerical suggestions. A signalling task presents biographical and curatorial information after the anchor. A public-art perception probe tests whether physical exposure translates into conscious visual encounter. Finally, a personalised feedback loop returns the results to participants through a pseudonym-based dashboard, allowing them to compare their own responses with the distribution of the group. The survey is therefore not only a measurement instrument. It is also a pedagogical device.

This article contributes neither a new theory of the art market nor a representative survey of public knowledge about contemporary art. Its contribution is a replicable educational protocol. We describe one implementation of that protocol in Warsaw, using Fangor as the local case. We then show how the same structure can be adapted elsewhere: with different artists, different public-art sites, different price scales, and different institutional signals. The protocol functions as a boundary object. It is structured enough to be repeated, compared, and analysed; but open enough to be interpreted differently by economics educators, art educators, curators, and students.

We write as economists working at the border of art education, and therefore from outside part of the field we address. Our aim is not to instruct art educators in their own discipline, nor to argue that economic valuation should replace aesthetic judgement. The opposite is closer to our intention. We propose a design through which both sides can examine what they usually bracket out. Economists can see how fragile price formation becomes when symbolic and institutional context dominate physical attributes. Art educators can see how strongly market signals enter the conditions under which artworks are encountered, remembered, ignored, and valued. The purpose is not to dissolve the difference between art and economics, but to create a door in the wall between them.

The article proceeds as follows. Section 2 locates the double pedagogical gap in economics education and art education. Section 3 introduces Fangor as the local implementation case and explains the choice of works and sites. Section 4 presents the three-layer framework used to distinguish physical attributes, visual-aesthetic features, and narrative-reputational capital. Section 5 describes the pedagogical design and the method. Section 6 reports four diagnostic stories. Section 7 presents three experimental cuts: anchoring,

the portrait test, and public-art invisibility, with a methodological note on what we did not randomise and a coda on convergent evidence from machine learning. Section 8 sets out the transferable protocol and the open call for replication. Section 9 offers a reflexive note on the authors' surprises, limitations, and a closing return to Warsaw.

2. Two blind walls: locating the gap

It is worth correcting a stereotype before we proceed. Contrary to a common image of the discipline - sustained, in part, by the very critical tradition we engage with in this section - economists do not, on the whole, treat the market as a source of truth or as a machine for generating correct value. As Mirowski (2002, 2013) has argued, that image is partly a Cold-War residue: a public simplification of a discipline whose everyday working culture is often more interested in market failure than in market success. Much of the most visible economics of the last forty years - Akerlof on lemons, Spence on signalling, Stiglitz on screening, Kahneman on judgement, Thaler on nudges - is a catalogue of situations in which markets fail to do what they appear to do. Economists, in this respect, are not only theorists of efficiency. They are also scholars of malfunction.

Inefficiencies of valuation, signalling, and information asymmetry are routinely taught at the undergraduate level. Yet they are rarely taught through the art market, despite the fact that this market offers one of their most concentrated empirical forms. Standard introductory courses present anchoring through Tversky and Kahneman's roulette wheel, signalling through Spence's labour market, asymmetric information through Akerlof's used cars. This impression is corroborated by external content analysis: a recent automated lexical study of major introductory textbooks shows that the word *art* appears densely across chapters of standard sociology textbooks (Schaefer, Macionis, Andersen) and only sparsely, if at all, in mainstream microeconomics - Mankiw, Krugman and Wells, Arnold, Hubbard and O'Brien (Letycja, 2026). Cultural economics exists as a developed subfield (Throsby, 1994; Velthuis, 2005; Beckert & Rössel, 2013), and there is a robust empirical literature on anchoring at art auctions (Beggs & Graddy, 2009) and on art as an asset class (Renneboog & Spaenjers, 2013). Almost none of this enters the ordinary teaching repertoire of microeconomics. Becker and Watts (2001), and Walstad and Siegfried (2015) two decades later, describe an undergraduate economics curriculum that remains narrow not only in its formal apparatus, but also in its examples.

The omission on the other side has a different character. It is not incidental, but principled. Many influential traditions within art and design education have taken the autonomy

of the artistic field as a starting condition (Bourdieu, 1993, 1996), maintained a critical distance from the commodification of culture (Adorno, 1991), and inherited from Read and Eisner a vocabulary in which aesthetic value cannot be reduced to instrumental measure. In contemporary debates, this suspicion of market capture remains alive. Within and around iJADE, Jagodzinski (2010, 2015, 2020) has repeatedly warned against the absorption of creativity into designer capitalism. These positions are not naive, and we do not treat them as something to be dismissed. They are well-founded responses to real pressures on the field. Their cost, however, is educational. Students may learn to make, interpret, and critique artworks without acquiring a vocabulary for the market life of the objects they study.

Placed side by side, the two omissions describe a single situation from opposite directions. Economics has tools for analysing uncertainty, reputation, signalling, attention, and price formation, but rarely uses art as a central pedagogical field for those tools. Art education has strong reasons to protect aesthetic judgement from market reductionism, but this protection can also leave students underprepared to understand how artworks acquire prices, narratives, institutional legitimacy, and public visibility. The art market is therefore not a place where economists arrive with a tidy model and discover that art breaks it. It is a place where the problems economics already studies appear with unusual clarity, but in a form that cannot be understood by economics alone.

This is the gap our design addresses. We do not propose to merge art education and economics education into a single language. We propose a boundary protocol: a structured but adaptable sequence of tasks through which educators on either side can examine the same object differently. For economics educators, the protocol turns artworks into cases of asymmetric information, anchoring, signalling, attention, and reputational capital. For art and design educators, it turns market valuation into an object of critical reflection rather than a replacement for aesthetic judgement. The protocol is therefore not a market lesson disguised as art education, nor an art appreciation exercise disguised as economics. It is a shared instrument for making visible what each discipline tends to bracket out.

The first concrete site at which both blind walls meet, in our own data, is what we call the *names problem*. When we asked our 131 respondents - students from different faculties of the University of Warsaw and participants in the open seminar - to name young promising Polish artists, sixty-six percent answered that they could not name any (95% CI [57%, 74%]). Most of those who did respond supplied either deceased masters or established international brands. From an economics-of-information perspective, this is a clean case: when respondents

lack recognised signals for the cohort the question asks about, they fall back on names with accumulated reputational capital. From an art-education perspective, it is a different problem with the same structure: the contemporary scene remains invisible to much of the public it seeks to address. The fact and the interpretation are coextensive. A student able to read the situation through both lenses would understand more than either discipline can show alone.

This section has not attempted to construct a new theory of either field. It does not settle whether market analysis is an intrusion into aesthetic education, nor does it argue that economics curricula owe art a central place. Its purpose has been narrower: to locate one point at which two defensible disciplinary habits leave students without the vocabulary needed to describe what is in front of them. The remainder of the article reports an implementation of a pedagogical design built for that point: a sequence of diagnostic questions, valuation tasks, signalling treatments, public-art perception probes, and feedback devices that can be reused with another artist, in another city, and in classrooms on either side of the wall.

3. Why Fangor: an art-historical case for the experimental stimulus

The choice of artist for the Warsaw implementation of the protocol was not made on aesthetic grounds. Fangor was chosen because he simultaneously satisfies four conditions that the design requires of any artist around whom it is to be run, and he satisfies them in a configuration unusual enough to make the result of the experiment legible. The conditions are: a recognisable name; a body of work whose visual surface has been institutionally legitimised but is unevenly known to the public the protocol is asking; a market presence with auction records that allow the numerical anchors to be calibrated against real transactions; and physical presence in the city's everyday environment in a form that the local population passes through without being asked to attend. The case for Fangor on each of these points was first developed at length by Laskowska (2025), on whose biographical and market account the present section draws.

Wojciech Fangor (1922–2015) trained at the Academy of Fine Arts in Warsaw under conditions of post-war reconstruction. His early figurative period - to which the two portraits used in the experiment of Section 7.2 belong - was painted in oil between 1948 and 1949, before the imposed turn to socialist realism that produced his canonical works of the early 1950s (*Postaci*, *Lenin w Poroninie*, *Matka Koreanka*). He moved through poster design (the Polish School of the Poster) and a brief experimental abstract phase (*Studium przestrzeni*, 1958, with Stanisław Zamecznik - three-dimensional installation work that anticipated, by several

years, similar pieces by Robert Morris in New York). In the early 1960s he completed the mosaics for the public hall of the Warszawa Śródmieście railway station that we describe in Section 7.3. He left Poland in the mid-1960s for Vienna and then the United States, where he developed the op-art series of soft-contour circles and waves for which his international market is now built. In 1970 he received a solo exhibition at the Solomon R. Guggenheim Museum in New York - the only Polish artist to have done so. He returned to Poland in the late 1990s. His final commissioned work was the visual identity for seven stations of Warsaw's second metro line, completed posthumously in 2022. In October 2024 a street in central Warsaw was named after him, alongside the new Museum of Modern Art.

His market is concentrated in two segments. The op-art canvases - to which M78 (1968), used in the signalling task, belongs - anchor the upper end: the painting M22 sold at DESA Unicum in 2020 for 7.3 million złoty, the auction record. The early figurative period, less visible to the contemporary art-historical narrative that built the op-art reputation, is priced selectively: the 1948 portrait used in Section 7.2 reached approximately 600,000 PLN, while the 1949 portrait - visually adjacent - reached approximately 64,000 PLN. The factor-of-ten gap between two superficially similar works by the same hand, in the same medium, in consecutive years, is the kind of price discontinuity that the Layer 2 / Layer 3 distinction described in Section 4 is designed to make legible.

The four works selected for the experimental tasks (the 1972 abstract used as the anchoring stimulus, the 1948 and 1949 portraits used in the paired comparison, and M78 used in the signalling task) cover three of Fangor's stylistic periods and the full range of his market: the early hand most viewers cannot place stylistically, the canonical op-art object that defines his contemporary recognition, and an abstract from the early 1970s whose price respondents are asked to estimate without scaffolding. A replication elsewhere does not require a Fangor; it requires an equivalent - an artist whose presence in the city's everyday physical environment is similarly mismatched against the recognition the public would give if asked.

4. The three-layer framework

The three experimental cuts in Section 7 work, in different ways, on different layers of what determines the price of a painting. To make the cuts legible we need a framework that distinguishes those layers cleanly. We adopt the one developed in detail by Okhrimenko (forthcoming), which integrates three established economic and sociological literatures into a single structure: hedonic pricing on the physical layer (Rosen, 1974), aesthetic theory on

the visual layer, and the sociology of cultural fields on the narrative-reputational layer (Bourdieu, 1993; Velthuis, 2005; Beckert & Rössel, 2013). The framework is empirical rather than conceptual: it organises the price-relevant features of an artwork by where they live and how they are observable to a buyer, a researcher, or a model.

The first layer is physical. It comprises the measurable attributes of the object: dimensions, medium, age, signature, condition, provenance documentation. These are the variables that hedonic pricing models since Rosen (1974) have used to decompose price into the implicit valuations of underlying characteristics. Czujack (1997), studying Picasso, found that surface area was among the strongest predictors of price; Schönfeld and Reinstaller (2007) confirmed the effect in primary-market data. The physical layer is the layer most accessible to ordinary observation and the one that machine-learning models, including those discussed in Section 7.5, can readily detect. It is also the layer that explains the smaller part of the variance in actual transaction prices.

The second layer is visual-aesthetic. It comprises features that are in principle observable on the surface of the work - composition, palette, gestural quality, formal coherence - but resist quantification by either human or machine. Two viewers may agree that a painting is *successful* or *resolved* without being able to articulate the criterion; convolutional neural networks trained on artwork images, as Greene (2020) and Lee et al. (2024) report, capture some of this layer but only weakly. The visual-aesthetic layer is the one on which non-expert valuation most often relies when no other information is available, and it is the layer that Section 7.2's portrait test directly probes by presenting two visually adjacent works whose market values diverge by an order of magnitude.

The third layer is narrative-reputational. It comprises everything that does not live in the canvas and does not live in the listing: artist biography, institutional affiliations, exhibition history, gallery prestige, critical reception, peer endorsement, presence in public collections, position in the structure of the cultural field. Bourdieu (1993) called this the *field of cultural production*; Velthuis (2005), in his ethnography of contemporary art dealers, showed how galleries actively construct it; Beckert and Rössel (2013) generalised it to a theory of singular-good valuation in which prices are stabilised through *narratives of worth*. Radermecker (2019), studying anonymous Flemish paintings reattributed under provisional names, demonstrated how acutely price depends on this layer: the addition of a credible *Master of...* designation can shift valuation severely without altering anything visible on the canvas. The narrative-reputational layer is dominant in price determination but structurally absent from

cross-sectional listing data - the central observation that Section 7.5 returns to under machine learning, and that Section 7.2 returns to under human valuation.

The three experimental cuts of Section 7 each manipulate or bypass one of these layers in a controlled way. The anchoring task of Section 7.1 introduces a Layer-3 input - a numerical reference detached from the work - and observes how respondents incorporate it. The portrait test of Section 7.2 holds Layer 1 and Layer 2 approximately constant across the two paintings and asks whether respondents can read the Layer-3 difference that the market reads between them. The invisible-Fangor probe of Section 7.3 removes Layer 3 entirely (no curatorial framing, no biography, no signal) and observes whether physical exposure alone - pure Layer 1 - registers as encounter with art. The structure of the section, in this sense, is not three unconnected experiments. It is three coordinated probes of the same three-layer system, designed to make each layer visible by isolating it.

5. Method: a Know Thyself design

The instrument described here is not a conventional survey designed for population inference. It is a structured occasion for collecting and verifying stories that people tell about themselves - about what they know, how they value, what they remember, what they ignore - and for returning each story to its author against the background of the others. The model is an old one. In 1979, the National Museum in Kraków mounted the exhibition *Polaków portret własny - Poles' Self-Portrait* - in which the figure of the Pole was reconstructed entirely from images Poles had made of themselves. Visitors did not learn who they were from a curator. They saw themselves in what their own culture had brought, gathered in one room. What we describe below is a much smaller and more specific instance of the same gesture, reduced to one corner of one city and adapted to a teaching environment.

The instrument has been in development since 2025; the present version was finalised at the start of 2026 and has not been altered since. From the first cohort onwards, each new group enters the dataset as a session-token; the dashboard allows analysis at the pooled level and at the level of individual sessions. For the analyses reported in this article, the dataset was frozen at N = 131 completed responses. The respondents were students from various faculties of the University of Warsaw and visitors to the open seminar. Later sessions remain part of the live teaching environment and of the future replication archive, but they are not included in the present empirical results. The questionnaire is delivered through the ProfiTest platform; respondents are randomly assigned to one of three anchoring conditions. Participation is

anonymous and voluntary. We report effect sizes as η^2 for Kruskal–Wallis tests and rank-biserial correlation r for Wilcoxon and Mann–Whitney tests. Randomisation produced balanced groups: the three anchoring conditions did not differ significantly in self-reported art familiarity ($H(2) = 0.39$, $p = .822$), general interest in art ($H(2) = 0.41$, $p = .815$), age ($H(2) = 1.26$, $p = .532$), or sex distribution ($\chi^2(6) = 6.11$, $p = .411$).

The instrument has six linked elements, and each is built to make a particular assumption about the respondent visible to the respondent. The diagnostic battery presupposes a person who declares knowledge they may not have: it measures self-reported familiarity with five art movements, one of which - *Transmaterializm* - does not exist. The dual willingness-to-pay question presupposes a person who, without knowing it, divides one market into two: it asks first what they would pay for a contemporary painting they liked, then what they would pay for the same hypothetical painting treated as a financial investment. The *names* probe presupposes a person whose cultural categories are present but empty: it asks them to name young promising Polish artists, with the option to answer that there are none. The anchoring task presupposes a person who believes their valuation is their own, while in fact it is partly borrowed: it presents a Fangor work alongside a numerical price suggestion - 5,000 PLN, 500,000 PLN, or no anchor - before asking for the respondent's own valuation. The signalling task presupposes a person who trusts the authority of a description without necessarily knowing they trust it: it presents a second Fangor work after a biographical signal (including the Guggenheim exhibition and an auction record) and a curatorial-aesthetic description. The public-art probe presupposes a person who believes they have seen art they did not, in fact, register: it asks whether the respondent has encountered Fangor's work directly in the last two weeks, and then asks whether the metro graphics and the Śródmieście mosaics counted as such an encounter.

What happens after the survey closes is, pedagogically, the point of the whole design. Each respondent chooses a nickname and is then directed to a public Shiny dashboard built around their data and the data of the group. The dashboard is not only a report. It is a teaching environment built on the principle that theory should arrive after a cognitive tension, not before. The respondent enters and is addressed in the second person - "you took part in an experiment in which..." - and is led, step by step, through what they themselves did. They see, on each screen, where their own answer sat in the group's distribution, located by their nickname. The dashboard names the pattern they have just produced, and only then introduces the concept that explains it - asymmetric information, signalling, anchoring, attention as a scarce resource - in language designed for a reader who has never taken an economics course. The respondent

does not read about other people's reactions to art prices; they encounter their own, and learn the theory as the answer to a question they cannot now un-ask. The instrument is therefore both the data-collection device and the lesson - designed in the spirit of a method we have used in economics teaching for several years and described separately as the *Know Thyself* approach (Kopczewski, forthcoming).

Anonymized data and analysis code are deposited at Zenodo (DOI to be assigned at submission). New sessions are added under successive session-tokens to the same deposit, clearly separated from the data freeze used in this article. The dashboard remains publicly accessible at the project URL for protocol inspection and for use by other educators. The survey instrument, in its full Polish original and in an English working translation, is included in the supplementary materials.

6. Cognitive infrastructure: four diagnostic stories

Before respondents reached the experimental valuations of Fangor's paintings, they had already told us - and themselves - four short stories about how they relate to art. We report those stories here, in the order in which they were collected, because each sets the conditions for the experiments that follow. The figures and proportions reported in this section describe the data freeze used in this article. As new cohorts enter the dashboard, later values may move; the present sample gives the diagnostic shape of the design, not a population portrait of Poland.

Story 1. What we know, and what we say we know

Respondents were asked how familiar they were with five movements of modern art on a five-point scale. Two of those movements - *abstrakcjonizm* and *socrealizm* - are central to the visual culture in which our respondents live. Two others - *konceptualizm* and *op-art* - are the movements within which the artist whose works they were about to value, Fangor, is most often placed. The fifth, *Transmaterializm*, does not exist; we constructed it.

Self-reported familiarity averaged 2.77 out of 5 for *abstrakcjonizm*, 2.59 for *socrealizm*, 2.00 for *konceptualizm*, and 1.66 for *op-art* - the movement most directly associated with Fangor's internationally recognised work. The two movements most directly relevant to Fangor's market and reputation were the least known. The first story respondents told us about themselves, before they ever saw a painting, was therefore that the cultural categories they brought to the encounter were not symmetric. They had words for categories that already

populate ordinary public memory, and very few words for categories within which the work of major contemporary artists actually circulates.

Thirty-four percent of respondents declared at least minimal familiarity with *Transmaterializm* (95% CI [26%, 43%]). The figure is not a footnote. It measures what happens when a person is asked, in a professionally formatted question, whether they know something whose name they have not heard before. Some shrugged honestly and said no. A meaningful third filled the gap. This is the second half of the same story: declarative knowledge of art, in our sample, is not a stable record of what people have learned. It is also a response to the conditions under which the question is asked.

Story 2. One painting, two valuations

Respondents were asked, in two separate questions, what they would pay for a contemporary painting that genuinely appealed to them, and what they would pay for the same hypothetical painting treated purely as a financial investment. Same person, same hypothetical work, two answers.

For each respondent, we computed the ratio of investment to aesthetic willingness-to-pay, excluding eleven cases in which the aesthetic figure was zero ($n = 120$). The mean ratio was 6.91 ($SD = 16.85$; median = 2.18), indicating that on average respondents valued the same hypothetical painting nearly seven times more as a financial instrument than as an aesthetic object. Fifteen percent of respondents produced a ratio of ten or more (95% CI [9%, 23%]). An equal proportion - also 15% - inverted the relation, naming a higher figure for aesthetic than for investment value. The difference between the two valuations was statistically robust (Wilcoxon signed-rank test, $V = 4674$, $p < .001$, $r = 0.62$, 95% CI [0.47, 0.73]).

The story respondents told here was not simply that art is worth more as an investment than as an object of aesthetic regard, although that is the average shape of the data. The story was that the same person, asked twice about the same hypothetical painting, returned two different valuations without prompting and without instruction - and the difference, on average, was large enough to become visible as a world-view. Most respondents did not announce the divergence as an explicit decision. They produced it, and the data made it visible.

One of the present authors, on first seeing the magnitude of the gap, recognised that the world he had been taught to think about - in which beauty was understood as a value held against pricing rather than translated into it - is no longer the world his respondents inhabit.

The dashboard returns this finding to each respondent in the form of two figures placed side by side, with their nickname, and asks them what to do with the gap.

Story 3. The names that come, and the names that do not

Respondents were then asked, in two parallel questions, to name young promising Polish artists whose work appealed to them aesthetically, and young promising Polish artists whose work could be treated as a financial investment. The format was identical; the option "there are none" was offered explicitly.

Sixty-six percent of respondents answered, on each of the two questions, that there were none. Of those who did supply names, most did not supply young names. The list, condensed, includes Banksy, Sasnal, Beksiński, Fangor himself, Dwurnik, Nowosielski, Warhol, and Malewicz. These are not wrong names in the abstract. They are wrong answers to this question. They are deceased masters, global brands, or already established reputational objects. A small number of actually-young Polish contemporary artists appeared - Hubert Dolinkiewicz, Zhao Xiaoli, Gosia Herba, Agata Kus, Paweł Baśnik, Marcin Janusz, Jędrzej Bieńko - but as a small fraction of the total response.

The story respondents told here is that the cultural category *young promising Polish artist* is, for most of them, structurally empty: it has a name, but no contents. Asked to populate it, they reach for the nearest cognitively available substitute - a deceased master, a global brand, a name from auction headlines. From the side of economics this is a clean information story: in the absence of a recognised market signal for the cohort the question asks about, respondents fall back on signals that have accumulated reputational capital. From the side of art education it is a different story with the same shape: the contemporary scene that art education aims to address is, for our respondents, not present in the form of names that come to mind unprompted.

Story 4. What we have seen

Respondents were asked whether they had encountered Fangor's work directly in the last two weeks. Most answered no. They were then asked how often per week they pass the metro graphics designed by Fangor at the Uniwersytet, Świętokrzyska, and other central stations of the second metro line, or the mosaics he completed at the Warszawa Śródmieście railway station. The complete results of this probe are taken up in Section 7.3, where they form the empirical heart of the *invisible Fangor* experiment.

We mention the question here because, as a diagnostic, it operates on the same logic as the previous three. What respondents say they know, value, and remember is not always identical with what their environments have given them, and the gap is itself the data. A respondent who walks through Fangor every weekday and reports, on the survey, that they have not encountered his work in two weeks is not lying. They are reporting accurately on what they registered as encounter. The gap between physical exposure and registered encounter is what the experiment in Section 7.3 sets out to make visible.

Taken together, the four stories form a portrait: of cultural categories declared but not always populated; of valuation logics held without awareness of their division; of memory reaching for accessible substitutes when asked for unfamiliar specifics; and of a perceptual field in which physical presence does not guarantee registered presence. The three experimental cuts in Section 7 work within this portrait. They are not tests of population-level hypotheses about Polish public knowledge of contemporary art. They are structured occasions on which respondents - having told us, and themselves, who they are - discover what happens to those self-portraits under specific informational and perceptual conditions.

7. Three cuts through the wall

The four diagnostic stories of Section 6 set the conditions; the present section reports what happened when those conditions were placed under specific informational and perceptual pressure. We did not undertake three independent randomised experiments - that would have required a sample considerably larger than ours and a study designed for population inference. We undertook three different cuts through the same instrument, of three different methodological kinds, from which different things can responsibly be said. The first is a randomised comparison; the second is a paired test inside the data; the third is a structural observation about the relation between physical exposure and conscious perception. A short fourth subsection records, in the same spirit of *Know Thyself* turned on the project itself, what the design did not test, and why what we found there is itself a finding. The section closes with a brief coda from machine learning, by way of triangulation.

7.1 Anchoring: when a number that means nothing means everything

The first cut is the closest to a classical experiment. Each respondent was randomly assigned to one of three conditions before being shown a Fangor abstract from 1972: *no anchor* (no price suggestion before valuation, $n = 52$); *low anchor* (the question "do you think

the auction price was higher or lower than 5,000 PLN?" preceded the valuation, $n = 40$); *high anchor* (the same question with 500,000 PLN as the reference number, $n = 39$). The reference number was, in each case, decoupled from the actual auction outcome and from any informational basis for valuation. It was, by design, only a number.

The valuations that respondents then produced moved, predictably, but unevenly. In the *no-anchor* condition the median valuation was 4,250 PLN (IQR: 500–10,000). In the *low-anchor* condition the median was 3,250 PLN (IQR: 1,375–10,000). In the *high-anchor* condition the median was 50,000 PLN (IQR: 10,000–175,000). A Kruskal–Wallis test returned $H(2) = 25.96$, $p < .001$, $\eta^2 = .187$. Pairwise comparisons corrected with the Holm method confirmed that *no-anchor* and *low-anchor* did not differ ($p = .52$, $r = .07$), while *high-anchor* differed from *no-anchor* ($p < .001$, $r = .49$) and from *low-anchor* ($p < .001$, $r = .47$).

Two features of this result deserve to be named, because each of them works against the most common reading of "anchoring effects." The first is that the *low anchor* moved nothing. Showing respondents the figure 5,000 PLN and asking whether they thought the work sold for more or less left their distribution of valuations indistinguishable from the *no-anchor* distribution. This does not mean that anchoring failed; it means that anchoring's reach is bounded by the implausibility - or perhaps the unhelpfulness - of the suggested value. A number that respondents could already easily generate from their own intuition does not work as an anchor, because they had been generating values like it on their own. The second feature is that the *high anchor* of 500,000 PLN did not, on average, pull respondents up to that figure. The high-anchor median was 50,000 PLN. The anchor exerts gravity, but the gravity is partial: respondents adjust away from the anchor in the direction of their own prior, while remaining structurally attached to it. This is the *anchoring-and-adjustment* pattern that Tversky and Kahneman (1974) described in their original studies, and that Northcraft and Neale (1987) reproduced among professional real-estate appraisers. Beggs and Graddy (2009) found the same shape in actual art-auction data, where pre-sale estimates partially anchor realised hammer prices. Our students reproduced the shape inside half a minute in front of a computer screen.

What is being measured here is not aesthetic judgement and not market knowledge. It is something more elementary: the willingness of an uncertain valuation to attach itself to whatever quantitative reference is offered first. In the dashboard, the respondent learns to recognise this in their own response, located by their nickname against the distribution of the others.

7.2 *The portrait test: name as brand, style as filter*

The second cut is not a randomised experiment but a paired comparison built into the instrument. Every respondent, regardless of anchoring condition, was shown two oil-on-canvas portraits painted by Fangor in 1948 and 1949 - early figurative works from the period before the abstraction for which he is now known, before his American career, and before the op-art series that anchors his contemporary market. The two paintings are of comparable size (91.5×73 cm and 65×50 cm), comparable medium, and comparable apparent technique. Visually they belong to the same world: a young figurative painter working under the constraints of late-1940s Polish art education. Respondents were asked to value each.

The actual auction outcomes for the two works diverge by a factor of nearly ten. The 1948 portrait reached approximately 600,000 PLN; the 1949 portrait reached approximately 64,000 PLN. The market reads them as different objects: one is a moment of art-historical transition and is priced as such; the other is the work of a young painter not yet distinguishable from his cohort and is priced accordingly. Whatever generates this tenfold difference, it is not anything visible on the surface of the canvas to a non-expert viewer.

The respondents' valuations did distinguish between the two works - but not in the same magnitude. The 1948 portrait received a median valuation of 70,000 PLN; the 1949 portrait received a median valuation of 50,000 PLN. The directional comparison was statistically robust (Wilcoxon paired test, $V = 4879.5$, $p < .001$, $r = 0.41$, 95% CI [0.24, 0.56]): of the 131 respondents, 82 valued the 1948 work higher, 35 reversed the order, and 14 valued the two equally. Respondents, in other words, *felt* a difference between the two paintings. They simply did not feel a tenfold difference. The market saw a factor of ten; respondents, on the median, saw a factor of about 1.4.

This is the cleanest test in our data of the difference between what Okhrimenko (forthcoming) calls *Layer 2* - the visual-aesthetic surface of the work - and *Layer 3* - the narrative-reputational capital that surrounds it. To a non-expert viewer the two portraits live almost entirely on Layer 2, and on Layer 2 they look similar. The market lives, for these works, almost entirely on Layer 3, where one is "the early Fangor before the turn" and the other is "another youthful painting." The respondents could not see Layer 3 because Layer 3 is not present in the canvas; it is present only in a body of art-historical and market knowledge most of them had not been given. The test demonstrates the layer-mismatch directly: the gap between respondent valuation and market valuation is not noise. It is the visible footprint of a layer of meaning that operates structurally outside the image.

The same paired comparison also reveals an anchoring carry-over: respondents in the high-anchor condition valued both portraits considerably higher than respondents in the no-anchor condition (1948 median: 400,000 vs 19,000 PLN; 1949 median: 320,000 vs 8,000 PLN). The earlier numerical anchor did not stop with the work to which it was attached. It continued to operate across screens and across paintings, several minutes later, on works it was never about. We return to this in Section 7.4.

7.3 *The invisible Fangor: physical exposure without perception*

The third cut is the simplest measurement we made and produced the result we were least prepared for. After all the valuations were complete, respondents were asked: "In the last two weeks, have you encountered Fangor's work directly?" Ten of the 131 said yes; 121 said no. They were then asked, separately, how often per week they pass the metro graphics designed by Fangor or the mosaics he completed at the Warszawa Śródmieście railway station. Sixty respondents declared that they pass these works at least once a week.

The two answers can be cross-tabulated. Of the 60 respondents who reported passing Fangor's public works at least once a week, every single one - without exception - answered "no" to the question of whether they had encountered his work in the previous two weeks. Sixty out of sixty. The proportion does not require statistical machinery. It is what we have, and it is unambiguous.

The ten respondents who answered "yes" were given a follow-up question. Six confirmed they had been thinking of the metro graphics and the Śródmieście mosaics; four had been thinking of something else - leaving four respondents (3.1% of the sample) who had encountered Fangor's work in a context other than the public works they pass daily.

The result is not a statement about Fangor and not a statement about his audience. It is a statement about the relation between physical exposure and conscious perception under conditions in which the work is not framed as art. The metro graphics are, by Warsaw City Council's own description, "a gallery of world op-art." The students of the University of Warsaw are, in our data, a population to whom that gallery is invisible *as a gallery* even as they pass through it daily. This is what we mean, in this article, by the *invisible Fangor*. It is not an idiom. It is the proportion 60/60.

The pedagogical implication is the heart of what we hope this article transmits to colleagues in art and design education. Visibility, in the sense that matters to art education,

is not a function of presence. It is a function of the framing devices that turn presence into encounter - the institutional thresholds, the curatorial language, the wall labels, the catalogues, the queuing, the silence, the reverence - none of which are operative when the same object is a wall one walks past on the way to a class. Crawford (2015) and Davenport and Beck (2001) have written about the broader attention-economic conditions under which this happens. Our small dataset adds one local, specific, and methodologically clean instance: the city has named a man, the city has designed its trains around his work, and the city's most educated young population, when asked, did not see him. The dashboard returns this to each respondent in the second person, by their nickname. Several, on first encountering the result, have written to ask whether the question was about the same Fangor whose name they know.

7.4 *The anchor that ate the signal*

The fourth subsection is a methodological note that is also a finding. Our instrument, in addition to the three cuts above, presented every respondent with a second valuation task: a later Fangor work (M78, from 1968), introduced after a biographical signal - including the Guggenheim exhibition of 1970 and an auction record - and a curatorial-aesthetic description drawn from the actual Polswiss Art catalogue. The intent, when the design was first written, was to test whether the curatorial-aesthetic signal would shift respondents' valuations differently from the market-institutional signal. In the version of the instrument used for this data freeze, respondents were not randomised across signal conditions. Every respondent saw the biography, the institutional history, and the curatorial description, in that order. What did differ, across respondents, was the anchoring condition from the earlier task with the 1972 work.

The valuations of M78 broke down by anchoring group - the only assignment variable that distinguished respondents - as follows. Respondents in the no-anchor condition gave a median valuation of 50,000 PLN (IQR: 11,500–212,500). Respondents in the low-anchor condition gave a median of 100,000 PLN (IQR: 50,000–312,500). Respondents in the high-anchor condition gave a median of 600,000 PLN (IQR: 300,000–1,000,000). A Kruskal–Wallis test returned $H(2) = 21.21$, $p < .001$, $\eta^2 = .150$. Pairwise comparisons (Holm correction) showed that *no-anchor* and *low-anchor* did not differ significantly ($p = .054$), while *high-anchor* differed from *no-anchor* ($p < .001$) and from *low-anchor* ($p = .003$).

The anchor that respondents had encountered minutes earlier - attached to a different painting, in a different period of Fangor's career, with no informational relationship to the work currently in front of them - continued to organise their valuations after the biography,

the Guggenheim, and the curatorial language had been read. Whatever effect the additional information may have had, it did not displace the anchor. To phrase it directly: in this design, the anchor ate the signal.

This is, on the face of it, a methodological limitation. A more ambitious version of this study would cross anchoring conditions with independently randomised signal conditions and test each separately. We did not do this. Within the constraints of a small pedagogical sample - three anchoring conditions already produce three groups of approximately thirty-five - adding a second three-level randomisation would produce nine cells of around twelve respondents, below any threshold for stable inference. The decision to keep the design simple was made in service of the pedagogy: we wanted every respondent to leave the instrument with a clear reading of one manipulation (the anchor), and to confront its consequences in the dashboard. A future iteration of the protocol, run on a sufficiently large sample by colleagues willing to forgo the within-subject clarity, could randomise signal independently.

We mention this not as an apology but as a finding. A pedagogical instrument designed to make the operation of the anchor visible to its respondents has shown, on its own data, that the anchor is more persistent than the version of theory that imagined the signal as a competing or additive force on top of it. Spence's (1973) signalling, Akerlof's (1970) lemons, and Velthuis's (2005) symbolic price all assume an informational substrate in which signals are read against the absence of an anchor. Our respondents, having received the anchor first, never returned to that absence. What they did to subsequent information - biography, institution, curatorial voice - is what people do to subsequent information when their first reading has already taken place. They incorporated it into the structure they had already begun to build.

7.5 What painting still resists: a coda from machine learning

A fifth observation, brief but deserving its own place in the section, comes from the contrast between domains. Over the last two decades, machine-learning systems have become remarkably effective at predicting cultural success in three of the four major creative industries. In music, *Hit Song Science* - the field formalised by McFee et al. (2012) and developed in dozens of subsequent studies - now produces classifiers that distinguish Billboard hits from non-hits with accuracies in the seventy-five to ninety percent range using audio features alone (Georgieva, Suta, & Burton, 2018; Vavaroutsos, 2024). In film and series, Netflix's recommendation engine drives more than eighty percent of viewing on the platform

and shapes the production decisions of one of the largest content studios in the world. In literature and text, book-recommendation systems and Spotify-style personalisation have become standard infrastructure for online reading platforms. Whatever exactly these algorithms learn, they have learned enough about songs, films, and books to direct hundreds of millions of people through them every day.

Painting, alone among the major art forms, has not yielded to this treatment. The pigment, the canvas, the dimensions, the medium, and the visible composition - what one might call the physical description of a painting, in the sense relevant to a machine-learning model - are not enough to predict the price the work will reach. Aubry et al. (2019), in a comparison of hedonic regression and gradient-boosted machines on hundreds of thousands of auction records, found that the most predictive variables were almost all non-visual: artist identity, auction house, prior sales, signing, dating, and exhibition history. Greene (2020), reviewing the field for the *Harvard Data Science Review*, summarised the methodological situation directly: convolutional neural networks, applied to the images of paintings alone, do a poor job of predicting auction prices; models based on textual and numerical metadata perform substantially better. Lee et al. (2024), with the most explicit comparison to date - running an XGBoost on visual features only versus an XGBoost on social-and-market features only - found that visual features played "a marginal role" in price prediction, while social signals approached expert-level accuracy. Their result is sharper in emerging markets, where the buy-the-artist logic dominates over the buy-the-artwork logic.

A parallel study by the second author of this article confirms the same pattern in a different sample. Okhrimenko (forthcoming), working with approximately 21,000 contemporary artworks scraped from the Artsper marketplace, applied a sequence of model architectures - tree-based ensembles, multilayer perceptrons, and convolutional networks combining image and tabular features - to the task of predicting price from observable attributes. Across all model families and feature combinations, predictive performance was poor by conventional standards. SHAP analyses revealed that what little signal the models could find was carried almost entirely by physical attributes: dimensions, medium, age, and time on platform. The narrative-reputational layer - the one that, as Section 7.2 just demonstrated, the human respondents in our Warsaw sample also could not see - is structurally absent from the data such cross-sectional listings provide. Models cannot learn what is not in the data.

The contrast with music, film, and literature is what we want to leave the reader holding. In those domains, observable structural features of the artefact - the audio signal, the visual

frames, the text - turn out to carry a great deal of the information that determines popular reception. A song's tempo, key, and timbral envelope are statistically related to whether listeners will play it again. A film's pacing, dialogue density, and visual style are statistically related to whether viewers will finish it. A novel's vocabulary and sentence rhythms are statistically related to whether readers will recommend it. Painting does not yet behave this way. The arrangement of pigment on canvas - the physical fact of the work, in the strictest sense - has not, so far, been shown to determine the price the work commands. Whatever does determine it sits outside the canvas. It sits in a layer of cumulative cultural transaction: who painted, when, for whom, in what institutional context, with what subsequent endorsement, in what auction sequence. The triangulation between our human respondents and the algorithms is therefore straightforward. Humans without a narrative cannot value paintings stably; algorithms without a narrative cannot price them reliably. The two failures are not failures of the same kind, but they describe the same absence - an absence that this article has tried, in its small Warsaw way, to make visible to its respondents, and through the protocol of Section 8, to make available for replication elsewhere.

8. The protocol as transferable design

The instrument described in this article is not, properly speaking, ours alone. From its first running it was designed to be repeated - by other educators, in other cities, with other artists, under conditions different enough from ours to make the comparison interesting. The empirical findings reported in Sections 6 and 7 are useful in their own terms, but their largest value lies in the protocol that produced them: a structured but adaptable sequence of tasks through which a teacher in either economics or art education can, in a single seminar or workshop, give students access to their own valuation behaviour and to the literatures that explain it. This section sets out what is invariant in that protocol, what is local, and how an educator elsewhere might begin.

What is invariant is, first, the gesture. The respondent must arrive without instruction and produce - in answer to specific, structured prompts - a set of valuations, recognitions, and recollections that they themselves will later be asked to confront. The protocol must address them in the second person. The mirror in the data must be located by their own nickname. Theory must arrive after, not before, the cognitive tension that the instrument creates. These conditions are non-negotiable not because we hold them rigidly, but because the protocol stops working without them. A version that begins by teaching the theory of anchoring before asking

respondents to value a painting will produce different data, and - more importantly - a different student.

What is invariant, second, is the sequence of six probes described in Section 5: the diagnostic battery with its fictitious-movement control; the dual willingness-to-pay question separating aesthetic from investment value; the *names* probe; the anchoring task; the signalling task; the public-art exposure-versus-perception question. These six together generate the four diagnostic stories of Section 6 and the experimental cuts of Section 7. Removing any of them collapses part of the portrait. Adding to them, by contrast, is straightforward. Local educators may want to include questions specific to their cultural context, and the dashboard can accommodate them without disturbing the core.

One reason we expect different results in different cities is that the cultural infrastructure surrounding contemporary art varies between countries in ways that this protocol can make visible. In Poland, the school subject *plastyka* - comparable to *visual arts* in the Anglophone curriculum, or *Bildende Kunst* in the German - has been substantially reduced in hours and standing across the last twenty years. One possible consequence appears in our data as the *names problem* described in Section 6. We do not know - and this is precisely the empirical question the protocol asks - whether the same thinness of contemporary art knowledge holds among students in cities whose curricula have not been similarly reduced, or among publics whose urban art is more deliberately staged.

We suspect, however, that even where the school subject is healthier, the issue will not vanish. Those of us who teach and write about contemporary art tend to overestimate how much of the contemporary scene reaches the people around us, because the people around us are selected for the same orientation. This is the insider version of the representativeness bias described by Tversky and Kahneman (1974): we mistake the small, culturally selected world around us for the larger public to which art education claims to speak. It is one reason a 1979 exhibition like *Polaków portret własny - Poles' Self-Portrait* - was needed in the first place, and why a smaller contemporary version of the same gesture - gathering portraits people make of themselves through what they value, ignore, and remember - may be useful in any city. Every city has at least one contemporary artist whose work is materially present in its streets, stations, churches, squares, museums, or public commissions - present often enough to be passed, recognisable enough that the city has chosen to name, preserve, commission, or display it, yet not necessarily visible as art to those who live among it. Whether students see these works,

whether they see them as deserving attention, or whether they see them only once a name, price, or institution has been attached to them, is the question this protocol can ask.

In Kraków, the question could be asked through Tadeusz Kantor's *Cricot 2* sites or Jerzy Nowosielski's sacred interiors. In Berlin, it could be asked through artists whose works carry the weight of post-war memory, division, and reunification. In Vienna, Hundertwasser would offer an almost opposite case: an artist whose public visibility is already staged, touristic, and strongly branded. These are not prescriptions. They are examples of the kind of local substitution the protocol requires. The Warsaw case used Fangor because he combines public presence, market recognition, art-historical weight, and perceptual invisibility. A replication needs its own equivalent.

What the protocol then asks, in its public-art probe, is the same question in each city: do students notice the work that is there? In its anchoring and signalling probes, it asks a related question: when students do notice, what do they see - an art-historical figure, a market name, an investment, a public ornament, a fragment of urban background? The answers may vary by city. The variation is the data.

The protocol functions, in this sense, as what Star and Griesemer (1989) called a boundary object: an instrument sufficiently structured to be replicated and compared across sites, yet sufficiently open to be reinterpreted by educators with different disciplinary commitments. An economics educator will read it as a sequence of natural experiments in asymmetric information, anchoring, and signalling. An art and design educator will read it as an investigation of perception, attention, and the conditions under which artworks become visible to their publics. The same dataset, the same dashboard, the same six probes - used by both. We do not require either side to abandon its language for the other's. We propose only that the same instrument can host both.

Practical infrastructure for replication consists of three elements. The Shiny dashboard source code, including the templated explanatory texts that introduce concepts in plain language, is available on GitHub under a permissive licence. The questionnaire instrument, in its full Polish original and an English working translation, is included in the supplementary materials. Anonymized data, cumulative across all sessions to date, are deposited at Zenodo with a citable DOI; new sessions are added to the same deposit under successive session-tokens, so that comparative analysis across cohorts and cities remains possible. We invite educators considering a replication to contact the corresponding author. We do not require co-authorship of resulting publications; we ask only that replications, whether successful, partial,

or contradictory, be deposited in the shared archive. The protocol has no completion date by design. We hope that this article marks the moment at which the instrument leaves our hands and begins to carry the shape of cities other than Warsaw.

9. Coda: surprises, limitations, and an open ending

9.1 Four shocks

The figures reported in this article were, for the four of us, more startling than we had expected them to be. The shocks were not the same.

Tomasz K. I am the only one who can remember the late socialist economy. I was born in the late 1960s and entered adult cultural life in conditions of material scarcity but cultural saturation. We did not understand, at the time, why young people in the West complained about capitalism. They had everything we wanted. What was the problem?

The answer came late, through this project. What changed was not only income, institutions, or access. What changed was time. That is why the result that struck me most was not simply the low declared knowledge of contemporary art, but the nearly 7× ratio between investment and aesthetic willingness-to-pay. Respondents valued the same hypothetical object far more as an asset than as something they might want to live with, look at, or enjoy.

I came to call this the Beuys paradox. Joseph Beuys placed fat, felt, decay, and repulsive materiality into the frame of art, declaring a social sculpture beyond the ordinary market object. Decades later, the work that was meant to resist commodification has become expensive to preserve, insure, narrate, and circulate. More importantly, the postulate of art without object has been realised not outside the market, but inside it. The art market does not need pigment, canvas, or labour in the old sense. It needs name recognition, auction signals, institutional networks, and engineered attention. Beuys postulated idea without aesthetic pleasure. The market answered with price without the need to look. Our 7× finding is a small Warsaw instance of that larger outcome.

Justyna. In my previous study, my hypothesis was that people with greater knowledge of art would be less susceptible to anchoring and signalling - that their own knowledge would protect them from external influence. The opposite turned out to be true: the greater the familiarity with art, the stronger the response to the signal. My interpretation is that knowing about art, in our sample, means above all knowing whose opinion matters. An expert is more

attuned to the signals of their environment than anyone else - and for precisely that reason, more susceptible to them.

The second observation is personally disappointing. I work in the cultural sector and for years I assumed that art is present in public life - that people pass it, notice it, that something reaches them even without an institutional frame. Our data suggest otherwise. Sixty respondents declared that they pass Fangor's works at least once a week. Not one of them - literally not a single person - counted that as an encounter with art. That is a result that is hard to dismiss. Art can be physically everywhere and at the same time mean nothing - if nothing around it says that it is worth a moment of your attention.

Jan. From the side of economics, the mechanisms here are unsurprising. Anchoring works. Signals shift valuations. Without recognised market signals, respondents fall back on accumulated reputational capital. These are standard results. What remains surprising is that economics teaching so rarely uses the art market to illustrate them, given that it concentrates asymmetric information, signalling, anchoring, and status consumption in unusually legible form.

What the data made visible more sharply than theory alone is the feedback loop between ignorance and price. Our respondents knew very little about the works in front of them - and precisely because of that, price became not one signal among many, but the only signal. This is rational behaviour under radical uncertainty, but it has a structural consequence. If price is the primary vehicle through which most people encounter and interpret art, then the reputational layer - critics, curators, auction houses - acquires an influence on public perception that is almost entirely self-reinforcing: reputation shapes price, price shapes what the public believes a work is worth, and that belief shapes the conditions under which reputation is next constructed. Our respondents did not know they were inside this loop. Making that loop visible - to them, through their own data - is precisely what the instrument was built to do.

Kostiantyn. My surprise came from the data-science side. We can measure size, medium, age, visual features, and listing metadata. We can train models on thousands of artworks. But the signal that seems to matter most is often the one that is absent from the dataset: the narrative-reputational layer. Machines fail where humans also hesitate. Without biography, provenance, institutional endorsement, and auction history, both the algorithm and the respondent are asked to price a surface whose market meaning is elsewhere.

Tomasz P. In the research, I stood between generations. I was not formed entirely by the older world in which beauty was defended against price, but neither do I belong to the younger cohort for whom art, investment, branding, and market signals seem to coexist almost naturally. I had watched the development of the Polish art market with fascination, but also with a certain scientific distance: from the outside, the market often appeared opaque, difficult to enter analytically, and governed by signals visible to insiders but hard to translate into a researchable language. What the experiment changed was not my view of art, but my view of that opacity. The protocol turned a closed world into an observable sequence of small decisions: recognition, first impressions, price cues, reputation, hesitation, adjustment. My first reaction was therefore not theoretical but methodological: the experiment made the art market researchable without first pretending that it is transparent. I also saw, immediately, possible connections with fast-and-frugal heuristics and the Gigerenzer tradition. I leave that connection for a future paper. Here, the important point is simpler: the design did not remove the opacity of the art market; it made the opacity visible enough to teach with.

A generational conversation. These five reactions are not symmetrical, and they are not meant to be. Each of us saw a different first problem: memory, field knowledge, microeconomic theory, prediction, or attention. Together, however, they point to the same mechanism: when direct looking is weak, signals and narratives do the work. The Fangor case is therefore not only about art-market valuation. It is also about a changed ecology of attention. When time becomes scarce, attention becomes expensive. When attention becomes expensive, price, name, auction history, and institutional status become shortcuts through a world in which nobody has enough time to learn everything slowly.

From Tomasz K.'s perspective, cultural access once looked like a problem of material scarcity. From the experience of the younger authors, the sharper constraint is time. The contrast is not simply between poverty and affluence, or between closed socialism and open capitalism. It is a change in the structure of attention. Students today live in a wealthier country than the one Tomasz K. knew as an adult and Tomasz P. knew while growing up, but they often have less unclaimed time. Their attention is pulled in one direction by networked media and algorithmic feeds, and in another by the need to work while studying.

In Poland, this is not a marginal condition. Recent EUROSTUDENT data show that student employment during the lecture period is a normal part of student life, not an exception. Recent Polish data on cultural participation point in the same direction from the other side: lack of time is one of the main declared reasons for not using cultural institutions. Time poverty is

therefore not decorative language added to the data after the fact. It is part of the social background from which these valuations were produced.

This changes the meaning of cultural access. The barrier is not primarily the price of a museum ticket. The deeper barrier is that art now has to compete for a scheduled slot in a life already fragmented by work, study, commuting, screens, and fatigue. In the late socialist city, at least in memory, art could be absorbed almost incidentally: one passed through museums, galleries, posters, exhibitions, and public interiors. Today art is still present in the city - and also in the smartphone - but presence no longer guarantees encounter.

The Warsaw Fangor is the perfect example. Students pass through his metro graphics and near his mosaics, but the works do not become cultural experience unless attention has enough slack to attach to them. This is why the invisible Fangor result is more than a local curiosity. It records a change in the ecology of attention. The result is not cultural emptiness. It is a different form of cultural participation: more continuous, more mediated, more private, more interwoven with work and routine, and less often organised by the slow public time of institutions.

The nearly 7× ratio of investment to aesthetic willingness-to-pay made this shift visible. When time becomes scarce, attention becomes expensive. When attention becomes expensive, signals replace looking. Price, name, auction history, and institutional status become shortcuts through a world in which nobody has enough time to learn everything slowly.

9.2 Limitations

Several caveats follow from what has just been said. The two blind walls were located in one corner of a city, with a sample dominated by one university and one open seminar. The 131 respondents in the data freeze do not represent Poland; they represent the kind of person who, in early 2026, walked into a Warsaw seminar room or an open seminar booth and chose to spend a quarter of an hour answering our questions. Order effects are partly built into the within-subject design: anchoring exposure precedes the biographical and curatorial signal for each respondent, and we cannot fully untangle the two without a between-subjects replication of the kind described in Section 7.4. Self-classified expertise is not professional expertise, and the surprise reported above by Justyna should be read with that caveat - among trained valuers and curators, the patterns may operate differently or not at all. Single-artist generalisability is a limitation of design rather than oversight: Fangor's particular configuration of public presence, market recognition, art-historical weight, and perceptual invisibility produces results that may not transfer to artists with different configurations. The 34%

prevalence of declared familiarity with the fictitious *Transmaterializm* is striking but rests on a single measurement; before being treated as a stable feature of cognitive infrastructure under our diagnostic conditions, it requires independent reproduction. Replications elsewhere are precisely the response to all of these limitations, which is why Section 8 took the form it did.

9.3 Closing

The three Warsaw sites with which this article opened - the metro station where students enter and leave the city, the railway-station mosaics that are slowly fading, the new street alongside the Museum of Modern Art - remain in place after this article is written. They will continue to be visible, ignored, and visible again, in whatever rhythm the city allows. We do not propose a remedy for that rhythm. We have proposed a small instrument with which a teacher in either economics or art education can ask students to look once more, and to look slightly differently. The instrument has no completion date; the dashboard will continue to grow as new cohorts enter under new session-tokens, in Warsaw and, we hope, elsewhere.

The two walls of our title are not walls. They are the absence of doors. Where one is cut, both sides can pass through, and what each finds on the other side is part of what their students were already trying to see.

Data and code availability: The survey instrument, anonymised dataset for the present data freeze, the cumulative dataset across session tokens, the R analysis code, and the Shiny dashboard source will be deposited on Zenodo with a DOI assigned at submission. The dashboard remains live at the project URLs: https://microeconomics.shinyapps.io/art_kotwica/ [https://microeconomics.shinyapps.io/art_anchor/] for protocol inspection. Replication enquiries: tkopczewski@wne.uw.edu.pl.

AI statement: This paper was written in conversation with an AI assistant that contributed no ideas and originated no arguments. Its role was what the poet demanded of language itself: that the supple tongue say all the head can think - no more, no less. The authors supplied the head. The machine supplied the tongue. Responsibility for both remains entirely with the authors. [see appendix: <https://doi.org/10.33138/2957-0506.2026.14.508>]

"I chodzi mi o to, aby język giętki powiedział wszystko, co pomyśli głowa."

"I want the supple tongue to say all that the head can think."

- Juliusz Słowacki, Beniowski, Pieśń V

References

- Adorno, T.W., 1991. *The culture industry: Selected essays on mass culture* (Bernstein, J.M., Ed.). London: Routledge.
- Akerlof, G.A., 1970. The market for lemons: Quality uncertainty and the market mechanism. *Quarterly Journal of Economics*, 84(3), 488–500.
- ARC Rynek i Opinia, 2021. *Uczestnictwo Polaków w kulturze: Raport z badania*. Warszawa: ARC Rynek i Opinia.
- Aubry, M., Kräussl, R., Manso, G., Spaenjers, C., 2019. *Machines and masterpieces: Predicting prices in the art auction market*. Working paper, Northeastern University.
- Becker, W.E., Watts, M., 2001. Teaching economics at the start of the 21st century: Still chalk-and-talk. *American Economic Review*, 91(2), 446–451.
- Beckert, J., Rössel, J., 2013. The price of art: Uncertainty and reputation in the art field. *European Societies*, 15(2), 178–195.
- Beggs, A., Graddy, K., 2009. Anchoring effects: Evidence from art auctions. *American Economic Review*, 99(3), 1027–1039.
- Bourdieu, P., 1993. *The field of cultural production*. Cambridge: Polity.
- Bourdieu, P., 1996. *The rules of art: Genesis and structure of the literary field*. Stanford: Stanford University Press.
- CBOS, 2025. *Pokolenie Z a media społecznościowe: Komunikat z badań*. Warszawa: Centrum Badania Opinii Społecznej.
- Crawford, M.B., 2015. *The world beyond your head: On becoming an individual in an age of distraction*. New York: Farrar, Straus and Giroux.
- Davenport, T., Beck, J., 2001. *The attention economy: Understanding the new currency of business*. Boston: Harvard Business School Press.
- Eisner, E.W., 2002. *The arts and the creation of mind*. New Haven: Yale University Press.
- Eurostudent IV, 2011. *Polish students working hours, reported in Rzeczpospolita, 2011*.
- Georgieva, E., Suta, M., Burton, N., 2018. *HitPredict: Predicting hit songs using Spotify data*. Stanford CS229 project report.
- Greene, J.P., 2020. Can machine learning predict the price of art at auction? *Harvard Data Science Review*, 2(2).
- Jagodzinski, J., 2010. *Visual art and education in an era of designer capitalism: Deconstructing the oral eye*. New York: Palgrave Macmillan.

- Jagodzinski, J., 2015. The challenges of art education in designer capitalism: Collaborative practices in the (new media) arts. *International Journal of Art & Design Education*, 34(3), 282–295.
- Jagodzinski, J., 2020. *Pedagogical explorations in a posthuman age*. Cham: Palgrave Macmillan.
- Kahneman, D., 2011. *Thinking, fast and slow*. New York: Farrar, Straus and Giroux.
- Kopczewski, T. (forthcoming). *The Know Thyself method in microeconomics teaching*. Working paper, University of Warsaw.
- Krajewski, M., Szlendak, T., 2014. *Domesticated culture: Polish everyday cultural participation*. Warsaw: Narodowe Centrum Kultury [working translation; original published in Polish].
- Laskowska, J., 2025. *Anchoring effect and the popularity of artworks: A case study of Wojciech Fangor in the international art market*. Master's thesis, University of Warsaw, Faculty of Political Science and International Studies.
- Lee, M., et al., 2024. Social signals predict contemporary art prices better than visual features, particularly in emerging markets. *Scientific Reports*, 14, 11578.
- Letycja, E., 2026. *Lexical analysis of introductory economics and sociology textbooks*. <https://doi.org/10.33138/2957-0506.2026.10.504>
- McFee, B., Bertin-Mahieux, T., Ellis, D.P.W., Lanckriet, G.R.G., 2012. *The Million Song Dataset challenge*. Proceedings of the 21st International Conference on World Wide Web.
- Mirowski, P., 2002. *Machine dreams: Economics becomes a cyborg science*. Cambridge: Cambridge University Press.
- Mirowski, P., 2013. *Never let a serious crisis go to waste: How neoliberalism survived the financial meltdown*. London: Verso.
- MKiDN, 2015. *Uczestnictwo Polaków w kulturze: Raport z badania*. Warszawa: Ministerstwo Kultury i Dziedzictwa Narodowego.
- Northcraft, G.B., Neale, M.A., 1987. Experts, amateurs, and real estate: An anchoring-and-adjustment perspective on property pricing decisions. *Organizational Behavior and Human Decision Processes*, 39(1), 84–97.
- Okhrimenko, K. (forthcoming). *Painting a price: A machine learning approach to art valuation*. Master's thesis, University of Warsaw, Faculty of Economic Sciences.
- Radermecker, A.-S., 2019. *The Master of...: Creating names for art history and the art market*. *Journal of Cultural Economics*, 43(1), 57–95.
- Read, H., 1943. *Education through art*. London: Faber and Faber.

- Renneboog, L., Spaenjers, C., 2013. Buying beauty: On prices and returns in the art market. *Management Science*, 59(1), 36–53.
- Rosen, S., 1974. Hedonic prices and implicit markets: Product differentiation in pure competition. *Journal of Political Economy*, 82(1), 34–55.
- Spence, M., 1973. Job market signaling. *Quarterly Journal of Economics*, 87(3), 355–374.
- Star, S.L., Griesemer, J.R., 1989. Institutional ecology, "translations" and boundary objects: Amateurs and professionals in Berkeley's Museum of Vertebrate Zoology, 1907–39. *Social Studies of Science*, 19(3), 387–420.
- Throsby, D., 1994. The production and consumption of the arts: A view of cultural economics. *Journal of Economic Literature*, 32(1), 1–29.
- Tversky, A., Kahneman, D., 1974. Judgment under uncertainty: Heuristics and biases. *Science*, 185(4157), 1124–1131.
- Vavaroutsos, P., 2024. HSP-TL: A deep metric learning model with triplet loss for hit song prediction using lyrics and audio features. *Intelligent Systems with Applications*.
- Velthuis, O. 2005. *Talking prices: Symbolic meanings of prices on the market for contemporary art*. Princeton: Princeton University Press.
- Walstad, W.B., Siegfried, J.J., 2015. Research on teaching economics to undergraduates. *Journal of Economic Literature*, 53(2), 285–325.



UNIVERSITY OF WARSAW
FACULTY OF ECONOMIC SCIENCES
44/50 DŁUGA ST.
00-241 WARSAW
WWW.WNE.UW.EDU.PL
ISSN 2957-0506