NBER WORKING PAPER SERIES

STATE OF THE ART:
ECONOMIC DEVELOPMENT THROUGH THE LENS OF PAINTINGS

Clément Gorin
Stephan Heblich
Yanos Zylberberg

Working Paper 33976
http://www.nber.org/papers/w33976

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
June 2025

We would like to thank Dionissi Aliprantis, Laure Athias, Sascha Becker, Karol Borowiecki, Nicolas Baumard, Davide Cantoni, Don Davis, Shari Eli, Oliver Falck, Oded Galor, Ed Glaeser, Claudia Goldin, Erik Hornung, Robert Johannes, Hubertus Kohle, Ruru Hoong, Michael Hutter, Bob Margo, Stelios Michalopoulos, Petra Moser, Nathan Nunn, Ömer Özak, Dominic Rohner, Stefanie Schneider, Munir Squires, Marco Tabellini, Mathias Thoenig, Daniel Trefler, Matt Turner, Hans-Joachim Voth, Fabian Waldinger, and participants at Brown, the Cleveland Fed, Cologne, Harvard/BU, LMU, Paris Saclay, Paris Dauphine, Ecole Normale Supérieure de Lyon, the University of Bristol, the University of Toronto, UBC, Warwick, and the ACEI 2023, the 2024 CUER-UBC workshop, ENS-PSL, the 2024 German Economists Abroad conference, the 2024 IEB Workshop on Urban Economics, the Microdata in Economic History Workshop, and the NBER SI for helpful comments. Leyi Wang provided excellent research assistance. We acknowledge the research support from the ORA Grant ES/V013602/1 (MAPHIS) and the Social Sciences and Humanities Research Council of Canada (SSHRC, Insight Grant 435-2025-0725). The usual disclaimer applies. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

State of the Art: Economic Development Through the Lens of Paintings
Clément Gorin, Stephan Heblich, and Yanos Zylberberg
NBER Working Paper No. 33976
June 2025
JEL No. C45, O10, R11, Z1

ABSTRACT

This paper analyzes 630,000 paintings from 1400 onward to uncover how visual art reflects its socioeconomic context. We develop a learning algorithm to predict nine basic emotions conveyed in each painting and isolate a context effect—the emotional signal shared across artworks created in the same location and year—controlling for artist, genre, and epoch-specific influences. These emotion distributions encode subtle but meaningful information about the living standards, uncertainty, or inequality characterizing the context in which the artworks were produced. We propose this emotion-based measure, derived from historical artworks, as a novel lens to examine how societies experienced major socioeconomic transformations, including climate variability, trade dynamics, technological change, shifts in knowledge production, and political transitions.

Clément Gorin
Paris 1 Pantheon-Sorbonne University
Sorbonne School of Economics
Sorbonne Economics Centre (CES)
clement.gorin@univ-paris1.fr

Stephan Heblich
University of Toronto
Department of Economics
and NBER
stephan.heblich@utoronto.ca

Yanos Zylberberg
University of Bristol
yanos.zylberberg@bristol.ac.uk

A data appendix is available at http://www.nber.org/data-appendix/w33976

Throughout history, painters have created images that reflect the social, political, and economic conditions of their world, offering visual records of ritual and daily life, power and protest, and poverty and prosperity, especially in contexts where other forms of data are sparse or absent. Yet despite their richness, paintings have rarely been used at scale to systematically study historical societies. This paper introduces novel, algorithmic methods to analyze the socioeconomic conditions of the past through the lens of 630,846 paintings by more than 29,000 artists.

Figure 1. Paintings as testaments to their times and contexts.

(a) Liberty Leading the People

(b) Guernica

(c) Dance at the Moulin de la Galette

(d) The Potato Eaters

Notes: Panel (a) shows *Liberty Leading the People* (Eugène Delacroix, 1830), which captures the short-lived July Revolution of 1830 in France. Panel (b) displays *Guernica* (Pablo Picasso, 1937). Panel (c) shows *Dance at the Moulin de la Galette* (Pierre-Auguste Renoir, 1876), which depicts a lively scene at a popular Parisian dance hall. Panel (d) shows *The Potato Eaters* (Vincent van Gogh, 1885), a portrayal of rural life in Nuenen, which van Gogh described in a letter to his brother Theo: "By continually observing peasant life, at all hours of the day, I have become so involved in it that I rarely think of anything else" (April 1885). Additional illustrations discussed in this introduction are provided in Appendix A.

The following examples illustrate how paintings recover societal contexts from the visual traces left by artists. Political turbulence and protest are portrayed in Eugène Delacroix's *Liberty Leading the People* (1830), which celebrates the fight for freedom by citizens from markedly different social classes (panel a of Figure 1), and in Pablo Picasso's *Guernica* (1937), which conveys the collective trauma and horror of the Spanish Civil War (panel b). Depictions of ritual and daily life reveal contrasting social realities: Pierre-

Auguste Renoir's portrayal of the Parisian bourgeoisie (panel c) reflects the emerging culture of urban leisure, growing prosperity, and the rhythms of modern life during the *Belle Époque*, while Vincent van Gogh's representation of peasant life (panel d) highlights the harsh conditions of rural poverty beyond urban centers. In some cases, the message is more symbolic or indirect. For instance, Hieronymus Bosch's *The Garden of Earthly Delights* captures fantastical creatures during a time of exploration and discovery, instilling both excitement and anxiety in a period marked by uncertainty. While the previous examples contrast different historical contexts, paintings can also offer multiple, distinct interpretations of the same subject, shaped by the artist's perspective: William-Adolphe Bouguereau and Gustave Courbet both depict rural labor in 19th-century France, yet Bouguereau presents a romanticized pastoral ideal, while Courbet delivers a realist critique of social hardship. These examples illustrate that paintings capture dimensions of historical experience typically absent from standard data sources—how societies perceived themselves, what they valued, and what they feared—enriching our understanding of past preferences, beliefs, and perceptions of living standards.

This paper leverages large, open-access repositories of artwork to develop an algorithm capable of predicting emotions conveyed through paintings—such as *sadness*, *fear*, *anger*, *awe*, *contentment*, *excitement*, *amusement*—and systematically analyzes artworks as reflections of their historical and cultural contexts. Our findings show that the emotional content of paintings fluctuates in ways that align with well-known economic and societal transformations, including those associated with Medieval Europe, the Renaissance, the Reformation, and the Enlightenment. For example, depictions of *fear* increase during turbulent periods, while *contentment* is frequently expressed during times of stability; *sadness* tends to be more prominent during times of hardship, at the expense of *amusement* or *excitement*. In addition, variation in emotional expression across contemporaneous artworks appears to follow the trend of economic inequality. The fine-grained nature of our data allows us to exploit subtle differences in dynamics across different locations and to identify expressive changes within the oeuvre of a single artist. This granularity, combined with a demanding empirical specification, enables us to detect localized emotional responses to historical events, including climatic fluctuations and food security, trade shifts, technological innovations, the dynamics of knowledge production, and political transitions. Together, these insights offer a new window into the lived emotional experience of societal transformation(s) across countries from 1400 onward.

To extract the contextual information embedded in historical paintings, we exploit open-access repositories from Google Arts and Culture (139,371 paintings), WikiArt (264,821 paintings), and Wikidata (226,654 paintings). Our empirical strategy develops a state-of-the-art learning algorithm based on Convolutional Neural Networks (CNNs) to identify the emotions conveyed through a painting. The model is trained using large-scale annotations comprising over 1.5 million emotion labels from prior studies (Achlioptas et al., 2021; Mohamed et al., 2022), covering 79,860 paintings from WikiArt. Annotators assigned a single dominant emotion to each painting from nine predefined categories: *anger*, *fear*, *sadness*, *disgust*, *awe*, *contentment*, *amusement*, *excitement*, and the residual category *other*. Alongside each label, annotators provided free-text descriptions justifying their choice. We rely on these justifications to understand the quality of annotations, notably (i) benchmarking emotion labels against a sample of artists' own descriptions of their work ("external validity") and (ii) analyzing the semantic alignment between each selected emotion and its textual justification ("internal validity"). In practice, these annotators can be more or less "trained" in interpreting artwork from certain regions or periods, depending on their own cultural background and their knowledge of art and of the historical context. We find that, while the 15-25 distinct annotators labeling a given artwork may disagree, there is a strong wisdom-of-the-crowd effect: the distribution of annotated emotions closely aligns with artists' stated intentions.

Our learning algorithm is optimized to best predict a probabilistic vector of emotions, for *any* given painting. The neural network architecture consists of two components: a pre-estimated "encoder" that extracts features from images—typically capturing patterns or objects—and a "regression head" that maps these extracted features into probabilistic emotion scores. The former is common to most image recognition models, usually trained to process "real-life" pictures. The latter is instrumental to our purpose: the regression head and its interaction with the encoder model, which is re-estimated in a sequential manner, help capture the nature of paintings and how they convey emotions through their composition or "texture." The minimization target—the Kullback-Leibler divergence, an entropy measure capturing distances between probability distributions—is hard to interpret, but we provide an array of evidence to show that (i) our model closely and consistently matches the targeted emotion probabilities across periods, exercises, or specific emotions and (ii) generalizes well beyond the subsample of paintings in our annotated dataset.[1]

We apply our learning algorithm to 630,846 paintings, generating a set of predicted emotion probabilities that reflect how a non-specialist viewer might perceive their emotional tone. This perceived tone may echo the more or less permanent mindset of the artist, the prevailing mood of their contemporary audience, and other influences shaped

[1] The validity of our procedure relies on two hypotheses: (H1) Can the model predict emotions, as actually (and hypothetically) provided by annotators? (H2) Do these contemporary predictions capture the intentions of the artist? We validate the model in three steps: (i) we analyze the model fit and over-fit, on average and across periods and exercises; (ii) we produce visuals and heatmaps for a set of images with a dominant emotion; and (iii) we benchmark our predictions against actual descriptions from artists.

by their (changing) socio-economic context. Our focus is the distribution of these perceived emotions across paintings produced within a specific context—defined as a particular location and year—adjusted for artist-specific characteristics and potential misperceptions tied to particular exercises or genres.[2] One concern is that preserved artworks in our data sources may reflect the tastes of both contemporary and later audiences, and historical contingencies, such as wealth accumulation, inheritance, museum donations, or destruction. To assess potential selection biases and the impact of commercialization on emotional expression, we analyze over one million art transactions from 1650 to 1950. We find a small price premium of approximately 10% for artists in our sample. Differences between artists whose works were or were not sold are limited, and we find no significant variation in emotional content between highly valued and lower-valued artists. While we cannot rule out selection into which artworks have survived, our findings suggest that these effects are unlikely to be the primary driver of our results.

We use this new measure—emotions extracted from historical artworks—as a novel lens to examine how societies perceived major socio-economic transformations. As a first step, we analyze aggregate patterns of emotions in 12 selected countries to capture both the prevailing mood and the emotional diversity within each context. Our aggregate-level analysis indicates that *excitement*, *amusement*, and *awe* are more common during times of improved living standards; *sadness* is more prominent during periods of hardship, particularly for vulnerable populations; and emotions such as *contentment* and *fear* tend to correspond to periods of stability and uncertainty, respectively. Importantly, while the average dynamics of emotions within a context are informative, the degree of disagreement across artworks produced at the same time and place also carries explanatory power, especially in relation to economic inequality. We formalize and validate these interpretations by correlating (i) the predicted probabilities of evoking each emotion, (ii) a *positivity* index, and (iii) a *disagreement* index, with conventional measures of economic development, economic uncertainty, and economic inequality, where and when such data are available (using, among others, data from Bolt and Van Zanden, 2025; Alvaredo et al., 2020).

In a second step, we examine emotional expressions at the level of individual artworks to investigate how societies responded to structural economic changes, development trajectories, short-term fluctuations in livelihoods, and geopolitical shocks. Conditioning on artist, location, year fixed effects, and 40 harmonized artistic genres, we systematically investigate five key response dimensions, denoted R1 through R5. (R1) We study how emotions respond to climatic variability in Europe between 1500–2000,

[2] To account for these factors, each painting is linked to rich metadata, including artist biography, year and place of production, artistic movement, genre, and physical attributes such as material and dimensions.

and to food insecurity, emphasizing the fragility of pre-industrial societies (otherwise illustrated in Nunn and Qian, 2011; Waldinger, 2022, focusing on the introduction of the potato or the Little Ice Age). (R2) We explore how trade shapes emotional expressions across time and space. (R3) We assess the emotional footprint of technological change, focusing on the diffusion of the steamship (Pascali, 2017) and radio (including its propagandistic uses; see Adena et al., 2015). (R4) We examine how the rise of scientific knowledge production and the concurrent decline of religiosity during the Enlightenment shaped collective sentiment. (R5) Finally, we analyze the relationship between depicted emotions and the political context, showing how moods shift in response to external conflicts, autocracy, political consolidation, and power transitions.

While well-studied drivers of living standards such as agricultural shocks (R1), trade (R2), and technological change (R3) are clearly reflected in our emotion metrics, we also find strong responses to dimensions that have received less attention in the empirical economic history literature. These include factors related to the rigidity of social order, such as the gradual shift in values associated with the Enlightenment—emphasizing freedom, rationality, experimentation, and collective progress (R4, echoing Mokyr, 2005)—as well as the stability of political regimes and the emotional impact of political transitions (R5). Together, these results provide new empirical evidence that emotional expression in art can trace both localized disruptions and long-run societal transformations.

The main contributions of our research are twofold: (i) we show the predictive power of artwork as measurement of socioeconomic contexts; and (ii) we use these measures to provide new insights into the dynamics of economic development, uncertainty, and inequality across time and space. More specifically, we construct a novel measure of emotions from a new source of data—historical paintings produced by diverse artists—to shed light on how societies perceived and responded to socioeconomic transformations. In doing so, we depart from well-established contributions in economic history that rely on more direct and objective indicators of past economic conditions (e.g., Bairoch, 1988; Bolt and Van Zanden, 2014; Bolt et al., 2018; Clark, 2005, 2007; Schularick and Taylor, 2012). Our work differs both in the nature of the data and in the nature of the measure itself—an indirect, subjective depiction of socioeconomic conditions.[3]

[3]Prominent contributions to our understanding of historical economic performance across regions come from Angus Maddison (Maddison, 2007) and the subsequent Maddison Project (Bolt and Van Zanden, 2014; Bolt et al., 2018; Bolt and Van Zanden, 2025), or work by Bairoch on city populations (Bairoch, 1988). These indicators have been instrumental in understanding the patterns of growth, structural change and urbanization in the very long run (see, e.g., Acemoglu et al., 2005a; North, 2010; Fujita et al., 2001; Galor, 2005, 2011). Other contributions have constructed economic indicators covering smaller regions in the long run (see, e.g., Clark, 2005, 2007; Bouscasse et al., 2024, in England) or a larger set of countries but in more recent times (see, for instance, Schularick and Taylor, 2012; Jacks et al., 2011; Jordà et al., 2017, providing various macroeconomic indicators or commodity prices for the most advanced economies).

The most closely related contributions are Safra et al. (2020), who use an algorithm to predict perceived trustworthiness in approximately 2,000 English portraits from 1505 to 2016, and Hills et al. (2019), who infer welfare trends in four countries through a textual analysis of books. Like these earlier contributions, our approach offers a proof of concept that indirect depictions of society, preserved in cultural artifacts, can be used as windows into their times and socioeconomic contexts. However, we differ in three key aspects: (i) the scale of our exercise, covering about 630,000 paintings, 29,000 artists, and 34 present-day countries with at least 1,000 artworks each; (ii) the richness of our measurement, using the overall composition of a painting, i.e., shapes, texture, content, and color, to predict complex emotional messages; and (iii) a wide-ranging analysis of perceived responses to shocks across five societal domains: rural livelihoods, trade dynamics, technology adoption, knowledge production, and political events.

Our study contributes to a growing body of research that infers economic development from indirect data sources, imagery in particular. The novel use of images as data is best illustrated in a recent contribution studying societal evolutions through photographs (Voth and Yanagizawa-Drott, 2024). This emerging literature also includes studies based on satellite imagery (Overman et al., 2006; Chen and Nordhaus, 2011; Henderson et al., 2012; Jean et al., 2016; Donaldson and Storeygard, 2016) or street views (Gebru et al., 2017; Naik et al., 2017). Our procedure relies on deep convolutional neural networks (see Dell, 2025, for a review), and the key contribution of our study is to extend these techniques to the measurement of economic activity in the distant past, rather than improving the measurement of contemporary economic activity (e.g., at more disaggregated levels or in setting with limited data coverage).

Our research also relates to the literature using arts and creative work in economics. One strand of research studies the life of artists to understand what drives their productivity and creativity (Galenson and Weinberg, 2000, 2001; Hellmanzik, 2010; Borowiecki, 2017, 2022). A second strand of research employs artistic production to better understand the importance of protecting novel ideas (Giorcelli and Moser, 2020; Whitaker and Kräussl, 2020); and finally, there is a strand of literature that studies art markets (Ginsburgh and Jeanfils, 1995; Ashenfelter and Graddy, 2006; Ursprung and Wiermann, 2011; Graddy, 2013; Renneboog and Spaenjers, 2013; Spaenjers et al., 2015; Etro, 2018; Etro and Stepanova, 2021) and art consumption (Scitovsky, 1972; Throsby, 1994). We contribute to this literature by highlighting another channel through which art informs our understanding of the economy: the preserved messages conveyed by artists themselves.[4]

[4]This complements the work of Logothetis et al. (2025), which shows how historical critical junctures can influence artistic expression across generations, shaping how music reflects and transmits cultural memory and identity.

Finally, our work indirectly relates to the literature discussing the measurement of economic activity or welfare (as surveyed in Hulten and Nakamura, 2022). Recent discussions have emerged to better account for aggregate welfare and find alternatives to the Gross Domestic Product (GDP), e.g., the Human Development Index and related measures of happiness (Benjamin et al., 2024). Our approach rests on the premise that conveyed emotions can offer novel insights into subtle societal responses—either complementing or, where absent, substituting for traditional data sources. For instance, expressions of *fear* (against *contentment*) capture periods of uncertainty in ways that parallel more conventional economic indicators (Bloom, 2014; Jurado et al., 2015).[5] Despite this potential, emotions have remained unconventional metrics in economics, likely due to long-standing measurement challenges. While early thinkers such as Bentham acknowledged that preferences—and thus choices—are rooted in emotions, these constructs have historically been more difficult to quantify than observable variables like prices, production costs, or output. However, recent interdisciplinary work has reasserted the importance of emotions in shaping behavior and decision-making (Elster, 1998; Damasio, 2006), influencing risk preferences (Meier, 2022), and informing policy attitudes (Algan et al., 2025).[6] A key contribution of this study is the construction of historical time series of emotions. These series offer insights into underlying societal values and beliefs, shedding light on the psychological foundations of collective behavior as reflected in policy preferences, political movements, and broader cultural transformations.

Our analysis is structured as follows. Section 1 provides insights from art history and psychology about our working hypotheses. Sections 2 and 3 introduce the data and discuss the modelling approach to produce a vector of emotions associated with each painting. Section 4 presents our strategy, sheds light on the dynamics of emotions, and studies the response to different socioeconomic transformations. Section 5 concludes.

1 Artwork and emotions as a window into socio-economic contexts

Our formal approach uses a learning algorithm to explore historical contexts through the lens of painters who documented their times. Before developing such an approach, this chapter will briefly summarize insights from the fields of art history, anthropology, evolutionary biology, and psychology that inform our working hypotheses.

[5]Our emotion categories align with the concept of *basic emotions* (Ekman, 1999), which are thought to emerge from evolutionarily adaptive tasks and social learning. Because these emotions are universal across cultures and relatively stable across time and space (Tooby and Cosmides, 1990), they serve as reliable indicators of how past societies responded to changing material conditions driven by economic, social, and political transformations.

[6]For more comprehensive discussions of the role of emotions in economics, see DellaVigna (2009); Wälde and Moors (2017).

The joint evolution of art and societies Research into the evolutionary roots of artistic behavior suggests that aesthetic practices were essential components of human cultural and cognitive evolution (Dissanayake, 1992; Dutton, 2009; Brown, 2021), and art has long been acknowledged as a critical factor in the evolution of societies and their cultures, acting both as a reflection of and a catalyst for social change. Recent contributions from economics echo this perspective, e.g., Michalopoulos and Xue (2021), or Michalopoulos (2025) for a more comprehensive discussion.

An important insight from the literature is that art serves not only as an aesthetic experience but also as a record of the emotional and social life of individuals and communities (a reflection of society; see, e.g., Albrecht, 1954; Alexander, 2020). Artworks aim to convey emotions, ideas, and knowledge. This reflective function underlies our own approach, which rests on the premise that the evolution of artwork can reveal the evolution of societies. However, artistic expression also shapes social interactions and helps societies articulate cultural patterns (a catalyst channel; see Dissanayake, 1992; Gell, 1998). Indeed, interdisciplinary research at the intersection of anthropology, evolutionary biology, neuroscience, and art history paints a complex picture in which art is not merely a byproduct of human civilization but an *active* participant in the cognitive and social development of humans (Henrich et al., 2016).

Quantifying the meaning of artwork Art historians seek to recover the meaning of artwork, employing methods that have influenced our work at different stages. We broadly categorize these methodologies into three groups: (i) those focusing on the visual and symbolic aspects of the artwork, such as formalism and iconography; (ii) those focusing on the artist, such as psychoanalytic art history, which explores artists' subconscious influences on art, and biographical approaches that consider the artist's personal life and intentions; and (iii) those focusing on the broader social context, which studies the social, political, and economic conditions surrounding the creation of artworks. We draw inspiration from all these areas in our attempt to decode the information that artists embedded in their work and to relate it to societal developments.

Our algorithmic approach emulates formalism and iconography to predict emotions expressed in the artwork. The algorithm flexibly accounts for the *Elements of Art*, e.g., color, form, line, shape, space, and texture, and for the *Principles of Art*, e.g., scale, proportion, unity, variety, rhythm, mass, shape, space, balance, volume, perspective, and depth. The attempt to link art elements and composition principles to emotions is perhaps best described by Leo Tolstoy's (1898) description of the artist's creative process: "To evoke in oneself a feeling one has once experienced, and having evoked it in oneself, then by means of movements, lines, colors, sounds, or forms expressed in words, so to transmit that feeling that others may experience the same feeling." Our approach essentially inverts this process: we observe the emotions associated with artworks by a large set of individuals and train an algorithm to learn how visual elements and principles of art translate into the feelings perceived by viewers.

Our procedure to quantify the meaning of art is closely related to the emerging field of digital art history, which explores how quantitative and computational techniques can enrich traditional art-historical inquiry. A 2013 special issue of the journal *Visual Resources* provided scholars with a forum to explore future directions for the field. Among the areas of interest identified were algorithmic computational methods, including machine learning and natural language processing. As noted by Jaskot (2019), "digital art history lets us address the tradition of the social history of art in new ways," because digital methods are particularly well suited to engaging with broader art-historical debates that examine artworks in their social context and require large bodies of evidence. This objective aligns with the core of the social history of art, which "is not satisfied with a social context for art, but rather reverses this equation by arguing that an analysis of art, artist, and audience must tell us something structurally about society." Building on this intuition, we propose a method to uncover what paintings can reveal about society and its historical circumstances, as we discuss next.

Societal influence, creativity, and personal circumstances Retrieving societal influence from the painter's creative process implies two difficulties. First, predicted emotions consist of individual, artist-specific effects and broader societal effects. For instance, there was a shift in Goya's way of painting after 1793 which was likely driven by personal circumstances (a near-fatal illness) and societal developments (political turmoil in Spain and the effects of the Napoleonic wars). This example illustrates how both personal and social factors find reflection in artwork. In this paper, our focus lies on the latter—the social effects. To disentangle social effects from personal idiosyncrasies, our empirical design controls for artist-specific circumstances and identifies group effects, i.e., the commonality behind artists' joint reflections on historical conditions and social changes. Second, a key working hypothesis is that the message and emotions perceived by a contemporary audience are a proper reflection of the message and emotions intended by artists of very different backgrounds.[7] The following sections will provide

[7] A related question concerns the stability of emotional expression across time and cultures. This question connects our work to (i) psychological research on the evolutionary basis of emotions, to (ii) work on the neuroscience of emotions (Damasio, 2006), and to (iii) psychoanalytic art history that explores emotions in art. While a full discussion of these literatures lies beyond the scope of this paper, we refer the interested reader to work by Ekman on basic emotions, Damasio on emotions as support to the reasoning process, and Robinson on emotions in art (Robinson, 2005). Our analysis focuses on the detection of basic emotions in historical paintings. Ekman (1999) argues that basic emotions are evolutionary

empirical support for these working hypotheses.

2 Data sources

This section describes our main data sources. Our baseline empirical analysis relies on two main sources: (i) large, open-access repositories of art collections, where each painting is associated with an artist, an estimated year of production, and a geographic location; and (ii) a smaller sample of paintings with manually annotated emotions. We provide further descriptive statistics on these data sources in Appendix B, including the available information, their coverage, and a detailed description of the annotations.

Figure 2. Temporal and spatial coverage of our image collection(s).

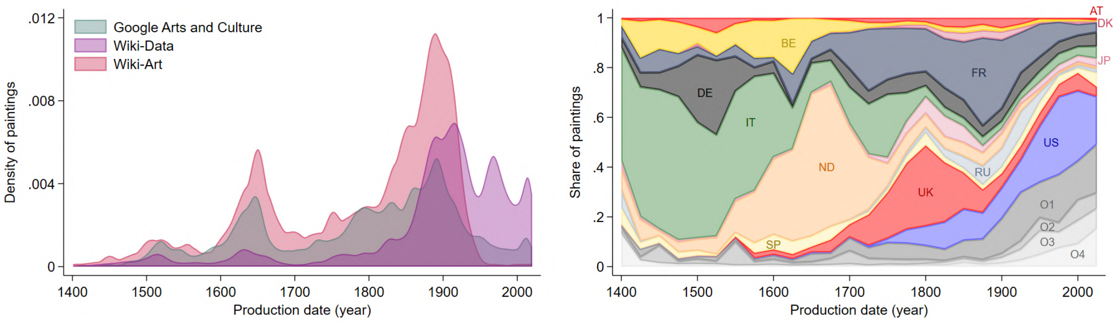

(a) Time of production (b) Selected countries

Notes: Panel (a) shows the distribution of production years in *Google Arts and Culture*, *Wiki-Art*, and *Wiki-Data*. *Google Arts and Culture* and *Wiki-Art* provide more consistent coverage of the Renaissance, the later Reformation, and the Enlightenment. The two main modes correspond to the Dutch Golden Age and the French Belle Époque. *Wiki-Data* primarily covers the more recent period from 1800 to 2000. Panel (b) shows the relative coverage over time for the following "harmonized" countries and country groups: Austria (AT, red), Belgium (BE, gold), Denmark (DK, light red), France (FR, navy blue), Germany (DE, black), Italy (IT, green), Japan (JP, pink), the Netherlands (ND, orange), Russia (RU, light blue), Spain (SP, yellow), the United Kingdom (UK, red), the United States (US, blue), other major European countries with more than 1,000 artworks (O1, e.g., Portugal, Sweden), other major Asian countries (O2, e.g., China, India), other major producers from other continents (O3, e.g., Argentina, Australia, Brazil, Mexico, South Africa), and smaller producers (O4, all countries with fewer than 1,000 artworks). We provide additional descriptive statistics on the coverage of these data sources in Appendix B.1, including their representation of different movements and types of artwork over time.

A harmonized collection of paintings Our primary dataset is a collection of high-resolution images, retrieved from three large repositories of art collections: Google Arts

adaptations to recurring life challenges. We adopt a commonly used framework that mostly relates to six such emotions: happiness; surprise; fear; disgust; anger; and sadness. According to Frijda (2007), each basic emotion is triggered by specific environmental stimuli and is associated with distinct behavioral responses. For example, disgust arises from contact with potentially harmful substances, prompting withdrawal and physiological reactions like retching; fear is evoked by perceived threats, leading to defensive responses like freezing or fleeing. These patterns reflect species-wide adaptive mechanisms that evolved to solve recurrent survival problems, and such emotions are universal and largely independent of cultural variation (Tooby and Cosmides, 1990).

and Culture; Wiki-Art; and Wiki-Data.[8] We collect information from *Google Arts and Culture* and *Wiki-Art* using a headless browser, which downloads images at the highest available resolution, along with tags (e.g., production year, style, or movement), and retrieves additional details from the artists' Wikipedia pages. Access to *Wiki-Data* is facilitated through an API. The data contain high-quality information on each painting's production date and physical characteristics (e.g., material and size).

Figure 2 presents the temporal coverage of these distinct data sources. *Google Arts and Culture* and *Wiki-Art* offer consistent coverage of the Renaissance, the later Reformation, and the Enlightenment, as well as the golden ages of painting (such as the Dutch Golden Age and the French Belle Époque). *Wiki-Data*, by contrast, is more focused on the period from 1800 to 2000 (panel a). Combined, these sources provide consistent coverage of the major European economies between 1400 and 2000 (panel b). Panel (b) of Figure 2 also illustrates how the geography of cultural production shifts over time in our curated sample: Italy is the dominant producer before 1600 (with the Holy Roman Empire also prominent around 1500); the Netherlands and the Low Countries lead from the late fifteenth to early seventeenth century; France emerges as a major center from 1700 to 1950; Great Britain peaks around 1800; and the United States, along with other world regions, gain prominence from 1900 onward.

Our data sources—*Google Arts and Culture*, *Wiki-Art*, and *Wiki-Data*—provide partial information about the artist or the artwork. To enrich and harmonize this information, we employ several strategies to associate each painting with approximately 100 harmonized artistic movements (e.g., post-impressionism, mannerism, fauvism), 40 harmonized exercises (e.g., allegorical painting, landscape, still life), and standardized artist identifiers, along with data on the artist's nationality, birth and death details, and possible main residences over the course of their career.[9]

This procedure yields a dataset of 630,846 paintings by more than 29,000 artists, for

[8]There is some overlap between these collections, e.g., for the most notable artwork. To identify duplicates, we rely on a fuzzy matching procedure and apply the algorithm described in Section 3 which produces—in an intermediary step—a vector of 1,260 embeddings for each painting. More specifically, we focus on paintings from the same artist within a given period, construct the squared difference between embeddings of separate pieces across different collections, and predict a probability to be "duplicates". The last step is informed by the constitution of a small training sample, in which we flag duplicates and non-duplicates for about 1,000 pairs of paintings from the same artists. The convenient representation of images into embeddings—allowing to create multi-dimensional indicators of distance—is used by a few recent contributions attempting to estimate demand for differentiated products (see, e.g., Han and Lee, 2025; Compiani et al., 2025).

[9]We rely on three main strategies. First, many artworks in *Google Arts and Culture*, *Wiki-Art*, and *Wiki-Data* include sufficient information to retrieve the artist's *Wikipedia* page, from which we request structured data via the Wikipedia API. Second, we use pre-trained Named Entity Recognition models to systematically clean and harmonize tags across sources (e.g., artist name, genre, art movement), which we complement with tagging information from over 500,000 paintings in the ART500K dataset (Mao et al., 2017). Third, we use a large language model (GPT-4o) to classify the remaining artists and paintings.

which we have reliable, comprehensive information on artwork characteristics (title, movement, genre, physical features), context of production (year, location), and creator identity. These geolocated and dated artworks span 110 contemporary countries, with 12 (and 34, respectively) accounting for more than 10,000 (resp. 1,000) distinct pieces. The leading contributors include Austria, Belgium, Denmark, France (115,286), Germany (40,792), Italy (60,375), Japan, the Netherlands (69,723), Russia, Spain, the United Kingdom (56,470), and the United States (68,454).

Figure 3. Annotations and the highest emotion scores.

(a) Contentment

(b) Amusement

(c) Excitement

(d) Awe

(e) Fear

(f) Anger

(g) Sadness

(h) Disgust

Notes: This figure shows the paintings with the highest emotion scores among 79,860 annotated artworks, by emotion (*amusement*, *anger*, *awe*, *contentment*, *disgust*, *excitement*, *fear*, *sadness*, and *other*—not shown). The most frequent emotions are: *contentment* (typically associated with landscapes, nature morte, and peaceful scenes); *awe* (typically associated with displays of power or opulence); *sadness* (often linked to melancholic portraits, depictions of poverty, or religious imagery); *fear* (typically associated with chaotic images, either through the subject or through texture—a ship engulfed by waves, for instance); *excitement*; *amusement*; and *disgust* (frequently associated with nudity). Sources: ArtEmis (Achlioptas et al., 2021, approximately 450,000 annotations), ArtELingo (Mohamed et al., 2022, approximately 1,000,000 annotations). We collapse the different emotion ratings for each painting into an empirical probability distribution, i.e., the probability that a labeler reports a given emotion.

A collection of *annotated* images Our empirical procedure requires labeled data, i.e., a subsample of paintings with associated annotations (e.g., emotion intensity), in order to map a painting j onto a vector of emotion $\mathbf{p}_j$. We rely on ArtEmis (Achlioptas et al., 2021, which provides approximately 450,000 annotations for a subset of 79,860 Wiki-Art paintings discussed above) and ArtELingo (Mohamed et al., 2022, which provides over 1,000,000 annotations for the same 79,860 paintings). These two sources collect emotion scores along nine dimensions: *amusement*, *anger*, *awe*, *contentment*, *disgust*, *excitement*, *fear*, *sadness*, and *something else* (which we hereafter label as *other*). In practice, individuals were shown paintings and asked to select one dominant emotion from the nine categories. ArtEmis includes an average of 5.6 annotations per painting, while ArtELingo includes an average of 15.3; each annotation corresponds to a different evaluator. We treat these multiple annotations as if they resulted from the probabilistic evaluation of a single individual: each painting j is allocated a probabilistic vector of emotions, $\mathbf{q}_j$, based on the 15-25 labels provided by the various annotators. This vector records the share of evaluators who selected each of the nine emotions, such that the entries sum to 1.

The most frequently mentioned emotions are: *contentment* (often associated with landscapes); *awe* (typically linked to beauty or displays of power and opulence); *sadness* (associated with melancholic portraits, depictions of poverty, or religious imagery); *fear* (often linked to chaotic scenes, either through subject matter or visual texture—such as a ship engulfed by waves); *excitement*; *amusement*; and *disgust*.

Figure 4. Understanding the reported emotion from the annotators' free descriptions.

(a) Contentment (0.302) (b) Amusement (0.101) (c) Excitement (0.071) (d) Awe (0.168)

(e) Fear (0.095) (f) Anger (0.018) (g) Sadness (0.101) (h) Disgust (0.048)

Notes: These figures rely on *Term Frequency-Inverse Document Frequency* (TF-IDF) to measure the importance or distinctiveness of a word within a corpus, and we cluster these words by emotion labels. TF-IDF captures the distinctiveness of a term in a document relative to the corpus, with $\mathrm{idf}(t) = \ln\left(\frac{1+n}{1+\mathrm{df}(t)}\right) + 1$ and $\mathrm{T}(t, d) = \mathrm{tf}(t, d) \times \mathrm{idf}(t)$, where $\mathrm{tf}(t, d)$ is the share of term t in document d, and $\mathrm{df}(t)$ is the number of documents—here, textual justifications by annotators—containing t in a corpus of size n. Intuitively, words that appear in a small number of documents receive higher scores than those that are common across all documents. We report the average probabilistic score of each emotion in parentheses.

How do annotators interpret the artist's intentions and categorize emotions in paintings? Different choices may reflect differences in the personalities of labelers, as well as in their training, cultural background, and knowledge of art and historical context. Conveniently, annotators select a dominant emotion but also provide a short description to justify their choice. We illustrate the mapping between paintings, annotators' perceptions, and their chosen emotions in Figures 3 and 4. Figure 3 shows paintings that received the highest scores among evaluators for each emotion: one can see that annotators respond to the subject (e.g., a clown for *amusement*), to color (e.g., the bright palette in the artwork associated with *excitement*), and to texture and composition (e.g., the sharp angles and layout in the artwork associated with *fear*). Another way to shed light on how annotators interpret emotion labels is to exploit the free text that they provide for each painting. Figure 4 visualizes the most frequently used terms in these justifications: *contentment* is linked to calm, peace, and nature; *sadness* to darkness and death; *fear* to both the subject (especially facial expressions or eyes) and the composition (e.g., darkness and shadows); and *disgust* appears to be triggered by bodily representations.

Figure 5. The internal and external validity of annotated emotions.

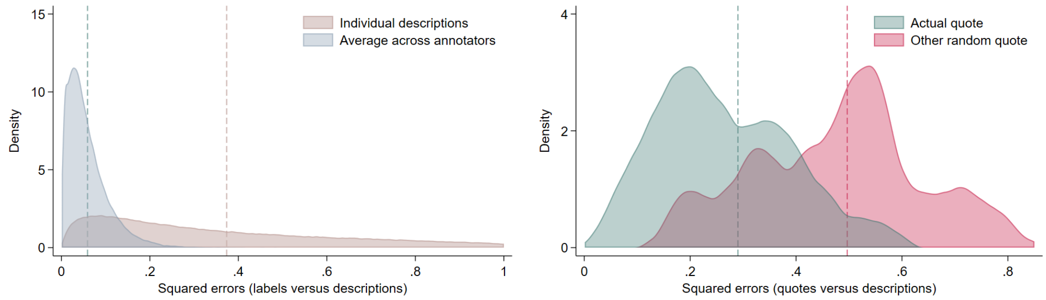

(a) Descriptions against labels

(b) Descriptions against quotes

Notes: Panel (a) displays the distribution of squared errors, $\sum_e (\vartheta^e - q^e)^2$, between the probabilistic vectors of emotions averaged across the dominant emotions chosen by annotators, $\{q^e\}_e$, and the probabilistic vectors of emotions associated with their textual justifications, $\{\vartheta^e\}_e$. The green line shows averages across textual justifications, while the beige line reports distances between the probabilistic vectors of emotions and each *individual* textual justification from each of the 15–25 annotators. Panel (b) displays the distribution of squared errors, $\sum_e (q^e - \rho^e)^2$, between the probabilistic vectors of emotions averaged across textual justifications provided by annotators, $\{q^e\}_e$, and the probabilistic vectors of emotions derived from quotes by the artists, $\{\rho^e\}_i$, across 34 distinct paintings (including: *The Potato Eaters* by Vincent van Gogh, depicted in Figure 1; *A Sunday Afternoon on the Island of La Grande Jatte* by Georges Seurat; *The Massacre of the Innocents* by Peter Paul Rubens; and *Moscow* (1916) by Wassily Kandinsky). The red distribution provides a benchmark in which quotes are randomly allocated across paintings. The textual classification of quotes into a probabilistic emotion score, $\{\rho^e\}_i$, relies on a transformer model (DistilBERT; see Sanh et al., 2019).

We further exploit these justifications to assess the suitability of the annotations for our purposes and ask: do these contemporary predictions accurately capture the artist's intentions? In a first step, we test the "internal validity" of these annotations and analyze the proximity between the average chosen emotions and the textual justification (panel a of Figure 5): the distribution of squared errors, $\sum_e (\vartheta^e - q^e)^2$, between the probabilistic vectors of emotions, as averaged across the dominant emotions chosen by annotators, $\{q^e\}_e$, and the probabilistic vectors of emotions associated with their textual justifications by a transformer model (DistilBERT, Sanh et al., 2019), $\{\vartheta^e\}_e$, features low values. One advantage of using a probabilistic target is that it allows for a wisdom-of-the-crowd effect, attenuating the influence of annotators' idiosyncrasies (e.g., their mood, cultural background, or knowledge of art and historical context).

In a second step, we benchmark these textual justifications against a selection of artists' own descriptions of their work, covering 34 distinct paintings (panel b of Figure 5). The average emotion vector, **q**, closely matches the artists' stated intentions.[10]

Socio-economic changes, shocks to livelihoods, and auctions The main empirical novelty of our research is to exploit paintings as data points. However, our analysis also requires access to more conventional data sources, most notably: indicators of economic development and inequality to validate and interpret fluctuations in measured emotions, such as GDP per capita (Bolt and Van Zanden, 2025), wealth concentration (Alvaredo et al., 2020), happiness (a valence index based on sentiment analysis of published books, see Hills et al., 2019), and life expectancy (Zijdeman and Ribeira da Silva, 2015); large, observable socio-economic changes, including the adoption of new technologies (Cross-country Historical Adoption of Technology, see Comin and Hobijn, 2004), the first wave of globalization (Jordà-Schularick-Taylor Macrohistory Database, see Jordà et al., 2017), and the intensification of knowledge production (The Rise of Universities in Medieval and Early Modern Europe, see De La Croix et al., 2024); shocks to livelihoods over the period of interest, such as political turbulence (Marshall et al., 2014), wars, climatic variability (Luterbacher et al., 2004), and short-term changes in the prices of basic commodities (Allen-Unger Global Commodity Prices Database, see Allen and Unger, 2019); and roughly 1,200,000 recorded transactions for a subset of artworks dated between 1650 and 1950, which we use to examine the selection of preserved artworks and the role of art commercialization. We describe these data sources in Appendix B.4.

3 An emotion prediction model

This section briefly describes the convolutional neural network architecture used to predict emotion scores from paintings. We defer a more detailed discussion to Appendix C.

3.1 Structure and optimization

Image pre-processing While convolutional networks typically require standardized input, for instance, square images of a uniform size, the raw images vary in both res-

[10] We provide supplementary material in Appendix B.2, where we discuss differences across movements and types of exercise (e.g., portraits, landscapes, religious or allegorical works), as well as measurement issues related to the number of ratings per image.

olution and aspect ratio. The canvas dimensions would, however, mostly fall within a longer-side-to-shorter-side ratio between 1:1 and 1:1.5 (see Figure C1 in the Appendix), and the nature of artwork usually implies that the center of the piece is its most important part. We therefore center-crop each image while preserving its aspect ratio, and then resize it to a $384 \times 384 \times 3$ array, where the last dimension represents the three color channels: red, green, and blue. For example, Figure 3 shows pre-processed paintings, demonstrating that cropping and resizing do not significantly affect visual perception.

A convolutional network structure We adopt a transfer learning approach to predict the emotions conveyed by paintings (i.e., our algorithm learns from initial, generic priors; see Bengio et al., 2013). Our model operates in two stages. First, a large, pre-trained image model—the encoder—extracts generic features from the image. These include visual elements such as edges, textures, motifs, shapes, and objects that capture the essential visual characteristics of the artwork. Second, an additional layer maps these features onto emotion probabilities.

First, the pre-trained encoder extracts multiple generic features from each image. This encoder consists of the convolutional layers of an EfficientNetV2-S model (Tan and Le, 2021), originally trained on the ImageNet dataset (Deng et al., 2009, a multinomial classification task involving 1.4 million photographs of objects from 1,000 different classes, with the output layer removed). The model produces an array of dimensions $12 \times 12 \times 1280$, where the 12×12 grid divides the input image into coarse spatial regions, and each of the 1,280 variables encodes the presence of a specific image feature within those regions. We then apply global average pooling by computing the mean value of each feature across all spatial locations, effectively ignoring any variation in importance between central figures and peripheral details.

Second, a regression head is added to the encoder, mapping the 1,280 average feature values into 9 scores associated with each emotion. This regression head—preceded by a dropout layer to mitigate overfitting (Tompson et al., 2014)—is more flexible than a basic multinomial logit model, allowing it to better capture the complexity of the encoder output, though it serves a similar function. The outcome is a distribution over emotions, with scores summing to 1 and interpretable as probabilities. Letting $j \in \mathcal{J}$ denote a painting, $\mathbf{x}_j$ the vector of extracted features from the encoder model, and y_j the latent emotion of painting j, the regression head assumes,

$$p_j^e = P\left(y_j = e|\mathbf{x}_j\right) = \frac{e^{z_{e,j}}}{\sum_{e \in \mathcal{E}} e^{z_{e,j}}},$$

where latent factors, $z_{e,j} = f_e(\mathbf{x}_j, \mathbf{b})$, are estimated through a non-parametric estimation, allowing to capture non-linearities and interactions between input variables.

The model minimizes the Kullback-Leibler divergence—or relative entropy—between the predicted and observed distributions of probabilistic scores, allowing us to account for the fact that the target is a probability distribution. The Kullback-Leibler divergence metric between the actual probability distribution, $\{q^e\}_{e\in\mathcal{E}}$, and the predicted probability distribution, $\{p^e\}_{e\in\mathcal{E}}$, is defined as,

$$\sum_{e\in\mathcal{E}} (q^e \ln(q^e) - q^e \ln(p^e)) .$$

Model parameters are estimated by maximum likelihood using the AdamW optimizer (i.e., with an adaptive learning rate for each parameter, dynamically adjusted learning rates, and sub-sampling of 32 images at each optimization step; see Loshchilov and Hutter, 2017). In practice, the model is implemented using the GPU-enabled version of the PyTorch library and can be run on a standard high-specification laptop with a properly configured Python environment; access to high-performance computing is not required.

Optimization The model optimization proceeds in two stages. First, we align the pre-trained encoder parameters with the randomly initialized parameters of the regression head. To do so, we optimize the model while keeping the encoder parameters fixed—effectively treating the encoder's output as given and updating only the regression head that maps these features onto emotion probabilities. Directly optimizing the entire model from the start could result in large gradient updates from the randomly initialized head, potentially disrupting the encoder's ability to extract meaningful features. Second, we fine-tune the full model by (i) unfreezing the encoder parameters—except for the batch normalization layers, which remain frozen to preserve the stable feature distributions acquired during pre-training—and (ii) training the model with a low learning rate to mitigate over-fitting. This step effectively adapts the model from the "photograph domain" (and more specifically, from ImageNet-style images) to the painting domain, encouraging it to extract features that are predictive of emotions.

The hierarchical structure of convolutional neural networks allows them to build increasingly complex representations by combining simpler features learned in earlier layers. As illustrated in Appendix C, early layers capture low-level characteristics such as brush strokes, color gradients, and textures; mid-level layers extract more abstract elements such as shapes and lighting; and later layers focus on high-level aspects including composition, subject matter, and expressive details such as facial expressions and posture.

Figure 6. Model validation—Kullback-Leibler divergence.

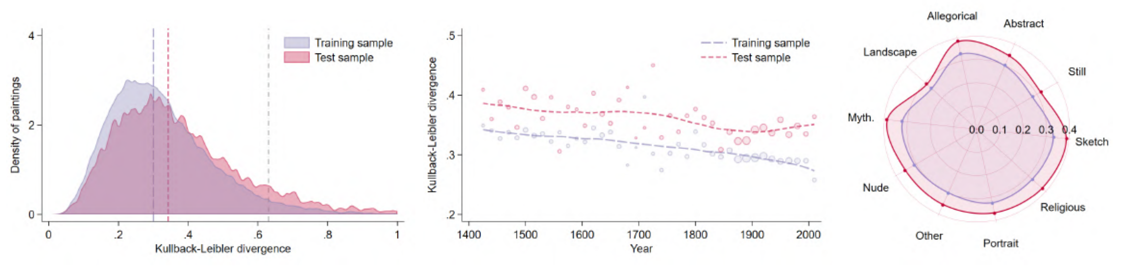

(a) Distribution (b) Period (c) Genre

Notes: These figures present statistics on the Kullback-Leibler divergence between the emotion ratings and the predicted ratings within the training sample and the test sample. The metric is defined as $\sum_{e\in\mathcal{E}}(q^e \ln(q^e) - q^e \ln(p^e))$, where $\{q^e\}_{e\in\mathcal{E}}$ is the target probability distribution and $\{p^e\}_{e\in\mathcal{E}}$ is the predicted distribution. Panel (a) shows the distribution of this divergence metric within the training sample (light blue) and the test sample (red); the median values for each are displayed as dashed lines (long dash for training, shorter dash for test). For interpretive purposes, we also report the median divergence between predictions and randomly reshuffled labels as a fully uninformed benchmark (dashed gray line). Panels (b) and (c) show median divergence values across sub-samples defined by: (i) production period, and (ii) genre or type of exercise. The following categories are used: *Abstract* (abstract, genre painting); *Landscape* (cityscape, landscape, marina); *Still* (animal painting, flower painting, still life); *Sketch* (sketch and study); *Nude* (nude painting); *Portrait* (portrait, self-portrait); *Religious* (religious painting); *Myth.* (mythological painting); and *Allegorical* (allegorical painting).

3.2 Validation

The validity of our overall procedure relies on two key hypotheses: (H1) the model predicts the emotions sensed by annotators, and (H2) the contemporary labels—and the resulting predictions—capture the artists' intentions. This section offers empirical evidence to assess model fit and potential over-fitting both overall and across genres and periods (H1) and provides a discussion of the less testable assumption that the predictions meaningfully reflect the emotions that the artist(s) intended to transmit (H2).

Fit and over-fitting We validate the model by splitting the annotated sample into three subsets: a training set (70%); a validation set (15%); and a test set (15%). The training set is used to estimate the parameters of the encoder and regression head; the validation set is used to interrupt training when the model starts overfitting and to select the parameters with the best generalization performance; and the test set is used to evaluate model performance.[11] Appendix C provides a comprehensive description of the model performance, and the main takeaways are illustrated in Figure 6.

First, the Kullback-Leibler divergence is similar across the training and test sets, indicating that the model is not subject to major over-fitting (panel a). The actual prediction

[11]To further limit over-fitting, random transformations, such as flipping, zooming, and rotation, are applied to each training batch, which increases the size of the training sample and provides invariance properties to these transformations. Since we also apply dropout layers to avoid over-fitting, we can use Monte-Carlo Dropout (Gal and Ghahramani, 2016) to compute standard errors for our predictions.

is four to five times more likely to outperform a reshuffled random prediction. Second, the divergence increases gradually with the time since production, in both the training and test sets (panel b). This appears to be driven by variation in sample size, as shown by the better fit for paintings produced between 1850 and 1950, which corresponds to the modal period in our harmonized collection of paintings (see Figure 2). Third, the divergence is slightly higher for sketches, nude, religious, and mythological paintings, reflecting their lower incidence in the annotated sample (panel c). Overall, however, the model successfully limits over-fitting across periods and artistic exercises.

The nature of the prediction, i.e., a probability distribution $\left\{\mathbf{p}_j, \forall j \in \mathcal{J}\right\}$, and its associated evaluation metric, namely the Kullback-Leibler divergence (or relative entropy), make it difficult to extract intuitive indicators of performance. To illustrate model accuracy in more accessible terms, we compute mean squared errors in Appendix C, where we show that the median mean squared error within the *training* sub-sample (0.056) is equivalent to allocating a probability of 0.830 to a dominant emotion and 0.170 to a secondary emotion, instead of a probability of 1 to the dominant emotion and 0 to all others. Given the limited over-fitting, the median mean squared error in the *test* set (0.067) is comparable—equivalent to assigning a probability of 0.815 to the dominant emotion and 0.185 to a secondary one.

To further illustrate the model's strengths and limitations, we extract eight additional images from the test sample, each selected to convey a specific emotion (see Appendix Figure C6). The model predicts well the most common emotions. For example, it identifies *awe* in depictions of opulent buildings or dramatic landscapes; *contentment* in peaceful natural scenes; *excitement* through warm colors, lighting, and human activity; *fear* through cold tones and indistinct textures that evoke uncertainty; and *sadness* via facial expressions or muted colors. The prediction is slightly weaker for less frequent emotions, as discussed below. The prediction also inherits the prejudices of evaluators. For instance, nudity is often associated with *disgust*, possibly reflecting moral codes and beauty standards of our times—maybe more than the artist's original intention.

The main prediction errors stem from two well-known issues. First, predicting a truly multidimensional emotion vector is difficult, especially given the skewed distribution of emotions in both the annotated data and the predictions. Subtle, secondary emotions are often captured less precisely than the primary, more salient emotional tone of an image. A second issue is neutrality versus distinct emotions. In effect, many images evoke *contentment* to evaluators, and the model tends to slightly over-predict such outcome at the expense of more distinctive emotions.

Identification One drawback of neural networks is that they are highly complex predictors, combining numerous parameters across connected layers. This complexity often makes it difficult to identify which sources of variation are most important for prediction. By contrast, regression models allow the statistician to measure the contribution of individual variables to explaining the overall variance in the outcome.

Figure 7. The identification of emotion in paintings.

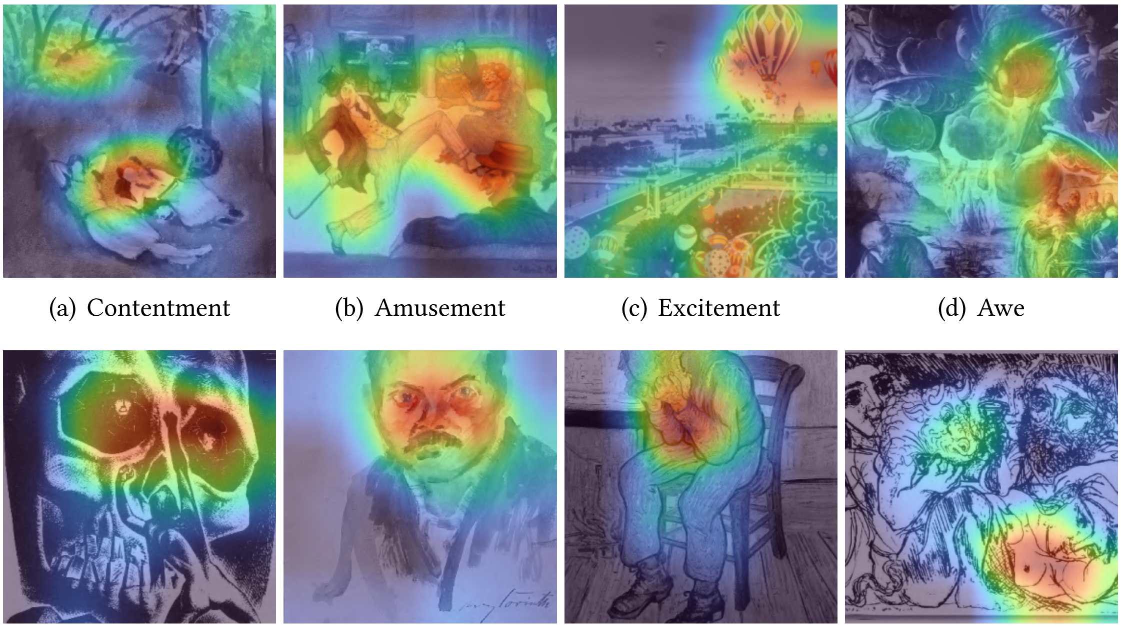

(a) Contentment (b) Amusement (c) Excitement (d) Awe

(e) Fear (f) Anger (g) Sadness (h) Disgust

Notes: This figure displays heatmaps generated using Class Activation Mapping (Zhou et al., 2016), which isolates the parts of the image most used by the network to predict a latent emotional probability. We apply this method to the paintings with the highest emotion scores—for each emotion—among 79,860 annotated artworks; these paintings were previously shown in Figure 3. The color scale, ranging from blue to green, yellow, and red, indicates the zones of activation associated with each dominant emotion. For example, *amusement* (panel b) is detected through the faces of the audience and the central dancing figure; the balloons in panel (c) evoke *excitement* to the trained algorithm; and depictions of naked bodies are predicted to elicit *disgust* (panel h). Additional illustrations are provided in Appendix C.3.

We cannot provide an equivalent decomposition. Instead, we illustrate the quality of the identification of probabilistic emotion scores through a set of examples. In Figure 3, we generate heatmaps using Class Activation Mapping (Zhou et al., 2016), which highlights the parts of an image most used by the network to predict a latent emotional probability. To this end, we use the eight paintings previously selected as most evocative of a dominant emotion (see Figure 3), and we report the heatmap for this dominant emotion. The results show that the algorithm often focuses on central, conspicuous elements, e.g., the angry face in panel (b), but can also infer emotion from background objects (balloons in panel c) or secondary characters (panel b). In general, bright colors and motion evoke *excitement*; sharp angles are associated with *fear*; and exposed body parts tend to signal *disgust.* In Appendix C.3, we provide additional evidence, including examples of Class Activation Mapping for different emotions within the same image.

Predictions and the intentions of artists In this section, we briefly summarize more indirect evidence that our contemporary predictions capture the artists' intentions. First, our predictions are essentially targeting the labeling process of annotators, such that any desirable or undesirable features of their annotation process will be projected onto the model. In Section 2, we showed that there is non-negligible heterogeneity across annotators, but that the average annotator appears to understand the emotion categories (see Figure 3), to be internally consistent in their classifications (see Figure 4 and panel a of Figure 5), and to broadly align with the stated intentions of artists, at least in the subset of examples where such intentions are available (see panel b of Figure 5). Second, the previous identification illustrations serve as sanity checks for a selection of artworks. Third, in Appendix C.3, we quantify the proximity between the predicted probabilistic emotion vectors and those derived from artist quotes.

Attrition, commercialization, and selection into conservation Finally, our harmonized collection of paintings reflects a selection of artworks produced within a specific historical context—a selection shaped by the tastes of contemporary and later audiences, the idiosyncrasies of wealth accumulation, inheritance patterns, museum donations, and the destruction of art through wars and revolutions. In Appendix B.4, we analyze digitized sales catalogs covering approximately 1.2 million transactions between 1650 and 1950 in the Low Countries, France, Great Britain, and Germany. This analysis allows us to examine the selection of artworks and artists into preservation and conservation, and to address two key questions: (1) Which artworks produced in a given period were selected for preservation and inclusion in today's digital collections? and (2) To what extent were preserved artists more highly valued around the time of production?

We find a moderate price premium—around 10%—for artists included in our sample, a difference that can be identified within the same auction. A related question concerns the role of commissioned art, and whether the artists in our sample—those whose works were commercialized or later appeared in digitized auctions—produced pieces that differed in technique or emotional expression from those of their peers. Here, we find limited differences between "selling" artists and others. Likewise, we find no evidence of a significant gap in emotional content between highly valued artists and those who commanded lower prices at the time.

4 The dynamics of emotions over time

The previous exercise produces a probabilistic vector of emotions, $\mathbf{p}_j$, for each painting $j \in \mathcal{J}$. This section describes how we exploit these painting-specific characteristics to examine the dynamics of emotions and their reflection of large socioeconomic changes.

4.1 Empirical strategy

Empirical strategy The data construction described in the previous Sections 2 and 3 yields a dataset of images with the following attributes: emotion scores, genre and movement, title, author, and year and location of production.

Let j denote a given painting and p_j^e an associated emotion indicator—for instance, the predicted probability to evoke emotion e. This image has been produced by a certain artist a in location ℓ and year t, characterized by socio-economic conditions $\mathbf{s}_{\ell,t}$. In this setting, the equivalent to a standard two-way fixed effects specification corresponds to the model,

$$p^e_{j,\ell,t} = \alpha \mathbf{s}_{\ell,t} + \beta \mathbf{X}_{j,\ell,t} + \nu_\ell + \eta_t + \varepsilon_{j,\ell,t}, \tag{1}$$

where $\mathbf{X}_{j,\ell,t}$ includes covariates such as fixed, unobserved artist effects, $\alpha_{a(j)}$, the artist's age, the type of exercise, or the adopted style; and ν_ℓ and η_t denote location and year fixed effects.

This specification is demanding, leveraging variation in conveyed emotions within an artist's production cycle. In Appendix D.1, we quantify the respective explanatory power of year, country, artist, and context fixed effects: country differences account for less than 2-3% of the overall variance in depicted emotions; common yearly dynamics across countries explain less than 8% for most emotions; unobserved, latent artist factors explain between 40 and 50% of the variance; and adding context-specific unobservables, (ℓ, t), to artist fixed effects increases the explained variance across our emotion indicators by 3-4 percentage points. The subsequent analysis shows that this "limited" within-artist variation is nonetheless sufficient to capture significant fluctuations in emotions across contexts. In Appendix D.1, we further leverage our empirical strategy using a sample of 30 prominent artists, highlighting variation in the content of their work and how their emotional expression may reflect broader societal changes.

One variant of the previous model would replace observable measures of socio-economic conditions, $\mathbf{s}_{\ell,t}$, with a set of dummy variables, $\psi_{\ell,t}$, associated with a context (ℓ, t). This alternative specification identifies the average latent emotion, $\psi_{\ell,t}$ in a setting that is well-known to the labor literature as an AKM decomposition (where artists play the role of "workers" and contexts correspond to "firms," with the objective of extracting firm productivity while isolating it from worker selection; see Abowd et al., 1999).[12]

Derivatives of the emotion scores Our empirical evidence primarily relies on the probabilistic vector of emotions, $\mathbf{p}_j$, associated with each painting j. In addition, we associate indices to this probabilistic vector of emotions. The first index combines the nine emotion scores into a single value,

$$l_j = \sum_{e \in \mathcal{E}^+} p_j^e - \sum_{e \in \mathcal{E}^-} p_j^e,$$

where l_j represents a *positivity* index ranging from -1 to 1, calculated by summing the positive emotion scores (*contentment*, *excitement*, *amusement*, *awe*) and subtracting the negative ones (*fear*, *anger*, *sadness*, *disgust*). This dimensionality reduction might help better visualize the general dynamics of emotions conveyed by different paintings, but such a simplistic index may overlook important nuances: both *fear* and *sadness* are negative emotions, but *fear* may be forward-looking and tied to uncertainty, while *sadness* is more often backward-looking; similarly, *contentment* and *excitement* are both positive, though *contentment* tends to reflect calmness. Emotions like *awe* and *disgust* are also more difficult to interpret or categorize.[13] We also construct an artwork-specific measure of *disagreement*, d_j, defined as the distance to the average emotion vector within a

[12]This decomposition relies on two assumptions: additivity in artist and context effects; and mean independence of the error term, i.e., the error is mean independent of previous (and future) contexts. Violations of these assumptions—discussed in recent contributions (see, e.g., Bonhomme et al., 2020)—can bias the estimated context-specific effects $\psi_{l,t}$. For instance, if there were complementarities between artist and context effects, the decomposition would fail because the effects are no longer additively separable. Similarly, if the error term was correlated with past or future contexts, the estimates would again be biased. Several factors may validate such concerns. First, there may be a "limited mobility bias." While artists *do* move across contexts in the sense that t varies across their works (which could theoretically be distinguished from life-cycle effects if age was modeled separately), they move less frequently across locations ℓ. Moreover, mobility patterns likely vary across historical eras, as apprenticeship systems were gradually replaced by formal training, and across the life-cycle, as apprenticeship gives way to more formal invitations. Second, mobility is not random. Artists are invited to specific places at specific times in response to their previous works, expectations about local economic conditions, or other unobserved factors. For example, François I, who directly influenced economic and cultural developments in France between 1515 and 1547, invited Leonardo da Vinci, Benvenuto Cellini, Andrea del Sarto, Rosso Fiorentino, Francesco Primaticcio, Joos van Cleve, and Godefroy le Batave, in part to disseminate the Italian and Northern Renaissance in France. Third, previous contexts might spill over to the future production of artists and introduce spurious time and spatial auto-correlation across context-specific effects $\psi_{\ell,t}$.

[13]To better understand the latent variation underlying our nine correlated variables, we perform a Principal Component Analysis (PCA) on the probabilistic vector of emotions, $\mathbf{p}_j$, and select the eigenvectors of the covariance matrix with eigenvalues above or close to 1. The loadings for each emotion are reported in Appendix D.1 and Appendix Table D1. The first principal component ($p_j^{c_1}$) loads positively on negative emotions (*fear*, *anger*, *sadness*) and negatively on *contentment*. The second principal component ($p_j^{c_2}$) loads positively on positive emotions (*amusement*) and *disgust*. The third component ($p_j^{c_3}$) captures *awe* (and *excitement*), while the fourth c($p_j^{c_4}$) predominantly reflects the residual category *other*.

context, $\bar{\mathbf{p}}_{\ell,t}$,

$$d_j = \frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} |p_j^e - \bar{p}_{\ell,t}^e|.$$

Commonality, disparity, and aggregation In the next section, we present descriptive evidence at the context level, defined by (ℓ, t). We perform two conceptually distinct aggregations: first, we compute average emotion indices within a context, capturing its commonality; second, we measure within-context dispersion, capturing its disparity.[14] In practice, the number of observations within each year-location cell may be limited. For visualization purposes (but not in any of our formal exercises), we interpolate the residuals along the time dimension (see, e.g., Figure 8). Specifically, we apply a Hodrick-Prescott filter using the standard smoothing parameter for annual data, as in business cycle analysis. Unlike that literature, however, our focus is on the smoothed series itself rather than deviations from it.

4.2 The dynamics of emotions across countries

This section describes the dynamics of emotions over time and sheds light on the perception of large socioeconomic changes.

The average emotions over time We begin by presenting descriptive statistics on the dynamics of emotions over time. Figure 8 illustrates the evolution of *fear* f, $p_{\ell,t}^f$, within the borders of twelve contemporary countries that have contributed more than 10,000 artworks: France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. The figure reveals significant variation in the depiction of *fear* through paintings, with some fluctuations reflecting long-term secular trends and others being more transient. Likewise, some of these changes are common across countries, while others are more idiosyncratic.

First, changes in depicted *fear* appear to correlate with major historical events: *fear* increases markedly during periods of conflicts, often by 3-4 percentage points within a few years—whether these conflicts involve multiple countries (e.g., World Wars I and II) or are more localized (e.g., the War of the Austrian Succession or the Spanish Civil War). Second, many European countries exhibit relatively higher expressions of *fear* during the Little Ice Age and throughout much of the twentieth century. Third, there are substantial idiosyncrasies in these dynamics: less than 2% of the variance is explained by common yearly fluctuations across countries or fixed cross-country variation.

[14]To assess divergent views within a context, one could alternatively compute average emotions for different groups of artists, for example, based on socio-economic status. While artists reflect the spirit of their time, they may have experienced different living standards depending on their social position.

Figure 8. The dynamics of *fear* across countries and over time.

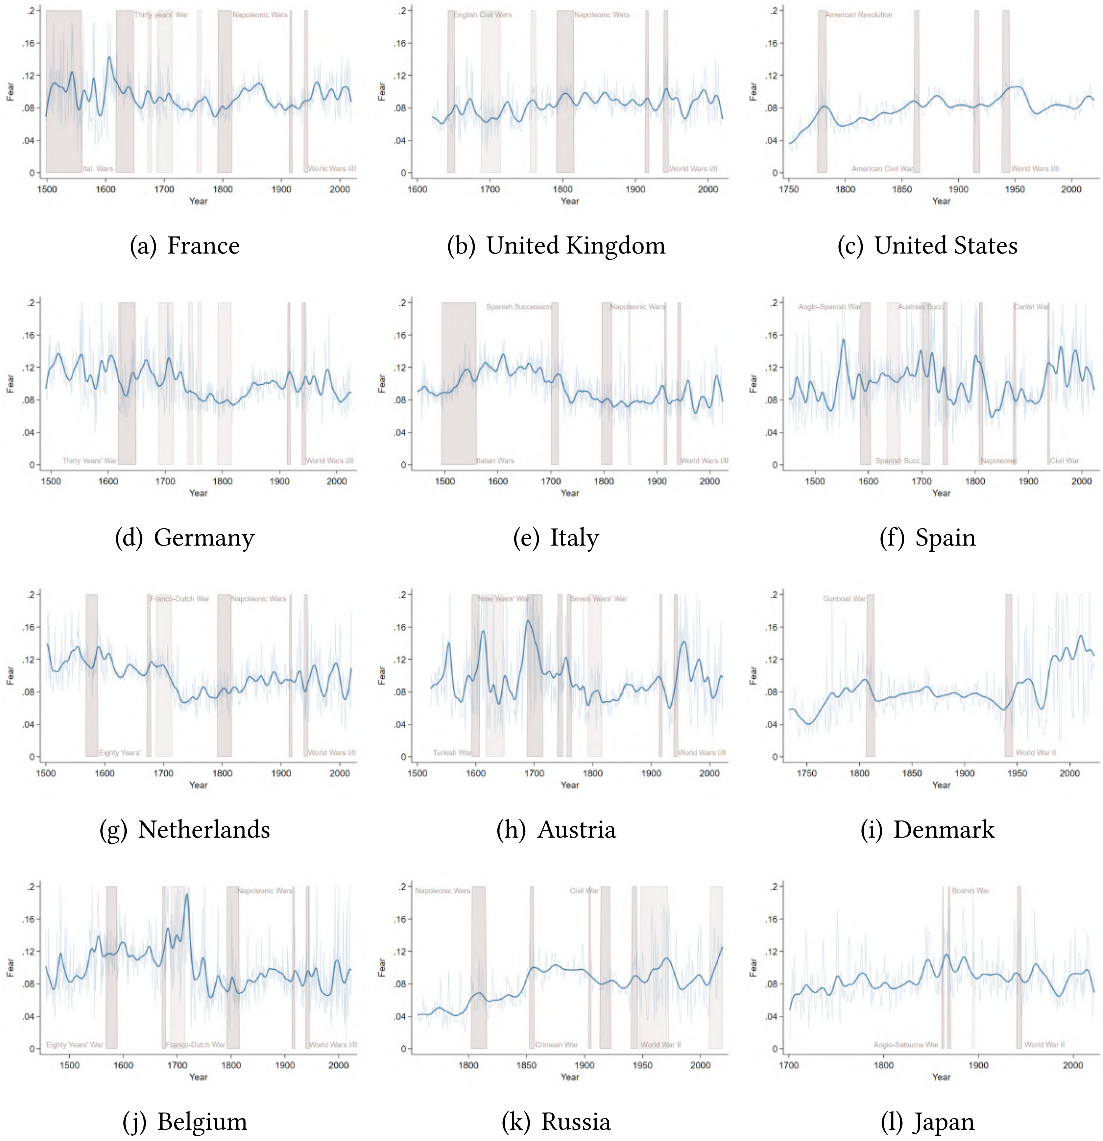

(a) France (b) United Kingdom (c) United States

(d) Germany (e) Italy (f) Spain

(g) Netherlands (h) Austria (i) Denmark

(j) Belgium (k) Russia (l) Japan

Notes: These figures illustrate the dynamics of the predicted probability of evoking *fear*, $p^{f}_{\ell,t}$, in France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. For illustration purposes, we represent the raw time series as a thin line and a smoothed time series—using a Hodrick-Prescott filter with a smoothing parameter of 100—as a thicker line. The respective starting dates are adjusted to exclude the first percentile of artworks in terms of production year. Each observation within a year is weighted to account for duplicates and the overall production volume of each artist. We also include shaded areas indicating periods of major wars for affected countries: e.g., the Wars of the Roses (1455–1487), the Italian Wars (1494–1559), the Eighty Years' War (1568–1648), the Anglo-Spanish War (1585–1604), the Thirty Years' War (1618–1648), the English Civil Wars (1642–1651), the Franco-Dutch War (1672–1678), the Nine Years' War (1688–1697), the Great Northern War (1700–1721), the War of the Spanish Succession (1701–1714), the War of the Austrian Succession (1740–1748), the Seven Years' War (1756–1763), the French Revolutionary Wars (1792–1802), the Napoleonic Wars (1803–1815), the American Civil War (1861–1865), World War I (1914–1918), the Spanish Civil War (1936–1939), and World War II (1939–1945).

In Appendix D.2, we present the long-term evolution of all negative emotions (*fear*, *sadness*, and *anger* in Appendix Figure D5), positive emotions (*excitement*, *amusement*, and *contentment* in Appendix Figure D6), and the remaining emotions (*awe*, *disgust*, and *other* in Appendix Figure D7). We also aggregate all emotion probabilities into a *positivity* index, $l_{\ell,t}$, and analyze its behavior across countries. While there are clear differences in levels (e.g., Great Britain, the United States, and Denmark are persistently more positive before 1900, whereas Italy, Spain, and Germany are persistently more negative) and in volatility (e.g., France, Spain, Belgium, and Austria exhibit greater volatility), a notable feature of the data is the marked negativity observed during the Little Ice Age and the twentieth century. This contrasts with a prolonged "positive" period spanning the eighteenth and nineteenth centuries. Interestingly, these broader fluctuations in *positivity* do not appear to be strongly driven by the rising living standards or the expansion of mass leisure at the turn of the twentieth century in many of the producing countries.

How do conveyed emotions compare to standard economic measures? Visual inspection of the time series of depicted emotions across major economies offers intuitive insights into their economic and social significance. First, Appendix Figure D5 shows peaks in *sadness* during periods of hardship for vulnerable populations, such as the Industrial Revolution or the Little Ice Age. In contrast, emotions like *excitement*, *amusement*, and even *disgust* tend to be more prevalent during times of rising living standards (Appendix Figures D6 and D7). Second, Figure 8 illustrates that depictions of *fear* correspond with periods of war and broader uncertainty, including civil conflicts (e.g., the Huguenot diaspora in France), regime transitions (e.g., the fall of communism), and major political shifts (e.g., the Meiji Restoration in Japan). Conversely, Appendix Figure D6 suggests that *contentment* moves inversely with *fear*, becoming more prominent during stable periods. Third, Appendix Figure D9 shows elevated levels of *disagreement* in times of high economic inequality—for example in France, Great Britain or Germany at the turn of the twentieth century.

In Figure 9, we formalize these interpretations by examining how the predicted probability of evoking each emotion $e \in \mathcal{E}$, denoted p_j^e, as well as the *positivity* index I_j and the *disagreement* index d_j, correlate with three key macroeconomic indicators (see Appendix Figure B12 for historical trends): (a) economic development, measured as log GDP per capita (Bolt and Van Zanden, 2025); (b) economic uncertainty, proxied by six-year volatility in GDP per capita (Bloom, 2014); and (c) wealth inequality, measured by the share of total wealth held by the top 1% (Alvaredo et al., 2020).

These relationships are estimated using Specification (1), which leverages within-artist variation while controlling for year, country, harmonized exercise, and artist age fixed effects. A 169% increase in GDP per capita—equivalent to sustained annual growth of 2% over 50 years—is associated with a 0.25 standard deviation decline in *sadness*, along with moderate increases in *amusement* (0.14 sd) and *excitement* (0.12 sd). Overall, the

Figure 9. How do conveyed emotions compare to standard economic measures?

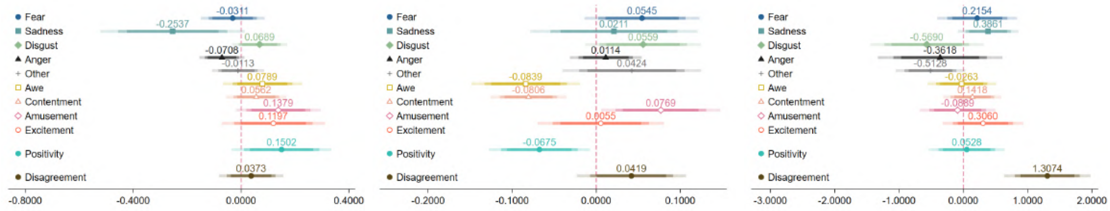

(a) GDP per capita (b) Uncertainty (c) Wealth inequality

Notes: These figures display the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to indicators of economic development and inequality. We consider Equation (1), with location, year, artist, age, and genre fixed effects, and replace the right-hand side variable $\mathbf{s}_{\ell,t}$ with: (a) recently revised estimates of (log) GDP per capita (Maddison Project Database, see Bolt and Van Zanden, 2025); (b) the volatility of GDP per capita calculated over a window of [-2,+4] years ($\sum_{\tau=-2}^{4}|\hat{y}_{t+\tau}|$, controlling for $\sum_{\tau=-2}^{4}\hat{y}_{t+\tau}$); and (c) the concentration of wealth among the top-1% wealthiest individuals, mostly available from the mid-nineteenth century onward (Alvaredo et al., 2020). For (b), the measures $\hat{y}_t$ are residualized using a Hodrick-Prescott filter with a smoothing parameter of 100. The coefficients represent separate regressions with the following left-hand side variables: the predicted probability of expressing emotion e, p_j^e for painting j, where e is *fear* (dark-blue, plain circle), *sadness* (light-blue, plain square), *disgust* (teal, plain diamond), *anger* (black, plain triangle), *other* (gray, cross), *awe* (gold, hollow square), *contentment* (salmon, hollow triangle), *amusement* (pink, hollow diamond), and *excitement* (orange, hollow circle); the *positivity* index, $l_j = \sum_{e\in\mathcal{E}^+} p_j^e - \sum_{e\in\mathcal{E}^-} p_j^e$, in light blue; and a measure of disagreement, defined as the sum of deviations from the mean emotion in a given context, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}}|p_j^e - \bar{p}_{\ell,t}^e|$. All left-hand side variables are standardized, so the reported estimates can be interpreted in terms of standard deviations. The bands represent 10%, 5%, and 1% confidence intervals, and standard errors are clustered at the level of country/half-century.

positivity index rises by 0.15 standard deviations, driven primarily by the reduction in *sadness* and the presence of more positive emotional tones (Figure 9, panel a).

Economic uncertainty, defined as the sum of absolute short-run output deviations ($\sum_{\tau=-2}^{4}|\hat{y}_{t+\tau}|$), triggers a distinct emotional pattern: artists depict higher levels of *fear* and *disgust*, lower levels of *contentment* and *awe*, and increased emotional variation within the same context (Figure 9, panel b).

Wealth inequality primarily influences the dispersion of emotions. For instance, the share of total wealth held by the top 1% in Europe declined by roughly 40 percentage points between the late 19th century and 1980, before rising again by about 10 points in recent decades. A 10-point increase in wealth concentration is associated with a 0.15 standard deviation increase in the *disagreement* index (Figure 9, panel c).

Together, these patterns offer a framework for interpreting emotional fluctuations: changes in *sadness*, *amusement*, or *excitement* broadly reflect material living standards and well-being; shifts in *contentment* or *fear* capture transitions between stable and uncertain times; and variation in *disagreement* signals rising or falling inequality, or heterogeneous emotional responses to external shocks.[15]

[15] In Appendix D.2, we benchmark our main emotion-based indicators—p_j^e, I_j, and d_j—against independent measures of well-being, including sentiment trends in books from select countries (Hills et al., 2019), and historical life expectancy at birth (Zijdeman and Ribeira da Silva, 2015). These benchmarks support the interpretation of *sadness* as a particularly informative proxy for overall living standards.

4.3 The perception of large socio-economic changes

We now turn to a systematic analysis of variation in our emotion indices over time and across countries. The analysis proceeds in five steps: (i) we examine the impact of climatic variability and basic subsistence costs, emphasizing the vulnerability of earlier societies near subsistence levels, as illustrated by the effects of the Little Ice Age (Waldinger, 2022) and the introduction of the potato (Nunn and Qian, 2011); (ii) we investigate the connection between conveyed emotions and trade dynamics; (iii) we analyze the influence of technology adoption, focusing on the steamship (Pascali, 2017) and the radio (including its potential negative effects; see Adena et al., 2015); (iv) we discuss the rise of integrated knowledge production, the decline of religiosity, and the emergence of the scientific revolution; and (v) we examine the relationship between conveyed emotions and (geo-)political events, highlighting the effects of power shifts, political consolidation, and autocratic episodes.

Across all exercises in this section, we estimate specification (1) using the following dependent variables: the predicted probabilities of depicting each of the nine emotions, $\left\{p_j^e\right\}$; the *positivity* index, l_j; and the measure of *disagreement*, d_j. This demanding specification, which isolates variation within the output of the same artist, shows that the emotions conveyed through paintings are closely linked to contemporaneous societal and historical conditions.

Climatic variability and basic subsistence costs Many studies have examined the long-run impact of major historical shocks on European economies prior to the Industrial Revolution. For example, Waldinger (2022) investigates the effects of climatic variability during the Little Ice Age on agricultural productivity and mortality, while Oster (2004) explores scapegoating and witch trials in response to the same climatic shock. Some paintings directly reflect these climatic fluctuations—most notably the increased prevalence of winter landscapes, such as *The Hunters in the Snow* by Pieter Bruegel the Elder. However, the broader societal effects of these fluctuations (see panel a of Figure 10 for variation across a subset of countries) may have been more subtly and systematically embedded in artistic expressions.

To investigate this, Figure 10 (b) presents regression estimates relating the predicted emotions conveyed in a painting, $\{p^e\}_{e\in\mathcal{E}}$, to the annual temperature (in degrees Celsius) during its year of production, available from 1500 onward for all European economies. A one-degree Celsius increase is associated with a 0.12 standard deviation decrease in the predicted probability of evoking *fear* (consistent with reduced uncertainty) and a 0.17 standard deviation decrease in *sadness* (consistent with a lower incidence of poverty).

Figure 10. Climatic variability and the emotions conveyed through paintings.

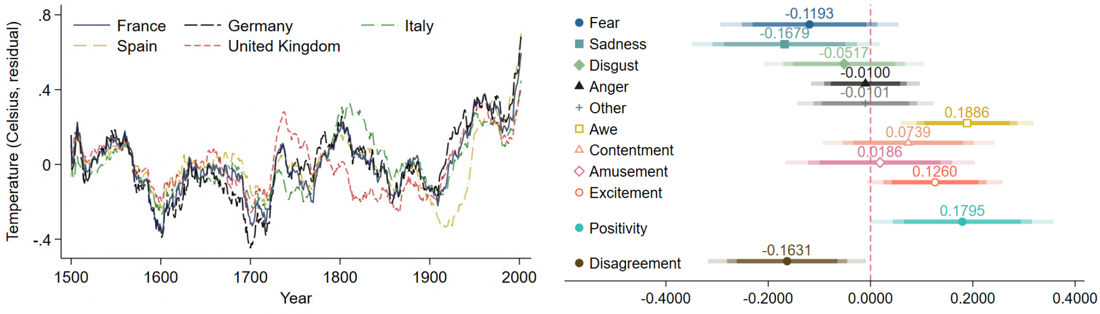

(a) Temperature (evol.) (b) Temperature (effect)

Notes: Panel (a) illustrates variation in climate, defined as the (residualized) 30-year moving average of predicted annual temperature in Celsius degrees in France, Germany, Italy, Spain, and the United Kingdom between 1500 and 2000 (Luterbacher et al., 2004). Panel (b) displays the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to climatic variability across all countries and periods. We consider Equation (1), with location, year, artist, and genre fixed effects, and replace the right-hand side variable $\mathbf{s}_{\ell,t}$ with the 30-year moving average of predicted annual temperature in Celsius degrees (Luterbacher et al., 2004). The coefficients represent separate regressions with the following left-hand side variables: the predicted probability of expressing emotion e, p^e_j for painting j, where e is *fear* (dark-blue, plain circle), *sadness* (light-blue, plain square), *disgust* (teal, plain diamond), *anger* (black, plain triangle), *other* (gray, cross), *awe* (gold, hollow square), *contentment* (salmon, hollow triangle), *amusement* (pink, hollow diamond), and *excitement* (orange, hollow circle); the *positivity* index, $l_j = \sum_{e\in\mathcal{E}^+} p^e_j - \sum_{e\in\mathcal{E}^-} p^e_j$, in light blue; and a measure of disagreement, defined as the sum of deviations from the mean emotion in a given context, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} |p^e_j - \bar{p}^e_{\ell,t}|$. All left-hand side variables are standardized, so the reported estimates can be interpreted in terms of standard deviations. The bands represent 10%, 5%, and 1% confidence intervals, and standard errors are clustered at the level of a country/decade.

Conversely, the likelihood of conveying positive emotions rises—for example, by 0.19 standard deviations for *awe* and 0.13 for *excitement*.

Climatic variation may have had such strong effects because much of the population at the time lived just above subsistence level. To provide more direct evidence of subsistence shocks, we incorporate data on commodity prices across markets and over time (Appendix Figure D11) and examine the relationship between bread prices and conveyed emotions. A 50% increase in the price of bread raises the probability of evoking *anger* and *sadness* (often linked to depictions of poverty) by 0.023-0.033 standard deviations, at the expense of *awe* and *contentment*. Both climatic variability and basic subsistence costs also have a pronounced effect on our measure of emotional *disagreement*.

Trade and the first wave of globalization One key driver of economic development during our period of interest was the expansion of commerce and the opening of new trade routes (Acemoglu et al., 2005b; Pascali, 2017). This shift spurred the rise of a merchant class, increased demand for luxury goods such as silk and spices, and encouraged the adoption of new crops (Nunn and Qian, 2011). However, reliable measures of trade prior to the nineteenth century are limited. To address this, we draw on two sources: (i) a measure of trade openness from the mid-nineteenth century onward (imports as a fraction of output, from Jordà et al., 2017), and (ii) the prices of luxury goods in commodity markets before 1850, shown in Appendix Figure D11 (b) (Allen and Unger, 2019).

Figure 11. Trade and the first wave of globalization.

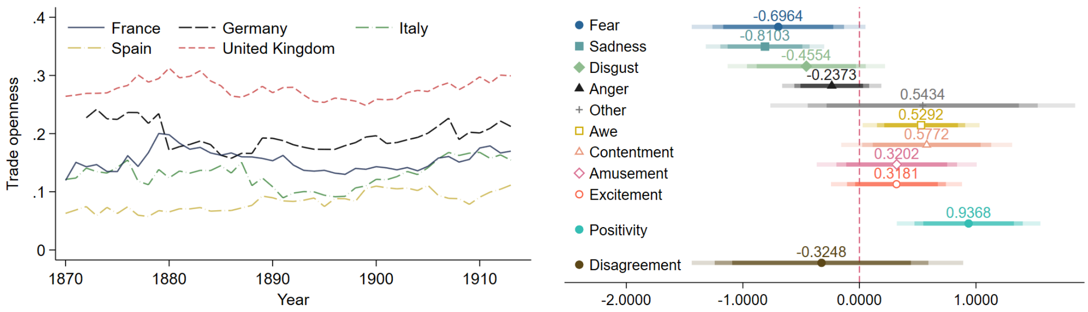

(a) Trade openness (evol.) (b) Trade openness (effect)

Notes: Panel (a) illustrates the variation in trade openness—defined as the ratio of imports to output—in France, Germany, Italy, Spain, and the United Kingdom (Jordà-Schularick-Taylor Macrohistory Database, see Jordà et al., 2017). Panel (b) displays the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to this indicator of trade across all countries and periods. We consider Equation (1), with location, year, artist, and genre fixed effects, and replace the right-hand side variable $\mathbf{s}_{\ell,t}$ with trade openness. The coefficients represent separate regressions with the following left-hand side variables: the predicted probability of expressing emotion e, p^e_j for painting j, where e is *fear* (dark-blue, plain circle), *sadness* (light-blue, plain square), *disgust* (teal, plain diamond), *anger* (black, plain triangle), *other* (gray, cross), *awe* (gold, hollow square), *contentment* (salmon, hollow triangle), *amusement* (pink, hollow diamond), and *excitement* (orange, hollow circle); the *positivity* index, $l_j = \sum_{e\in\mathcal{E}^+} p^e_j - \sum_{e\in\mathcal{E}^-} p^e_j$, in light blue; and a measure of disagreement, defined as the sum of deviations from the mean emotion in a given context, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} |p^e_j - \bar{p}^e_{\ell,t}|$. All left-hand side variables are standardized, so the reported estimates can be interpreted in terms of standard deviations. The bands represent 10%, 5%, and 1% confidence intervals, and standard errors are clustered at the level of a country/decade.

Figure 11 (b) shows that the first wave of globalization—along with access to cheaper food and raw materials—coincided with a markedly more positive emotional outlook in artwork across countries and time periods. A 20 percentage point increase in trade openness (as observed in England during the mid-nineteenth century) is associated with a 0.14–0.16 standard deviation decline in the probability of *fear* or *sadness*. More broadly, negative emotions give way to positive ones, and the overall *positivity* index very significantly rises by 0.19 standard deviations.[16]

Technology adoption The first wave of globalization was partly driven by technological advances in maritime transportation (Pascali, 2017). How societies perceive technological progress is a fascinating question (see, e.g., Mokyr et al., 2015, on the "cultural anxiety" often associated with major innovations). Studying this question is challenging, not only due to limited data availability before 1800, but also because major technological shifts typically unfold gradually over time, whereas our empirical strategy relies on "shorter-horizon" variation within an artist's production cycle.

[16] Appendix Figure D11 (b) shows that higher prices for luxury goods—potentially reflecting increased demand from elites *before* the first wave of globalization—are associated with a less positive outlook: more frequent depictions of *excitement*, but also of *anger*.

Figure 12. Technology adoption (steamship, radio).

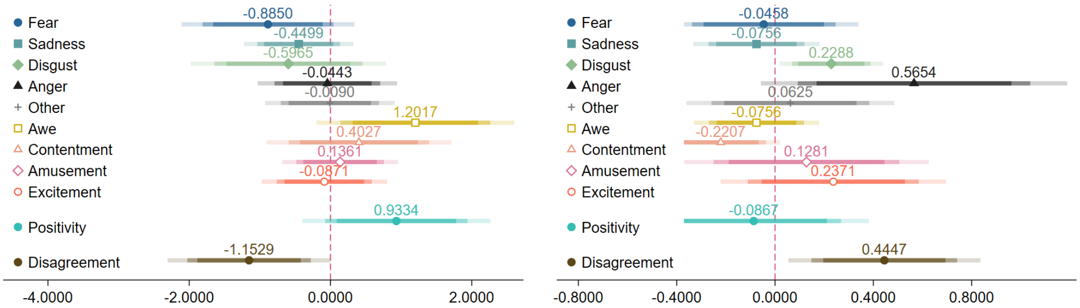

(a) Steamship (b) Radio

Notes: These figures display the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to indicators of technology adoption. We consider Equation (1), with location, year, artist, and genre fixed effects, and we replace the right-hand side variable $s_{\ell,t}$ with: (a) the annual ratio of tons of cargo transported using steamships over the population between 1788 and 1880; and (b) the ratio of radios per inhabitant between 1923 and 1950. Both variables are extracted from the Cross-country Historical Adoption of Technology (Comin and Hobijn, 2004). The coefficients represent separate regressions with the following left-hand side variables: the predicted probability of expressing emotion e, p_j^e for painting j, where e is *fear* (dark-blue, plain circle), *sadness* (light-blue, plain square), *disgust* (teal, plain diamond), *anger* (black, plain triangle), *other* (gray, cross), *awe* (gold, hollow square), *contentment* (salmon, hollow triangle), *amusement* (pink, hollow diamond), and *excitement* (orange, hollow circle); the *positivity* index, $l_j = \sum_{e\in\mathcal{E}^+} p_j^e - \sum_{e\in\mathcal{E}^-} p_j^e$, in light blue; and a sum of deviations from the mean emotion in a given context, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} |p_j^e - \bar{p}_{\ell,t}^e|$. All left-hand side variables are standardized, so the reported estimates can be interpreted in terms of standard deviations. The bands represent 10%, 5%, and 1% confidence intervals, and standard errors are clustered at the level of a country/decade.

In Figure 12, we examine the rapid and uneven diffusion of two innovations from the nineteenth and early twentieth centuries: the steamship, which transformed trade and mobility, and the radio, which influenced entertainment, information, and, in some cases, propaganda (Adena et al., 2015; Castiello, 2025). The figure shows that reactions to new technology can vary considerably. Steamship adoption is associated with positive emotional responses: *awe* increases by 0.036 standard deviations following an increase in steamship cargo equivalent to the rise observed in the United States during the nineteenth century—approximately 0.03 tons per capita per year. By contrast, the spread of radio corresponds to increased depictions of extreme emotions (such as *anger*, *disgust*, and *excitement*) at the expense of *contentment*. The 15 percentage point increase in radio ownership among the French or German populations between 1900 and 1950 would be associated with a 0.084 standard deviation increase in depicted *anger*. While *disagreement* declines following steamship adoption, it rises sharply in the case of radio.

Universities, the decline of religiosity, and the scientific revolution Whereas data on technological adoption in the distant past are limited, recent efforts have helped to document the substantial evolution of academic knowledge production from 1200 onward (De La Croix et al., 2024, see Appendix B.4)—a broad expansion with important implications for economic development (Cantoni and Yuchtman, 2014; Squicciarini and Voigtländer, 2015). We draw on this dataset of scholars and their host institutions to characterize knowledge production within a given context, focusing on a country's global academic influence, its openness to foreign scholars, the role of religiosity in its academic output (Cantoni et al., 2018; Blasutto and De La Croix, 2023; Cummins and Noble, 2025), and the emergence of scientific research.

Figure 13. Universities and the scientific revolution.

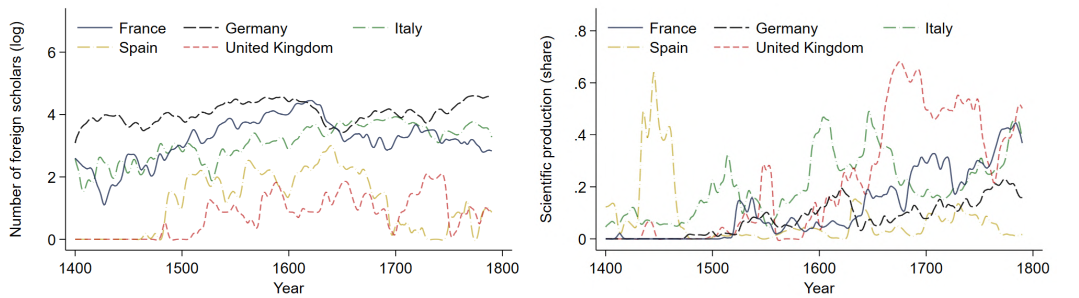

(a) Foreign scholars (evol.) (b) Sciences (evol.)

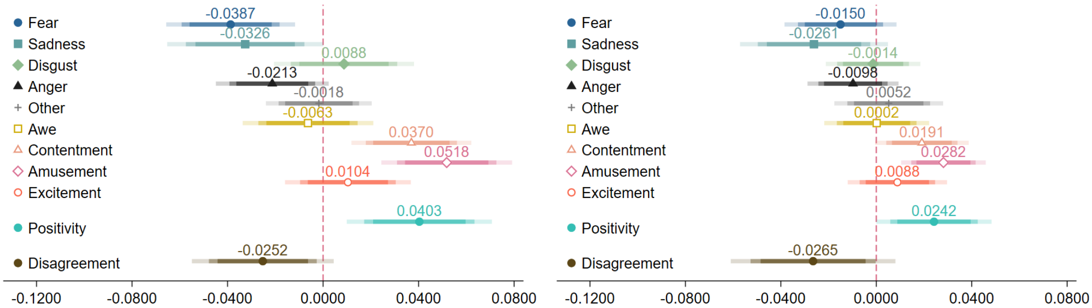

(c) Foreign scholars (effect) (d) Sciences (effect)

Notes: The top panels illustrate the variation in the intensity and nature of knowledge production in France, Germany, Italy, Spain, and the United Kingdom between 1400 and 1800 (De La Croix et al., 2024). Panel (a) reports a measure of academic openness, defined by the (log) number of foreign scholars; and panel (b) reports the weighted share of knowledge production in sciences. The bottom panels display the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to these indicators of knowledge production. We consider Equation (1), with location, year, artist, and genre fixed effects, and we replace the right-hand side variable $\mathbf{s}_{\ell,t}$ with: (c) the (log) number of foreign scholars; and (d) the weighted share of knowledge production in sciences. The coefficients represent separate regressions with the following left-hand side variables: the predicted probability of expressing emotion e, p^e_j for painting j, where e is *fear* (dark-blue, plain circle), *sadness* (light-blue, plain square), *disgust* (teal, plain diamond), *anger* (black, plain triangle), *other* (gray, cross), *awe* (gold, hollow square), *contentment* (salmon, hollow triangle), *amusement* (pink, hollow diamond), and *excitement* (orange, hollow circle); the *positivity* index, $l_j = \sum_{e\in\mathcal{E}^+} p^e_j - \sum_{e\in\mathcal{E}^-} p^e_j$, in light blue; and a sum of deviations from the mean emotion in a given context, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} |p^e_j - \bar{p}^e_{\ell,t}|$. All left-hand side variables are standardized, so the reported estimates can be interpreted in terms of standard deviations. The bands represent 10%, 5%, and 1% confidence intervals, and standard errors are clustered at the level of country/decade.

Figure 13 (panels a and b)—academic openness to foreign scholars and scientific research—and Appendix Figure D12 (panels a and b)—research quality and religiosity—illustrate the substantial variation in the extent and nature of academic production within the borders of five contemporary countries (France, Germany, Italy, Spain, and the United Kingdom) between 1400 and 1800. For example, Great Britain experiences a significant rise in scientific production during the seventeenth century (much larger than in France or Germany), which persists through 1800 and predates the Industrial Revolution. By contrast, Spain and Italy undergo episodes of enlightenment—such as the Scientific Revolution involving Galileo Galilei—but these are followed by abrupt counter-reformations.

Figure 13 (panel c) and Appendix Figure D12 (panel c) show that the overall quality of academic production tends to positively influence the emotional tone of contemporary artistic output across all countries and periods. However, the effect of "openness"—measured by the number of foreign scholars—is particularly pronounced. The increase observed in France between 1500 and 1600 (equivalent to a rise of 2 in the "foreign scholars" variable) is associated with a decline in depictions of *fear* (0.077 standard deviations), *sadness* (0.065), and *anger* (0.042), and a sharp increase in *amusement* (0.10) and *contentment* (0.074). The relative dominance of religiosity over scientific research in academic production yields a mirrored pattern (Figure 13, panel d, and Appendix Figure D12, panel d): a higher prevalence of religiosity is associated with more frequent depictions of sadness and less amusement, whereas the opposite pattern is observed in contexts with a stronger emphasis on scientific research.

Building on our claim that emotions offer insight into how societies perceived their historical moment, these results help illuminate the transformation brought about by the Enlightenment in Europe. Prior to this period, societies were hierarchically structured, with rulers often seen as governing by divine right. Religious doctrine reinforced these arrangements, and the widespread emphasis on the "fear of God" conveyed more than theological devotion—it helped legitimize authority and maintain the social order, often at the expense of intellectual inquiry and innovation. The Enlightenment marked a decisive break from this worldview. Thinkers rejected inherited hierarchies and advocated for a society grounded in reason, science, and individual rights. This vision elevated values such as freedom, rationality, and collective progress, fostering an environment conducive to discovery and experimentation. Scientific advancement became not just a practical pursuit but a cultural symbol of autonomy and forward-looking thought.

Our emotion-based metrics capture this transformation in ways that conventional indicators—such as the volume of scientific output—do not. Periods dominated by religiosity are emotionally characterized by heightened expressions of sentiments consistent with a constraining social structure (*fear* and *sadness*). In contrast, scientific production is associated with greater *amusement* and *contentment*, reflecting a climate of optimism, intellectual curiosity, and growing self-determination. This emotional shift resonates with Mokyr's concept of Industrial Enlightenment (Mokyr, 2005): a transformation not only in knowledge and institutions, but in the underlying disposition toward progress. This example illustrates the potential of our approach to uncover the affective contours of historical change as they were lived and expressed.

Regime characteristics, regime transitions, and wars Finally, we estimate Specification (1) using several short- to medium-term sources of fluctuation in the (geo)political environment, denoted $\mathbf{s}_{\ell,t}$ in Equation (1). These include: a binary indicator for participation in a major war; a standardized measure of "institutionalized autocracy" from *Polity IV* (a composite score capturing executive recruitment, constraints on the executive, and the regulation and competitiveness of alternative viewpoints); a standardized measure of current regime durability (measured as the number of years since the most recent regime change); and an indicator for the nature of regime transitions in years marked by discontinuities in regime characteristics.

First, turbulence during our period of interest often stems from inter-state conflicts, and wars exhibit ambiguous emotional effects (panel a of Figure 14). While depictions of *fear* and *anger* become more common—and the quiet emotion, *contentment*, declines—*excitement* also increases, leading to heightened emotional disagreement among artworks produced within the same context.

Second, regime characteristics (panel b), regime durability (panel c), and the nature of regime transitions (panel d) all significantly influence the emotions conveyed in artworks. A one standard deviation increase in the autocracy score is associated with modest increases in depictions of *fear*, *sadness*, and *disagreement*, each by approximately 0.01 standard deviations. Regime durability has a much stronger—and perhaps surprising—effect (given the literature comparing "stationary bandits" to "roving bandits," see Olson, 1993; Clague et al., 1996; Acemoglu and Robinson, 2006): a one standard deviation increase is associated with a decline in positive emotions (e.g., *amusement* decreases by 0.020 standard deviations), a rise in negative emotions (e.g., *sadness* increases by 0.024), and greater emotional disagreement across artworks. In years marked by regime transitions, positive democratic shifts are associated with more frequent portrayals of *awe* and *excitement*, while negative political shifts tend to elicit *sadness*.[17]

5 Conclusion

This paper shows how historical paintings can provide insights into societal moods and structural changes over time. By analyzing artworks, we complement the limited existing data on historical living standards at a high frequency across European societies

[17] In Appendix D.3, we complement this evidence with (i) a specific focus on labor and emancipation reforms as well as transitions to democratic regimes, and (ii) an event study approach that further examines the dynamics of conveyed emotions.

Figure 14. Regime characteristics, transitions, wars, and the emotions conveyed through paintings.

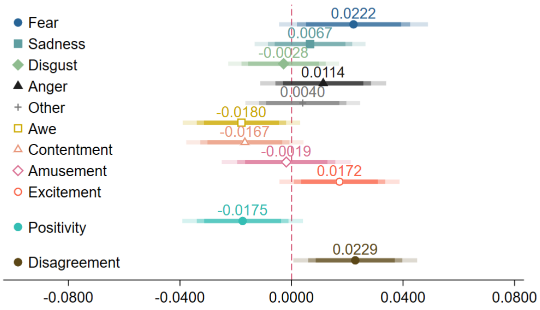

(a) Wars

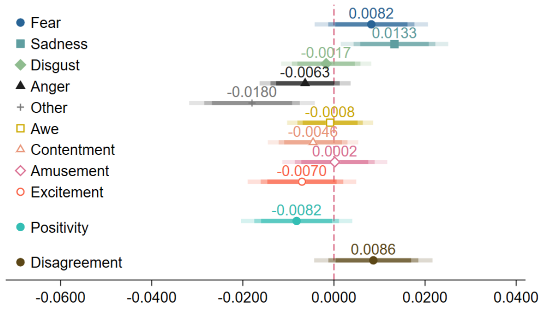

(b) Autocratic regime

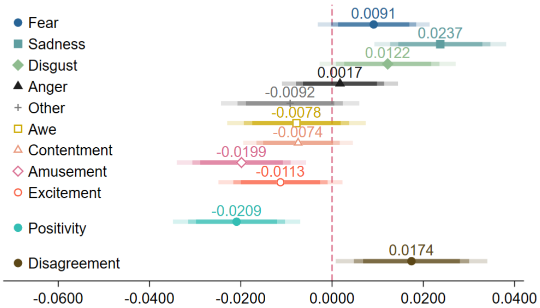

(c) Regime durability

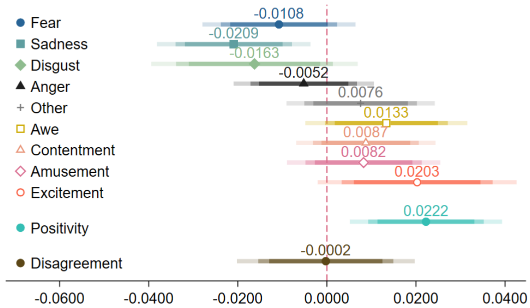

(d) Regime transitions

Notes: These figures display the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e \in \mathcal{E}}$, to indicators of political turbulence and conflicts. We consider Equation (1), with location, year, artist, and genre fixed effects, and we replace the right-hand side variable $\mathbf{s}_{\ell,t}$ with: (a) a dummy for being involved in a major war; (b) a standardized measure of "institutionalized autocracy" from Polity IV (a score combining the competitiveness/openness of executive recruitment, the constraints on the executive, and the regulation/competitiveness of alternative views); (c) a standardized measure of current regime durability (the standardized number of years since the most recent regime change); and (d) the nature of regime transition when there is a discontinuity in regime characteristics (a standardized score centered at 0). The coefficients represent separate regressions with the following left-hand side variables: the predicted probability of expressing emotion e, p^e_j for painting j, where e is *fear* (dark-blue, plain circle), *sadness* (light-blue, plain square), *disgust* (teal, plain diamond), *anger* (black, plain triangle), *other* (gray, cross), *awe* (gold, hollow square), *contentment* (salmon, hollow triangle), *amusement* (pink, hollow diamond), and *excitement* (orange, hollow circle); the *positivity* index, $l_j = \sum_{e \in \mathcal{E}^+} p^e_j - \sum_{e \in \mathcal{E}^-} p^e_j$, in light blue; and a sum of deviations from the mean emotion in a given context, $\frac{1}{|\mathcal{E}|} \sum_{e \in \mathcal{E}} |p^e_j - \bar{p}^e_{\ell,t}|$. All left-hand side variables are standardized, so the reported estimates can be interpreted in terms of standard deviations. The bands represent 10%, 5%, and 1% confidence intervals, and standard errors are clustered at the level of a country/decade.

before 1800. We employ a supervised learning approach to uncover the emotions conveyed in paintings. The underlying idea is to reverse the creative process, recognizing that painters are influenced by their contemporary environments and express these influences in their work. To validate our approach, we compare our extracted emotional readings with painters' own descriptions of their work. Reassuringly, our results closely align with their stated intentions, supporting the effectiveness of our method.

Our results demonstrate that the geography of emotions across locations provides valuable insights into how societies respond to structural economic changes, development, inequality, and geopolitical shocks. We establish the predictive power of expressed emotions by correlating their various dimensions with measures of economic development, political turmoil, and major societal shocks—such as wars, climatic shocks, trade expansion, and technological advances with uncertain long-term effects—while controlling for artist identity and the style of the artwork. Our predicted emotions show pronounced fluctuations around well-documented economic transformations in Europe from 1400 onward. However, they also reveal previously overlooked variations induced by less prominent historical events.

This research stands to benefit from ongoing and future efforts in art preservation, as more museums continue to digitize their collections. Extensions of the present work could enable broader coverage of earlier time periods, non-European regions, and less prominent or lesser-known artists. In particular: (i) refining the spatial allocation of paintings to exact locations would allow for the identification of more granular spatial patterns within countries; and (ii) artists may have faced vastly different living standards within the same location depending on their socioeconomic status, gender, or religion and ethnicity—dimensions that would be important to better understand the geography and dynamics of inequality.

References

Abowd, John M, Francis Kramarz, and David N Margolis, "High wage workers and high wage firms," *Econometrica*, 1999, *67* (2), 251–333.

Acemoglu, Daron and James A Robinson, *Economic origins of dictatorship and democracy*, Cambridge university press, 2006.

_ , Simon Johnson, and James A Robinson, "Institutions as a fundamental cause of long-run growth," *Handbook of Economic Growth*, 2005, *1*, 385–472.

_ , _ , and James Robinson, "The rise of Europe: Atlantic trade, institutional change, and economic growth," *American Economic Review*, 2005, *95* (3), 546–579.

Achlioptas, Panos, Maks Ovsjanikov, Kilichbek Haydarov, Mohamed Elhoseiny, and Leonidas Guibas, "ArtEmis: Affective language for visual art," *CoRR*, 2021, *abs/2101.07396.*

Adena, Maja, Ruben Enikolopov, Maria Petrova, Veronica Santarosa, and Ekaterina Zhuravskaya, "Radio and the Rise of the Nazis in Prewar Germany," *The Quarterly Journal of Economics*, 2015, *130* (4), 1885–1939.

Albrecht, Milton C, "The relationship of literature and society," *American Journal of Sociology*, 1954, *59* (5), 425–436.

Alexander, Victoria D, *Sociology of the arts: Exploring fine and popular forms*, John Wiley & Sons, 2020.

Algan, Yann, Eva Davoine, Thomas Renault, and Stefanie Stantcheva, "Emotions and policy views," Technical Report, Harvard University Working Paper 2025.

Allen, Robert C and Richard W Unger, "The Allen-Unger Global Commodity Prices Database: Social and Economic History," *Research Data Journal for the Humanities and Social Sciences*, 2019, *4* (1), 81–90.

Alvaredo, Facundo, Anthony B Atkinson, Thomas Blanchet, Lucas Chancel, Luis Bauluz, Matthew Fisher-Post, Ignacio Flores, Bertrand Garbinti, Jonathan Goupille-Lebret, Clara Martínez-Toledano et al., "Distributional national accounts guidelines, methods and concepts used in the world inequality database," Technical Report 2020.

Ashenfelter, Orley and Kathryn Graddy, "Art auctions," *Handbook of the Economics of Art and Culture*, 2006, *1*, 909–945.

Bairoch, Paul, *Cities and economic development: from the dawn of history to the present*, University of Chicago Press, 1988.

Bengio, Yoshua, Aaron Courville, and Pascal Vincent, "Representation learning: A review and new perspectives," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 2013, *35* (8), 1798–1828.

Benjamin, Daniel J., Kristen Cooper, Ori Heffetz, and Miles Kimball, "From Happiness Data to Economic Conclusions," *Annual Review of Economics*, August 2024, *16* (1), 359–391.

Blasutto, Fabio and David De La Croix, "Catholic censorship and the demise of knowledge production in early modern Italy," *The Economic Journal*, 2023, *133* (656), 2899–2924.

Bloom, Nicholas, "Fluctuations in uncertainty," *Journal of Economic Perspectives*, 2014, *28* (2), 153–76.

Bolt, Jutta and Jan Luiten Van Zanden, "The Maddison Project: collaborative research on historical national accounts," *The Economic History Review*, 2014, *67* (3), 627–651.

__ **and** __ , "Maddison-style estimates of the evolution of the world economy: A new 2023 update," *Journal of Economic Surveys*, 2025, *39* (2), 631–671.

__ **, Robert Inklaar, Herman De Jong, and Jan Luiten Van Zanden**, "Rebasing Maddison: new income comparisons and the shape of long-run economic development," 2018.

Bonhomme, Stéphane, Kerstin Holzheu, Thibaut Lamadon, Elena Manresa, Magne Mogstad, and Bradley Setzler, "How Much Should we Trust Estimates of Firm Effects and Worker Sorting?," Technical Report 2020.

Borowiecki, Karol Jan, "How Are You, My Dearest Mozart? Well-Being and Creativity of Three Famous Composers Based on Their Letters," *The Review of Economics and Statistics*, 2017, *99* (4), 591–605.

__ , "Good Reverberations? Teacher Influence in Music Composition since 1450," *Journal of Political Economy*, 2022, *130* (4), 991–1090.

Bouscasse, Paul, Emi Nakamura, and Jón Steinsson, "When did growth begin? New estimates of productivity growth in England from 1250 to 1870," *The Quarterly Journal of Economics*, 2024, p. qjae046.

Brown, Steven, *The Unification of the Arts*, Oxford University Press, 11 2021.

Cantoni, Davide and Noam Yuchtman, "Medieval universities, legal institutions, and the commercial revolution," *The Quarterly Journal of Economics*, 2014, *129* (2), 823–887.

__ **, Jeremiah Dittmar, and Noam Yuchtman**, "Religious competition and reallocation: The political economy of secularization in the protestant reformation," *The Quarterly Journal of Economics*, 2018, *133* (4), 2037–2096.

Castiello, Maxence, "Spread the Word: Mass Media, Language and Propaganda in Fascist Italy," 2025. Working paper.

Chen, Xi and William D Nordhaus, "Using luminosity data as a proxy for economic statistics," *Proceedings of the National Academy of Sciences*, 2011, *108* (21), 8589–8594.

Clague, Christopher, Philip Keefer, Stephen Knack, and Mancur Olson, "Property and contract rights in autocracies and democracies," *Journal of Economic Growth*, 1996, *1*, 243–276.

Clark, Gregory, "The condition of the working class in England, 1209–2004," *Journal of Political Economy*, 2005, *113* (6), 1307–1340.

—, "The long march of history: Farm wages, population, and economic growth, England 1209–1869 1," *The Economic History Review*, 2007, *60* (1), 97–135.

Comin, Diego and Bart Hobijn, "Cross-country technology adoption: making the theories face the facts," *Journal of Monetary Economics*, 2004, *51* (1), 39–83.

Compiani, Giovanni, Ilya Morozov, and Stephan Seiler, "Demand estimation with text and image data," *arXiv preprint arXiv:2503.20711*, 2025.

Croix, David De La, Frédéric Docquier, Alice Fabre, and Robert Stelter, "The academic market and the rise of universities in medieval and early modern Europe (1000–1800)," *Journal of the European Economic Association*, 2024, *22* (4), 1541–1589.

Cummins, Neil J. and Aurelius Noble, "Real and Nominal Belief in God: Evidence from Wills, England: 1300-1850," Technical Report 2025.

Damasio, Antonio R, *Descartes' error*, Random House, 2006.

Dell, Melissa, "Deep learning for economists," *Journal of Economic Literature*, 2025, *63* (1), 5–58.

DellaVigna, Stefano, "Psychology and Economics: Evidence from the Field," *Journal of Economic Literature*, June 2009, *47* (2), 315–72.

Deng, Jia, Wei Dong, Richard Socher, Li-Jia Li, Kai Li, and Li Fei-Fei, "Imagenet: A large-scale hierarchical image database," *IEEE conference on computer vision and pattern recognition*, 2009, pp. 248–255.

Dissanayake, Ellen, *Homo Aestheticus: Where Art Comes from and why*, Free Press, 1992.

Donaldson, Dave and Adam Storeygard, "The view from above: Applications of satellite data in economics," *Journal of Economic Perspectives*, 2016, *30* (4), 171–98.

Dutton, Denis, *The art instinct: Beauty, pleasure, and human evolution*, Bloomsbury Press, 2009.

Ekman, Paul, "Basic emotions," in T. Dalgleish and M. J. Power, eds., *Handbook of cognition and emotion*, John Wiley & Sons Ltd., 1999, pp. 45–60.

Elster, Jon, "Emotions and Economic Theory," *Journal of Economic Literature*, 1998, *36* (1), 47–74.

Etro, Federico, "The economics of Renaissance art," *The Journal of Economic History*, 2018, *78* (2), 500–538.

— **and Elena Stepanova**, "Art return rates from old master paintings to contemporary art," *Journal of Economic Behavior & Organization*, 2021, *181*, 94–116.

Frijda, Nico H., *The laws of emotion*, Mahwah, NJ: Lawrence Erlbaum, 2007.

Fujita, Masahisa, Paul R Krugman, and Anthony Venables, *The spatial economy: Cities, regions, and international trade*, MIT Press, 2001.

Gal, Yarin and Zoubin Ghahramani, "Dropout as a bayesian approximation: Representing model uncertainty in deep learning," in "international conference on machine learning" PMLR 2016, pp. 1050–1059.

Galenson, David W and Bruce A Weinberg, "Age and the quality of work: The case of modern American painters," *Journal of Political Economy*, 2000, *108* (4), 761–777.

— **and** — , "Creating modern art: The changing careers of painters in France from impressionism to cubism," *American Economic Review*, 2001, *91* (4), 1063–1071.

Galor, Oded, "From stagnation to growth: unified growth theory," *Handbook of Economic Growth*, 2005, *1*, 171–293.

— , *Unified growth theory*, Princeton University Press, 2011.

Gebru, Timnit, Jonathan Krause, Yilun Wang, Duyun Chen, Jia Deng, Erez Lieberman Aiden, and Li Fei-Fei, "Using deep learning and Google Street View to estimate the demographic makeup of neighborhoods across the United States," *Proceedings of the National Academy of Sciences*, 2017, *114* (50), 13108–13113.

Gell, Alfred, *Art and Agency: An Anthropological Theory* Art and Agency: An Anthropological Theory, Clarendon Press, 1998.

Ginsburgh, Victor and Philippe Jeanfils, "Long-term comovements in international markets for paintings," *European Economic Review*, 1995, *39* (3-4), 538–548.

Giorcelli, Michela and Petra Moser, "Copyrights and creativity: Evidence from Italian opera in the napoleonic age," *Journal of Political Economy*, 2020, *128* (11), 4163–4210.

Graddy, Kathryn, "Taste endures! The rankings of Roger de Piles (1709) and three centuries of art prices," *The Journal of Economic History*, 2013, *73* (3), 766–791.

Han, Sukjin and Kyungho Lee, "Copyright and Competition: Estimating Supply and Demand with Unstructured Data," *arXiv preprint arXiv:2501.16120*, 2025.

Hellmanzik, Christiane, "Location matters: Estimating cluster premiums for prominent modern artists," *European Economic Review*, 2010, *54* (2), 199–218.

Henderson, J Vernon, Adam Storeygard, and David N Weil, "Measuring economic growth from outer space," *American Economic Review*, 2012, *102* (2), 994–1028.

Henrich, Joseph, Robert Boyd, Maxime Derex, Michelle A. Kline, Alex Mesoudi, Michael Muthukrishna, Adam T. Powell, Stephen J. Shennan, and Mark G. Thomas, "Understanding cumulative cultural evolution," *Proceedings of the National Academy of Sciences*, 11 2016, *113.*

Hills, Thomas T., Eugenio Proto, Daniel Sgroi, and Chanuki Illushka Seresinhe, "Historical analysis of national subjective wellbeing using millions of digitized books," *Nature Human Behaviour*, 2019, *3* (12), 1271–1275.

Hulten, Charles R. and Leonard I. Nakamura, "Is GDP Becoming Obsolete? The "Beyond GDP" Debate," *NBER Working Paper 30196*, 2022.

Ioffe, Sergey and Christian Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," *CoRR*, 2015.

Jacks, David S, Kevin H O'Rourke, and Jeffrey G Williamson, "Commodity price volatility and world market integration since 1700," *Review of Economics and Statistics*, 2011, *93* (3), 800–813.

Jaskot, Paul B., "Digital Art History as the Social History of Art: Towards the Disciplinary Relevance of Digital Methods," *Visual Resources*, 4 2019, *35*, 21–33.

Jean, Neal, Marshall Burke, Michael Xie, W Matthew Davis, David B Lobell, and Stefano Ermon, "Combining satellite imagery and machine learning to predict poverty," *Science*, 2016, *353* (6301), 790–794.

Jordà, Òscar, Moritz Schularick, and Alan M Taylor, "Macrofinancial history and the new business cycle facts," *NBER macroeconomics annual*, 2017, *31* (1), 213–263.

Jurado, Kyle, Sydney C Ludvigson, and Serena Ng, "Measuring uncertainty," *American Economic Review*, 2015, *105* (3), 1177–1216.

Kingma, Diederik P. and Jimmy Lei Ba, "Adam: A method for stochastic optimization," in "International Conference on Learning Representations" 2015.

Logothetis, Vassilis, Stelios Michalopoulos, and Elias Papaioannou, "Songs of Refugees," Technical Report 2025.

Loshchilov, Ilya and Frank Hutter, "Decoupled Weight Decay Regularization," in "International Conference on Learning Representations" 2017.

Luterbacher, Jurg, Daniel Dietrich, Elena Xoplaki, Martin Grosjean, and Heinz Wanner, "European seasonal and annual temperature variability, trends, and extremes since 1500," *Science*, 2004, *303* (5663), 1499–1503.

Maddison, Angus, *Contours of the world economy 1-2030 AD: Essays in macro-economic history*, OUP Oxford, 2007.

Mao, Hui, Ming Cheung, and James She, "Deepart: Learning joint representations of visual arts," in "Proceedings of the 25th ACM international conference on Multimedia" ACM 2017, pp. 1183–1191.

Marshall, Monty G, Ted Robert Gurr, and Keith Jaggers, "Polity IV project: Political regime characteristics and transitions, 1800–2013," *Center for Systemic Peace*, 2014, *5.*

Meier, Armando N., "Emotions and Risk Attitudes," *American Economic Journal: Applied Economics*, July 2022, *14* (3), 527–58.

Michalopoulos, Stelios, "Ethnographic Records, Folklore, and AI," *National Bureau of Economic Research*, 2025.

_ **and Melanie Meng Xue**, "Folklore," *The Quarterly Journal of Economics*, 2021, *136* (4), 1993–2046.

Mohamed, Youssef, Mohamed Abdelfattah, Shyma Y. Alhuwaider, Feifan Li, Xiangliang Zhang, Kenneth Ward Church, and Mohamed Elhoseiny, "ArtELingo: A Million Emotion Annotations of WikiArt with Emphasis on Diversity over Language and Culture," Technical Report 2022.

Mokyr, Joel, "The intellectual origins of modern economic growth," *The Journal of Economic History*, 2005, *65* (2), 285–351.

_ **, Chris Vickers, and Nicolas L Ziebarth**, "The history of technological anxiety and the future of economic growth: Is this time different?," *Journal of Economic Perspectives*, 2015, *29* (3), 31–50.

Naik, Nikhil, Scott Duke Kominers, Ramesh Raskar, Edward L Glaeser, and César A Hidalgo, "Computer vision uncovers predictors of physical urban change," *Proceedings of the National Academy of Sciences*, 2017, *114* (29), 7571–7576.

North, Douglass C, "Understanding the process of economic change," in "Understanding the Process of Economic Change," Princeton University Press, 2010.

Nunn, Nathan and Nancy Qian, "The potato's contribution to population and urbanization: evidence from a historical experiment," *The Quarterly Journal of Economics*, 2011, *126* (2), 593–650.

Olson, Mancur, "Dictatorship, democracy, and development," *American Political Science Review*, 1993, *87* (3), 567–576.

Oster, Emily, "Witchcraft, weather and economic growth in Renaissance Europe," *Journal of Economic Perspectives*, 2004, *18* (1), 215–228.

Overman, Henry, Diego Puga, Matthew Turner, and Marcy Burchfield, "Sprawl: A Portrait from Space," *The Quarterly Journal of Economics*, 02 2006, *121*, 587–633.

Pascali, Luigi, "The wind of change: Maritime technology, trade, and economic development," *American Economic Review*, 2017, *107* (9), 2821–2854.

Renneboog, Luc and Christophe Spaenjers, "Buying beauty: On prices and returns in the art market," *Management Science*, 2013, *59* (1), 36–53.

Robinson, Jenefer, *Deeper than reason: Emotion and its role in literature, music, and art*, Oxford University Press, 2005.

Safra, Lou, Coralie Chevallier, Julie Grèzes, and Nicolas Baumard, "Tracking historical changes in perceived trustworthiness in Western Europe using machine learning analyses of facial cues in paintings," *Nature Communications*, 2020, *11* (1), 1–7.

Saleh, Babak and Ahmed Elgammal, "Large-scale Classification of Fine-Art Paintings: Learning The Right Metric on The Right Feature," *International Journal for Digital Art History*, 2016.

Sanh, Victor, Lysandre Debut, Julien Chaumond, and Thomas Wolf, "DistilBERT, a distilled version of BERT: smaller, faster, cheaper and lighter," *arXiv preprint arXiv:1910.01108*, 2019.

Schularick, Moritz and Alan M Taylor, "Credit booms gone bust: Monetary policy, leverage cycles, and financial crises, 1870-2008," *American Economic Review*, 2012, *102* (2), 1029–61.

Scitovsky, Tibor, "What's wrong with the arts is what's wrong with society," *The American Economic Review*, 1972, *62* (1/2), 62–69.

Spaenjers, Christophe, William N Goetzmann, and Elena Mamonova, "The economics of aesthetics and record prices for art since 1701," *Explorations in Economic History*, 2015, *57*, 79–94.

Squicciarini, Mara P and Nico Voigtländer, "Human capital and industrialization: Evidence from the age of enlightenment," *The Quarterly Journal of Economics*, 2015, *130* (4), 1825–1883.

Szegedy, Christian, Wei Liu, Yangqing Jia, Pierre Sermanet, Scott Reed, Dragomir Anguelov, Dumitru Erhan, Vincent Vanhoucke, and Andrew Rabinovich, "Going deeper with convolutions," in "Proceedings of the IEEE conference on computer vision and pattern recognition" 2015, pp. 1–9.

Tan, Mingxing and Quoc V. Le, "EfficientNetV2: Smaller Models and Faster Training," *Proceedings of the 38th International Conference on Machine Learning*, 2021, *139*, 10096–10106.

Throsby, David, "The production and consumption of the arts: A view of cultural economics," *Journal of Economic Literature*, 1994, *32* (1), 1–29.

Tolstoy, Leo, *What is Art?*, New York: Brotherhood Publishing, 1898.

Tompson, Jonathan, Ross Goroshin, Arjun Jain, Yann LeCun, and Christopher Bregler, "Efficient object localization using convolutional networks," *Computer Vision and Pattern Recognition*, 2014.

Tooby, John and Leda Cosmides, "The past explains the present," *Ethology and Sociobiology*, 7 1990, *11* (4-5), 375–424.

Ursprung, Heinrich W and Christian Wiermann, "Reputation, price, and death: An empirical analysis of art price formation," *Economic Inquiry*, 2011, *49* (3), 697–715.

Voth, Hans-Joachim and David Yanagizawa-Drott, "Image(s)," July 2024. Working paper.

Wälde, Klaus and Agnes Moors, "Current Emotion Research in Economics," *Emotion Review*, 2017, *9* (3), 271–278.

Waldinger, Maria, "The economic effects of long-term climate change: evidence from the Little Ice Age," *Journal of Political Economy*, 2022, *130* (9), 2275–2314.

Whitaker, Amy and Roman Kräussl, "Fractional equity, blockchain, and the future of creative work," *Management Science*, 2020, *66* (10), 4594–4611.

Zeiler, Matthew D. and Rob Fergus, "Visualizing and understanding convolutional networks," in "Computer Vision – ECCV 2014" 2014, pp. 818–833.

Zhou, Bolei, Aditya Khosla, Agata Lapedriza, Aude Oliva, and Antonio Torralba, "Learning deep features for discriminative localization," in "Proceedings of the IEEE conference on computer vision and pattern recognition" 2016, pp. 2921–2929.

Zijdeman, Richard and Filipa Ribeira da Silva, "Life Expectancy at Birth (Total)," 2015.