ONLINE APPENDIX—not for publication

A Motivation

This section complements our introductory remarks by: (i) displaying artworks referenced in the motivational discussion (see, e.g., p.2–3); and illustrating the dynamics of emotions in French paintings.

Figure A1. Paintings as a testament to their times and contexts—The Garden of Earthly Delights.

Notes: This Figure shows *The Garden of Earthly Delights* (Hieronymus Bosch, around 1500), which depicts novel fruits and imaginary creatures at a time of exploration and discovery of the New World.

Additional illustrations In the introduction, we refer to Hieronymus Bosch's The Garden of Earthly Delights as a prime example of (i) symbolic or indirect messaging in visual art, and (ii) the emotional duality provoked by uncertain times, simultaneously evoking excitement and anxiety. This triptych is reproduced in Appendix Figure A1: the left panel depicts "The Garden of Eden," the center shows "The Garden of Earthly Delights," and the right panel presents "Hell." While the triptych has been extensively studied, its meaning remains open to interpretation. Nonetheless, the imagery—featuring "exotic" fruits, animals, and characters—appears influenced by the era's global explorations (e.g., the visit to Egypt by Cyriac of Ancona) and the discovery of the New World. These themes resonated with contemporaries such as Albrecht Dürer, who shared a fascination for these discoveries.

In the introduction, we highlight the variation in emotional expression across contemporaneous artworks, even when addressing similar themes. For instance, William-Adolphe Bouguereau—born into a Catholic merchant family—portrays an idealized vision of pastoral life in The Shepherdess (panel a of Figure A2) and other works, where

Figure A2. Paintings as a testament to their times and contexts—different perspectives.

(a) The Shepherdess

(b) The Stone Breakers

(c) Napoleon Crossing the Alps

(d) Snow Storm: Hannibal and his Army Crossing the Alps

Notes: Panel (a) shows *The Shepherdess* (1889, William-Adolphe Bouguereau), an idealized vision of pastoral life. Panel (b) displays *The Stone Breakers* (1849, Gustave Courbet), a painting typical of the realism movement. Interestingly, Bouguereau was a very successful painter in his lifetime, whereas Gustave Courbet gained national prominence later in life and worldwide recognition after his death. Panel (c) shows one of the five versions of *Napoleon Crossing the Alps* (1801–1805, Jacques-Louis David), an allegorical portrait of Bonaparte leading his army (France) through the Col du Grand Saint-Bernard—an image of the early challenges of the French Consulate, seen from a French perspective. Panel (d) presents a response by William Turner, showing the struggle of Hannibal in his earlier crossing (as a possible metaphor for Napoleon's later difficulties), from an English perspective.

peasants often resemble goddesses or nymphs. By contrast, Gustave Courbet—raised in an anti-monarchical family in the Doubs—depicts rural life with what was, at the time, considered radical realism. The Stone Breakers (panel b) captures what he described as "the most complete expression of poverty." A further contrast appears in depictions of The Crossing of the Alps—a symbol of Napoleonic ambition. Jacques-Louis David and William Turner both painted this event, but from markedly different perspectives (Figure A2, panels c and d). David's portrayal casts Napoleon as an allegorical figure, im-

periously leading a nation on horseback—a distinctly French perspective from the early Consulate years. In contrast, Turner's version depicts Hannibal's crossing of the Alps as an army engulfed by a snowstorm—an English view painted during the waning years of Napoleon's reign.

The dynamics of *fear* and *disagreement* in French paintings In the main text, we provide a systematic visualization of how emotions in paintings vary across countries and over time (see Section 4.2 and Appendix D.2). In these introductory paragraphs, we briefly illustrate some of these meaningful fluctuations using the case of France (Appendix Figure A3). Panel (a) shows that depictions of *fear* increase by 3 percentage points during periods of historical turbulence, such as the Religious Wars, the end of the Napoleonic Wars, and World War II. Panel (b) highlights that within-context emotional variance appears to track long-run trends in inequality since 1807, suggesting a connection between social inequality and emotional diversity in artistic expressions within a given environment.

Figure A3. The dynamics of *fear* and *disagreement* in French paintings over time.

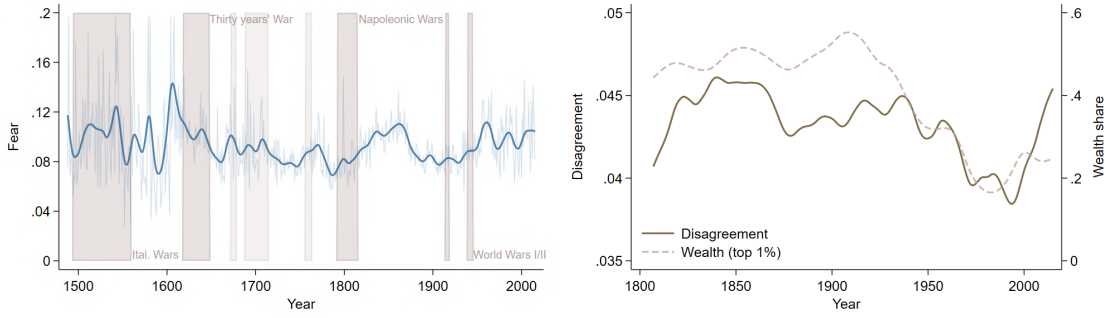

(a) Fear over time

(b) Disagreement over time

Notes: Panel (a) illustrates the evolution of the average emotion probability allocated to *fear*, $p^f_{\ell,t}$, for paintings produced at time t between 1491 and 2015 in France. The starting date is set to exclude the first percentile of French artworks in terms of production time, and each observation within a year is weighted to account for duplicates and the overall production of each artist. 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 coefficient of 100—as a thicker line. Shaded areas highlight the periods of major wars (e.g., 1792–1815 for the Revolutionary and Napoleonic Wars). Panel (b) displays the disagreement in portrayed emotions between 1807 and 2015, alongside the interpolated/smoothed share of wealth owned by the top 1%, extracted from the World Inequality Database (Alvaredo et al., 2020). The context-level *disagreement* index is computed as the average level of disagreement across all paintings produced in a given year t and location ℓ, $d_{\ell,t} = \frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} |r^e_j - \bar{r}^e_{\ell,t}|$, where both r^e_j and $\bar{r}^e_{\ell,t}$ are derived from the earlier measures, p^e_j and $\bar{p}^e_{\ell,t}$, but residualized by genre, movement, age, and artist fixed effects.

B Data description

This section provides complements to Section 2.

B.1 A harmonized collection of paintings

Google Arts and Culture *Google Arts and Culture* is a partnership between Google and numerous cultural institutions (museums, galleries, etc.) aimed at digitizing and publishing art collections in high resolution for online access. Many of the largest art museums collaborate with Google, resulting in the availability of most well-known paintings that are open to the public as online images through the platform. For example, the Musée d'Orsay (Paris) shares hundreds of items, including *The Church in Auvers-sur-Oise* by Vincent Van Gogh and *The Ballet Class* by Edgar Degas. The Museum of Modern Art shares its collection, including *Hope II* by Gustav Klimt and *Turning Road at Montgeroult* by Paul Cézanne. Overall, more than 9,000 artists are represented on *Google Arts and Culture*, with about 360,000 art pieces (approximately two-thirds of which are paintings). The most well-known artists, including Albrecht Dürer, Pierre-Auguste Renoir, Rembrandt, Vincent Van Gogh, Edgar Degas, Henri de Toulouse-Lautrec, and William Turner, typically have between 500 and 2,000 paintings featured. While the art collections primarily consist of European pieces, they also include Chinese calligraphy, Japanese prints and water-based works, "primitive" art, landscapes from the Hudson River School, and American realism.

Artists and paintings are accompanied by descriptions and various tags. Figure B1 provides an example of this information and how it is stored across the different, interlinked pages associated with an artist or a specific artwork. Data collection involves the use of a headless browser, followed by post-processing and cleaning procedures. The main challenge lies in correctly identifying the location of production and assigning a production date. First, we retain the completion year of the artwork or, when a time interval is provided (e.g., 1900–1906), we use the final year reported by *Google Arts and Culture*. Second, we use entity recognition models and large language models to clean the location data and associate a (geo)location. Location information is often noisy—partly because it is frequently unobserved by museum curators, art collectors, or auction houses—unless explicitly tied to a specific place, as is often the case with landscapes. When the exact location is unavailable, we retain the country of production. If the country is also missing, we reconstruct it from artist biographies.

Figure B2 shows the locations of a selection of exactly geolocated paintings across Northern Europe and Japan, where only about one-third of artworks are assigned to a precise location. The figure also illustrates that (i) the geography of artwork production

Figure B1. Artist biographies and artworks in *Google Arts and Culture*.

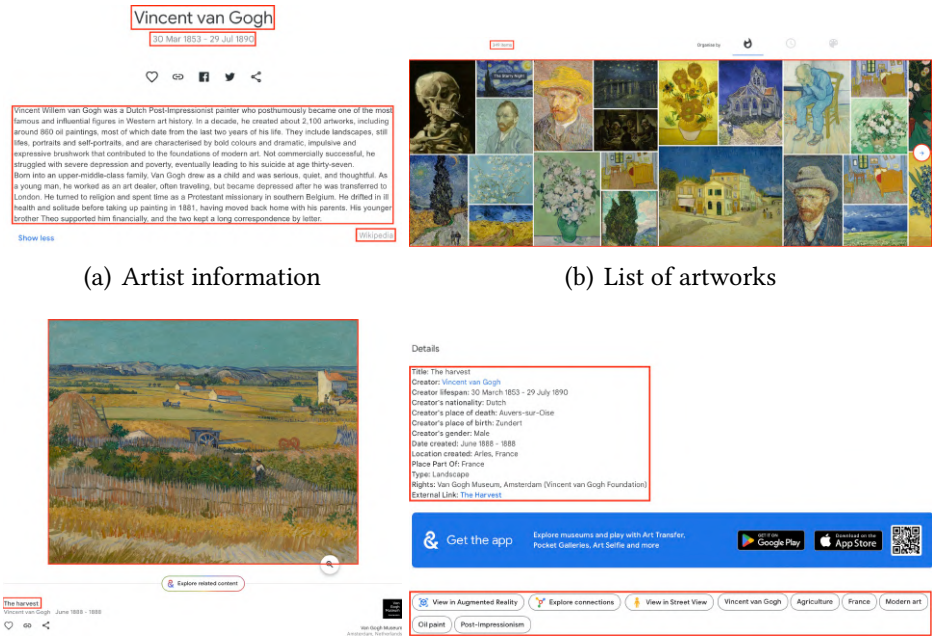

(a) Artist information (b) List of artworks

(c) Artwork image (d) Artwork attributes and tags

Notes: This figure shows an example of the information available in *Google Arts and Culture* for an artist (here, Vincent van Gogh). Panel (a) is the biographical page; panel (b) shows the list of artworks; panel (c) displays the image of *The Harvest*; and panel (d) shows the associated attributes and tags. The artist information includes: name, birth and death dates, description, list of artworks, and Wikidata ID for linking with Wikipedia information. The painting information includes: the image, title, date, location of creation, and medium.

closely aligns with the geography of economic activity, and (ii) the golden ages of painting (e.g., the Dutch Golden Age or the French Belle Époque) are clearly visible in the preserved artworks—with the Netherlands mostly represented by paintings from before 1700, and France predominantly covered by paintings from 1800–1900.

Wiki-Art *WikiArt* is an online repository of approximately 250,000 high-resolution images of paintings (Saleh and Elgammal, 2016). The selection of artworks and associated metadata follows the general principles of Wikipedia, i.e., it is a collaborative effort augmented by checks and moderation to ensure the quality and reliability of the information provided. As a result, the collection primarily focuses on the golden ages of painting, such as the Northern Renaissance, the Belle Époque, and on very well-known artists, often with near-complete coverage of their oeuvre (e.g., Pablo Picasso, Claude Monet, Salvador Dalí, Vincent van Gogh, Henri Matisse, Paul Cézanne, Wassily Kandinsky, Edward Hopper, Leonardo da Vinci, Rembrandt, Gustav Klimt, Egon Schiele, Frida Kahlo, Albrecht Dürer, and Caravaggio). *WikiArt* also contains high-quality information

Figure B2. The geography of (selected) paintings in Northern Europe and in Japan.

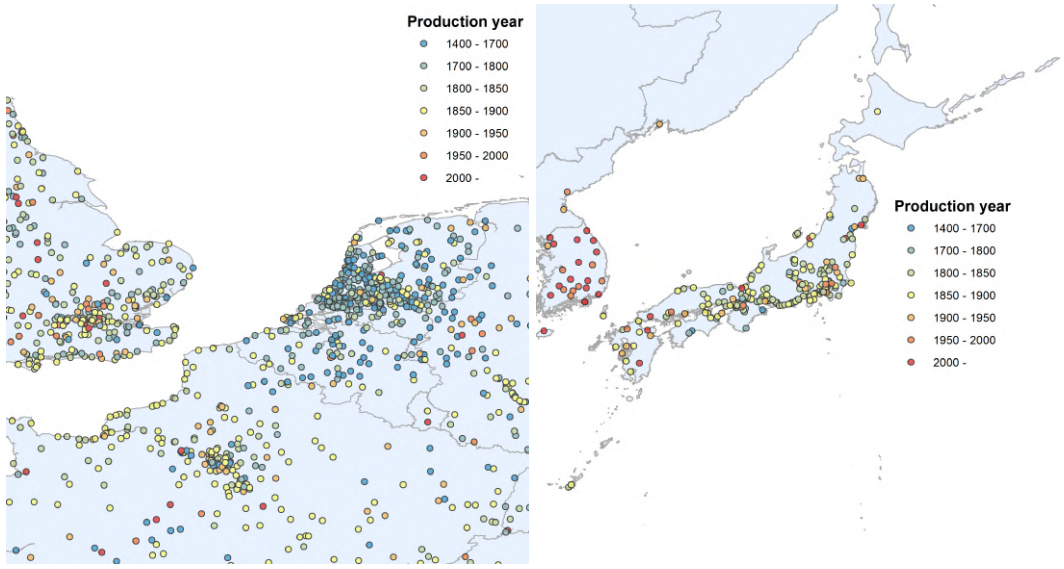

(a) Northern Europe (b) Japan

Notes: This figure shows the geolocation of paintings in a subset of *exactly* geolocated production locations in Northern Europe and Japan. The colors of the dots indicate the period: 1400–1700, 1700–1800, 1800–1850, 1850–1900, 1900–1950, 1950–2000, and 2000 onward.

on artists, most notably their networks of "influencers" and "influencees."[18]

This collection provides well-organized information about both the painting and the artist. It is further enriched by two annotation projects conducted on a subsample of artworks, in which individuals were asked to associate each painting with one of nine emotions. One advantage of *WikiArt* is its highly detailed and structured metadata. For example, artworks are classified into more than 90 main movements and a similarly rich set of exercises (e.g., abstract, genre painting, landscape, cityscape, animal painting, flower painting, marina, sketch and study, still life, nude painting, portrait, self-portrait, mythological painting, allegorical painting, religious painting).

To harmonize these classifications across sources, we combine structured queries to the Wikipedia API, tags from over 500,000 paintings in ART500K (Mao et al., 2017), and Large Language Models (GPT-4o) to match 100 harmonized movements (e.g., impressionism, mannerism, fauvism) and 40 harmonized exercises (e.g., allegorical painting,

[18]For instance, Paul Cézanne was influenced by: *Gustave Courbet, El Greco, Charles-Francois Daubigny, Nicolas Poussin, Pierre-Auguste Renoir, Eugene Delacroix, Jean-Baptiste-Simeon Chardin*; he had influence on: *Pablo Picasso, Amedeo Modigliani, Jackson Pollock, Fernand Leger, Chaim Soutine, Piet Mondrian, Francis Bacon, Man Ray, Vilhelm Lundstrom, Paul Gauguin, Wassily Kandinsky, Roman Selsky, Adalbert Erdeli, Michel Kikoine, Giorgio Morandi, Jozef Pankiewicz, Robert Falk, Harry Phelan Gibb, Marjorie Acker Phillips, Thomas Hart Benton, Beauford Delaney, William Balthazar Rose*; and he was friend and co-worker with *Paul Gauguin*.

landscape, still life). This information allows us to condition our analysis on genre and exercise, which may influence how paintings are interpreted by annotators. The temporal distribution of the main movements and exercises across all data sources is shown in Appendix Figure B3.

Figure B3. Movements and genre/exercise.

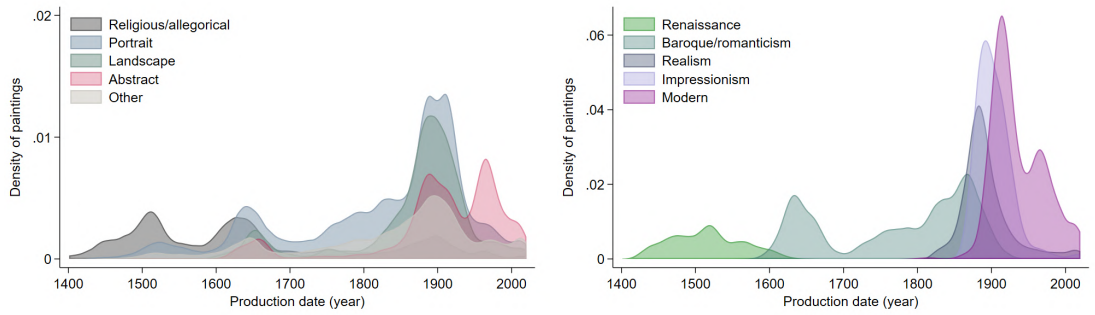

(a) Exercise/genre

(b) Movement

Notes: Panel (a) shows the distribution of production years across different categories of exercise (abstract: abstract, genre painting; landscape: landscape, cityscape, animal painting, flower painting, marina, sketch and study, still life, nude painting; portrait: portrait, self-portrait; religious/allegorical: mythological painting, allegorical painting, religious painting; other). Panel (b) shows the distribution of production years across different categories of movement (renaissance: early renaissance, northern renaissance, mannerism; modern: symbolism, cubism, expressionism, abstract expressionism, art nouveau, pop art; impressionism: impressionism, post-impressionism; baroque: baroque, rococo, romanticism; realism: realism). These distributions are calculated from artworks in *Google Arts and Culture*, *Wiki-Art*, and *Wiki-Data*.

Wiki-Data The WikiData collection is officially titled the "WikiProject Sum of All Paintings," and is dedicated to the systematic identification and linking of "notable paintings." More specifically, "[a] painting is considered notable if one of these options is true: the painting is in a notable collection of a museum, library, or archive (e.g., a painting in the Metropolitan Museum of Art); or the painting is made by a notable artist (e.g., a Van Gogh painting in a private collection)." One consequence is that the collection is primarily composed of artworks held in major national and regional museums.

A live summary of its coverage is available here, showing that the largest contributors include many of the world's most prestigious institutions—for example, the Louvre (10,129 pieces), the Metropolitan Museum of Art (13,007 pieces), the Museo del Prado (6,418 pieces), and the Rijksmuseum—as well as prominent national galleries such as the Finnish National Gallery, the Statens Museum for Kunst in Copenhagen, and the National Museum in Warsaw.

Duplicates The previous collections include more than 650,000 paintings, but overlaps exist due to the shared interest in notable artwork. Identifying duplicates is surprisingly

Figure B4. Detecting duplicates.

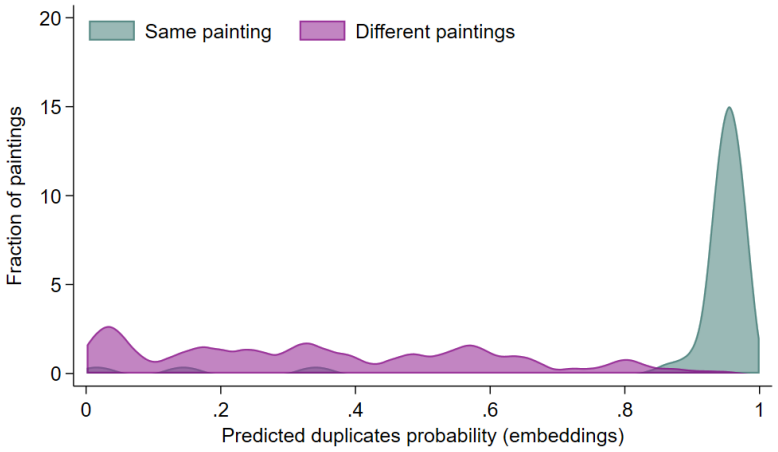

Notes: This figure displays the distributions of the embedding-based probability of being duplicates, $\widetilde{\varphi}_{i,j}$, between paintings in a "training" sample. We distinguish between pairs of paintings that have been identified as duplicates (green) and those that have not (purple).

challenging because titles often differ across sources—sometimes slightly, sometimes significantly—and many lesser-known paintings lack formal titles altogether. Rather than relying on fuzzy matching based on these short or missing textual descriptions, we use the algorithm described in Section 3, which generates a vector of 1,260 embeddings for each painting k, denoted as $z_{e,k}$ in Section 3. Letting $\mathbf{z}_i$ represent the vector of 1,260 embeddings associated with painting i, we filter all paintings from the same artist and period across different collections and calculate the distance to these paintings j as:

$$d_{i,j} = (\mathbf{z}_i - \mathbf{z}_j)'(\mathbf{z}_i - \mathbf{z}_j) = \sum_e \left(z_{e,i} - z_{e,j}\right)^2 .$$

We then identify duplicates and non-duplicates for about 1,000 pairs of paintings from the same artists, i.e., we construct an indicator $d_{i,j} \in \{0, 1\}$ within this sub-sample, and estimate the following probit specification,

$$\varphi_{i,j} = P\left(d_{i,j} = 1 \mid d_{i,j}\right) = \Phi\left(\alpha + \beta d_{i,j}\right).$$

The estimates $(\widetilde{\alpha}, \widetilde{\beta})$ allow us to compute the predicted probability of being a duplicate for any pair of paintings (i, j) across collections, $\widetilde{\varphi}_{i,j}$, which we subsequently use to re-weight our sample. Figure B4 illustrates the predictive power of the embedding-based distance in distinguishing duplicates from non-duplicates within the "training" sample: approximately 90% of true duplicates have a predicted probability above 0.90, compared to only 2% of false duplicates.

B.2 A collection of *annotated* images

A collection of *annotated* images Our training sample consists of two sources of annotations associated with *Wiki-Art* images: ArtEmis (Achlioptas et al., 2021, providing about 450,000 labels for 79,860 paintings); and ArtELingo (Mohamed et al., 2022, providing more than 1,000,000 labels for the same 79,860 paintings). Both datasets include textual explanations justifying the choice of emotion, with the latter dataset leveraging differences across annotators in language and culture (thereby covering many different countries). Figure B5 displays the average emotion scores and the cross-correlations for/between: *amusement*, *anger*, *awe*, *contentment*, *disgust*, *excitement*, *fear*, *sadness*, and *other*. As discussed in Section 2, the most frequent emotions are *contentment* (e.g., landscapes), *awe*, *amusement*, *sadness*, and *fear* (panel a). Panel (b) shows the correlation between the resulting scores: *fear*, *anger*, and *disgust* appear to function as substitutes within a well-defined cluster of negative emotions. At the other end of the spectrum, *contentment* is often treated as a distinct emotion and is rarely mentioned alongside other emotions in a systematic way.

Figure B5. Average emotion scores and cross-correlations.

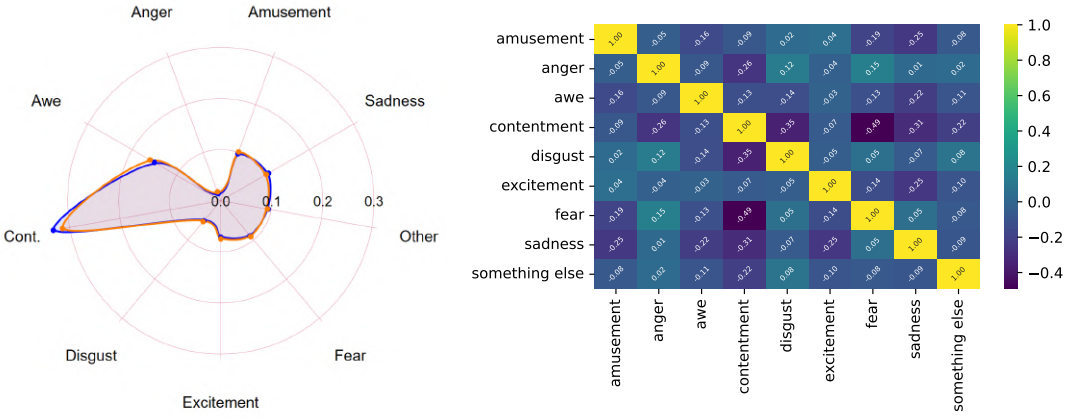

(a) Scores (b) Correlations

Notes: Panel (a) shows the distribution of emotion scores across 79,860 annotated paintings—the annotations in orange and the predictions in blue. The respective scores are (standard deviations in parentheses): *amusement*, 0.101 (0.058); *anger*, 0.018 (0.019); *awe*, 0.168 (0.074); *contentment*, 0.302 (0.123); *disgust*, 0.048 (0.048); *excitement*, 0.071 (0.039); *fear*, 0.095 (0.090); *sadness*, 0.101 (0.086); and *other*, 0.094 (0.044). Panel (b) shows the correlation between each dimension of the emotion score associated with the 79,860 annotated paintings.

A probabilistic vector of emotions In our baseline strategy, we construct a probabilistic vector of emotions, $\mathbf{q}_j$, from the 15-25 different labels associated with each painting j. Specifically, we treat these multiple labels as a probabilistic evaluation from a single individual. Each painting is assigned a vector of scores representing the share of

annotators who selected each of the nine emotions, normalized so that the scores sum to 1 for each painting. In the presence of measurement error and idiosyncratic traits or preferences of annotators, the predicted score could in principle depend on the number of labels provided. However, Appendix Figure B6 (a) shows that the average score does not appear to vary systematically with the number of annotators.

Figure B6. Average and dispersion in annotations, as a function of the number of labelers.

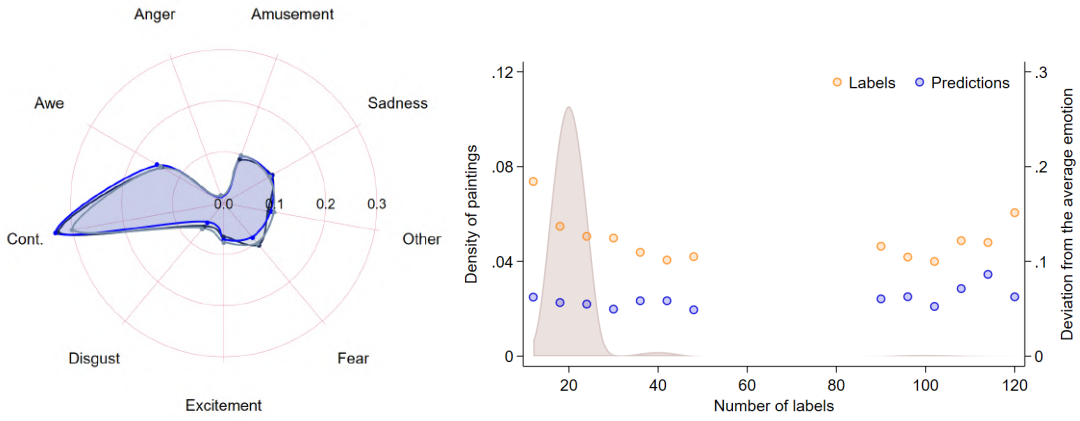

(a) Scores by labels

(b) Variance by labels

Notes: This figure explores the distribution of ratings and their number per artwork. Panel (a) replicates panel (a) of Figure B5 for artworks associated with fewer than 17 labels (dark blue), between 18 and 22 labels (blue), and more than 23 labels (light blue). Panel (b) shows a measure of deviation from the average emotion within the annotated sample, based on labels (orange) and predictions (blue), as a function of the number of provided labels. To construct these deviations from the mean, we collapse the different emotion ratings for a painting into an empirical probability distribution, i.e., the probability that one of the associated annotators reports a given emotion, and compute the squared difference from the overall sample average (represented as dots). Data sources: ArtEmis (Achlioptas et al., 2021) and ArtELingo (Mohamed et al., 2022).

The number of annotations may not affect the average emotion attributed to a painting, but it could influence the variance of scores across artworks. For instance, if each annotator provided a purely random, idiosyncratic assessment, we would expect the scores to converge to a stable value as the number of annotations increases. However, we find no evidence supporting this pattern. As shown in Figure B6, panel (b), artworks with fewer annotations are not significantly more likely to deviate from the average rating in the overall population of paintings.

B.3 Attrition, commercialization, and selection into conservation

Our main data sources, described in Appendix B.1, may be subject to attrition bias, driven not only by selection into these collections, but also by selection into preservation or conservation. For instance, one might expect more valuable artworks to be more likely to survive aggregate shocks and idiosyncratic events (e.g., inheritance). A potential concern for our analysis is that such selection might be systematically correlated with the

emotions conveyed through paintings, such that the average "color" of preserved artwork reflects the preferences of a later audience. To explore this possibility, this section draws on digitized sales catalogs from the Getty Research Institute—described in Appendix B.4—to shed light on the role of attrition.

Mapping artists across data sources While both our collection and the sales catalogs contain information on the artwork (e.g., title and physical characteristics), the data in the sales catalogs are too unstructured and incomplete to allow for reliable matching at the artwork level. For instance, a year of production is rarely available, and artworks are often described informally rather than by a formal title—titles that are often assigned retrospectively by curators or owners. Given these limitations, we instead link the two data sources using artist identity. We clean artist names in both datasets and apply a fuzzy matching procedure. Approximately 45% of artworks featured in the sales catalogs were produced by artists also present in our sample—a share that rises to 51% for transactions with a disclosed price.

This bridge across the two data sources allows us to examine two sources of selection: (i) the selection of artwork into commercialization, and (ii) the selection of artwork into preservation or conservation. For (i), we compare artworks in our dataset based on whether the artist also appears in the sales catalogs. For (ii), we examine differences in transaction values between "preserved" artists (those included in our data) and "forgotten" artists (those not covered).

Selection of artwork into commercialization Figure B7 sheds light on selection into commercialization by examining the extent to which paintings in our dataset can be attributed to artists whose work appears in the sales catalogs (versus "non-selling" artists), and whether commercialized artwork differs in terms of conveyed emotions. Panel (a) shows the share of potentially commercialized paintings by genre or type of exercise. Panel (b) presents the average emotion probabilities for artwork by artists featured in the sales catalogs (in green) versus those attributed to non-selling artists (in red). The only notable differences concern *amusement*/*sadness*: sad paintings are slightly less likely to be commercialized, while amusing paintings are slightly more likely to be. These differences, however, remain marginal.

Selection of artwork into preservation/conservation Another, possibly more important, question concerns selection into preservation or conservation: which artworks, produced at a given time and with a given "value," are preserved to the present day and ultimately selected for inclusion in digital collections? To address this, we proceed in

Figure B7. Selection of artwork into commercialization.

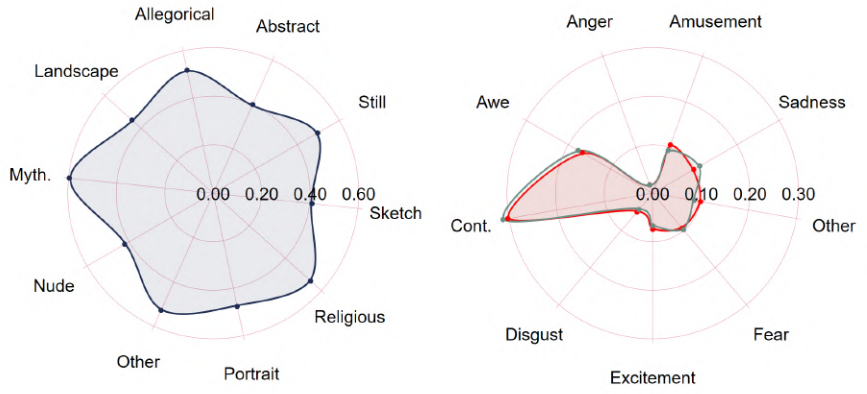

(a) Genre (b) Emotion

Notes: These figures provide statistics on the selection of artwork from our collection of paintings into commercialization, as captured by a sample of sales catalogs digitized by the Getty Research Institute (part of the so-called *Getty Provenance Index*). As the latter only covers a selection of countries at different points in time, the following statistics are generated for artwork produced in a given country and year, for which at least 10% can be attributed to an artist present in the sales catalogs. In simpler terms, we exclude contexts that are not covered at all by the sales catalogs. Panel (a) displays the attrition rate for our artwork (i.e., the share that can be attributed to an artist whose work features in the sales catalog), as a function of its genre or type of exercise—where *Abstract* stands for abstract, genre painting; *Landscape* stands for cityscape, landscape, marina; *Still* stands for animal painting, flower painting, still life; *Sketch* stands for sketch and study; *Nude* stands for nude painting; *Portrait* stands for portrait, self-portrait; *Religious* stands for religious painting; *Myth.* stands for mythological painting; and *Allegorical* stands for allegorical painting. Panel (b) displays the average emotion probabilities for artwork that can be attributed to an artist whose work features in the sales catalog (in green) versus artwork attributed to "non-selling" artists (in red).

Figure B8. Selection of artwork and selling prices.

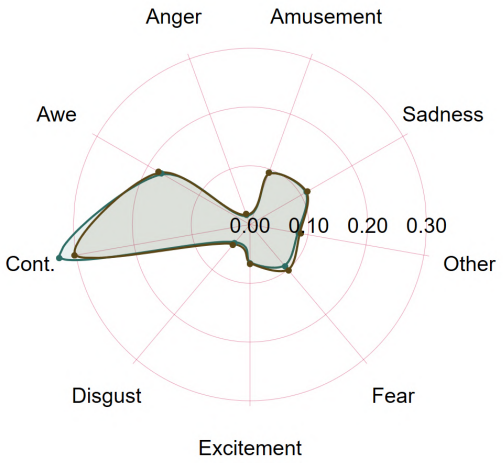

Notes: This figure is based only on artwork that can be associated with an artist present in the sales catalogs and displays the average emotion probabilities for artwork attributed to an above-median artist (in green, with an average selling price above the median of their contemporaries) versus a below-median artist (in red, with an average selling price below the median of their contemporaries).

two steps.

In a first step, we focus on preserved artworks whose creators appear in the sales catalogs and ask: do highly valued artists produce different kinds of works than their peers? Figure B8 displays the average emotion probabilities for artworks attributed to artists with an average selling price above (lighter green) or below (darker green) the median price for their contemporaries. The differences in conveyed emotions between these groups are minimal. Although not entirely comparable—sales may feature the same artwork multiple times and may occur long after the time of production—the number of auctions within a given context is roughly four times higher than the number of preserved paintings in our collection.

Table B1. Selection of artwork into preservation/conservation: estimating a value gap.

Price (log)	(1)	(2)	(3)
In our sample	0.2062	0.2027	0.1060
	(0.0457)	(0.0453)	(0.0276)
Fixed effects: currency	Yes	Yes	Yes
Fixed effects: country	Yes	Yes	Yes
Fixed effects: year	Yes	Yes	Yes
Fixed effects: month	No	Yes	No
Fixed effects: auction	No	No	Yes
Observations	509,424	509,424	509,424

Notes: A unit of observation is a transaction within the sales catalogs digitized by the Getty Research Institute (part of the so-called *Getty Provenance Index*), for which we observe: (i) a price and currency, (ii) the artist and their country of origin, and (iii) a consistent link between pieces by the same artist within the sales catalogs. The explaining variable is a dummy equal to 1 if the artist is present in our main sample of artwork (described in Appendix B.1), and standard errors are clustered at the artist level. In column (1), we add fixed effects at the level of the currency, the country (of the artist), and the year. In column (2), we further include a set of month fixed effects. In column (3), we additionally include a set of auction fixed effects to identify a premium from pieces sold during the same auction.

In a second step, to quantify the selection of artwork into preservation or conservation, we ask: how much more valuable were preserved artists at the time of production? (see Appendix Table B1). We focus on a subsample of transactions from our digitized sales catalogs that include consistent information on price, currency, and artist, and compare prices for artists represented in our main digitized, preserved collection of paintings versus "forgotten" artists. In the most stringent specification (column 3), we identify the preservation premium by comparing pieces sold during the same auction and find it to be around 10%. This moderate premium may suggest that the disproportionate value contemporary societies now place on notable artists was not fully reflected in their market value at the time of production.

B.4 Other data sources

We have assembled a range of data sources capturing historical economic conditions to help validate and interpret fluctuations in emotions, as discussed in Section 4. Below, we provide a more detailed description of these sources.

Figure B9. The production of knowledge over time.

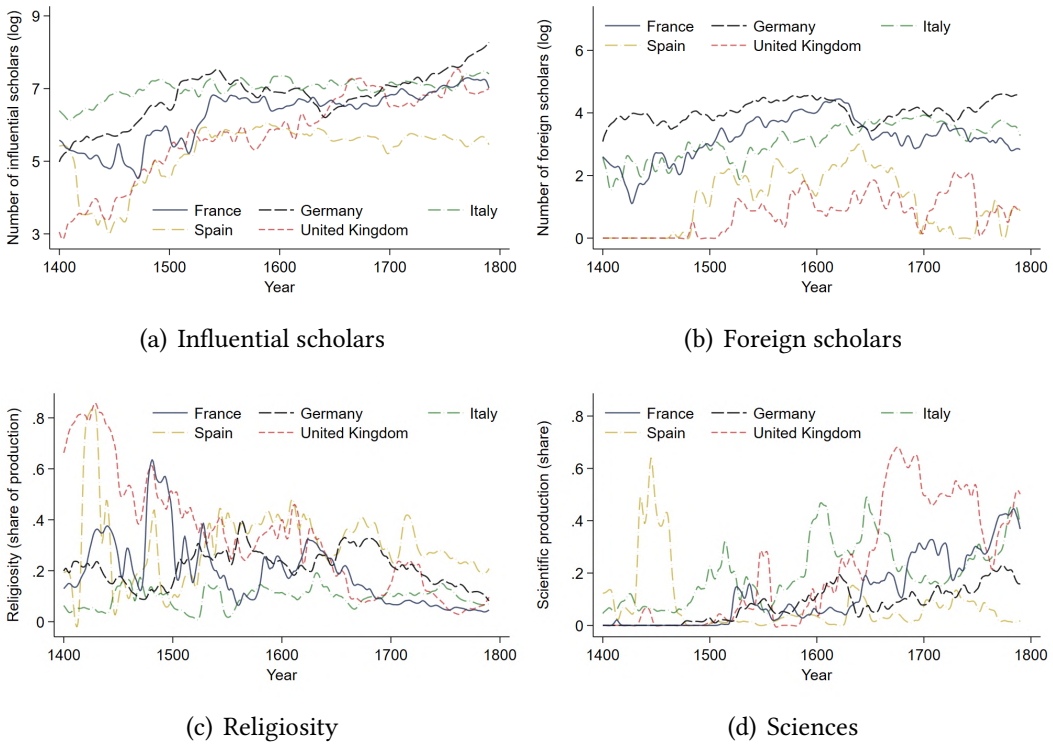

(a) Influential scholars

(b) Foreign scholars

(c) Religiosity

(d) Sciences

Notes: This figure illustrates 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 "quality" or influence, as captured by the (log) number of scholars with more than 10 Wikipedia pages (in different languages). Panel (b) reports a measure of openness, defined by the (log) number of foreign scholars. Panel (c) reports the weighted share of knowledge production in theology. Panel (d) reports the weighted share of knowledge production in the sciences. This variation is used more systematically across countries and periods in Figure 13 of the paper.

Large socioeconomic changes Our analysis aims to capture several key transformations in Medieval Europe, the Renaissance, the later Reformation, the Enlightenment, and the Industrial Revolution. To this end, we first draw on data documenting the adoption of new technologies (Cross-country Historical Adoption of Technology; see Comin and Hobijn, 2004), focusing on two technologies highlighted in Figure 12: (i) steamships, which contributed to the first wave of globalization, and (ii) radio (including its potential use as a tool of propaganda; see Adena et al., 2015; Castiello, 2025).

The Macro History Database provides more conventional measurements of macroeconomic variables from the mid-nineteenth century onward for 18 economies (Jordà-

Schularick-Taylor Macrohistory Database, see Jordà et al., 2017). One of these variables—trade openness (the ratio of imports and exports to GDP), which we analyze in Section 4 and Figure 11—allows us to capture the dynamics of early globalization.

This period is marked by the rise of universities, the democratization of knowledge, a decline in religiosity, and the onset of the scientific "revolution." We capture these long-run secular changes using an extensive database of scholars and their host institutions across Europe (The Rise of Universities in Medieval and Early Modern Europe; see De La Croix et al., 2024).

The scholars data are recorded at the individual level, and we aggregate them to the country-year level in three steps. First, we assign characteristics to each scholar, including measures of influence (such as the number of listed publications, Wikipedia page length, and number of language editions), whether the scholar is foreign (based on country of birth), and their main academic field (theology, law, arts, medicine, or sciences). Second, we construct a scholar-year panel to account for spells across different institutions. Third, we aggregate these data at the country-year level using contemporary borders. We focus on three main outcomes: (i) the number of influential scholars (defined as those with more than ten Wikipedia language editions); (ii) the number of foreign scholars; and (iii) the weighted share of knowledge production across fields. For this third measure, we restrict attention to influential scholars, though results are robust to alternative weighting schemes.

Figure B9 illustrates historical patterns in France, Germany, Italy, Spain, and the United Kingdom, based on the scholars database: (a) knowledge production rises over time in all countries; (b) France, Germany, and Italy heavily rely on foreign scholars—Great Britain less so—and this reliance reflects both secular trends and idiosyncratic movements; (c) religiosity follows markedly different trends across economies, declining sharply in Great Britain from 80% of academic output in 1400 to under 10% by 1700; (d) scientific production follows the opposite pattern, with a notable rise during the mid-seventeenth century in Great Britain (and to a more moderate extent in France)—a pattern not clearly evident in the Holy Roman Empire, which was in its Baroque period. Time variation in scientific production is particularly pronounced in Italy, with major peaks during the early Scientific Revolution (e.g., Leonardo da Vinci, Galileo Galilei), followed by a sharp decline after 1630 with the onset of the Counter-Reformation and church censorship. We use—and in some cases reproduce—this variation in Figure 13.

Shocks to livelihoods The previous section focused on longer-term, secular changes in livelihoods across the period of study. In this section, we turn to more transient fluctuations throughout the period of interest.

Figure B10. The price of commodities over time.

(a) Bread (b) Luxury

Notes: This figure illustrates the variation in the (residualized) price of commodities in the United Kingdom (mostly England) between 1550 and 1850. The procedure involves residualizing (log) prices of standardized measures and aims to account for differences in units and currencies, the selection of markets entering or exiting the database, and the possible aggregation of similar commodities. Panel (a) displays the variation in a "basic" commodity, bread, and panel (b) focuses on a "luxury" commodity based on several different goods (e.g., including silk, spices). This variation is used more systematically across countries and periods in Appendix Figure D11 and Appendix D.3.

First, we capture political turbulence using regime characteristics and transitions documented from 1800 onward by the *Polity IV* project (Marshall et al., 2014). Second, we identify key historical events that affected political entities corresponding to 26 contemporary countries, e.g., major wars.[19] Third, we capture short-lived changes in the price of basic commodities. To do so, we rely on a database of commodity prices at the market level (Allen-Unger Global Commodity Prices Database; see Allen and Unger, 2019). Constructing measures at the country-year level is challenging due to differences in units and currencies, variation in market coverage over time, and the possible aggregation of similar commodities. We consider the (log) price of a standardized measure of commodity c within a market m in year t, and consider,

$$p_{cmt} = \alpha_c + \beta_m + \gamma_t + \varphi_{cmt}.$$

Next, we aggregate the residuals $\{\varphi_{cmt}\}$ within commodity category d, a given country i, and year t, using the prevalence of a market/commodity over time as weights. We

[19]We define a sample of contemporary territories for which we track socio-economic changes and shocks: Argentina, Australia, Belgium, Brazil, Canada, China, Czech Republic, Denmark, France, Germany, India, Italy, Japan, Mexico, New Zealand, Norway, Poland, Portugal, Russia, South Africa, Spain, Sweden, Switzerland, Turkey, the United Kingdom, and the United States. In each of these countries/territories, we identify country-specific shocks (e.g., colonization, international and civil wars, pandemics, revolutions, political upheaval, new trade arrangements, famines, and financial crises), as well as the timing of major socio-economic transformations (e.g., the emancipation of slaves, female emancipation, the arrival of new communication and transportation technologies, industrialization, mass urbanization, and the emergence and regulation of labor unions). For each event and country, we record a unique event identifier, the start year (and, where applicable, the end year), a textual description, and whether and which other countries are involved.

illustrate the variation in those quantities for two "commodities", bread (basic, used in Appendix Figure D11, panel (a) and a luxury basket (e.g., including silk, spices, used in Appendix Figure D11, panel (b), in the United Kingdom (see Figure B10).

Figure B11. Climatic conditions over time.

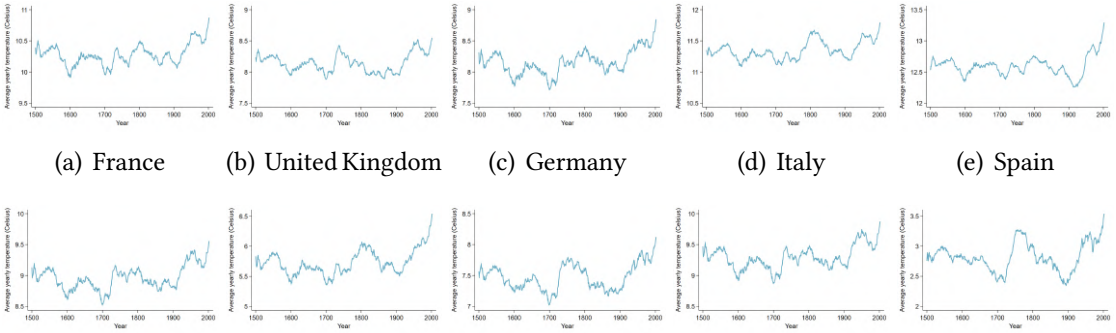

(a) France (b) United Kingdom (c) Germany (d) Italy (e) Spain

(f) Netherlands (g) Austria (h) Denmark (i) Belgium (j) Russia

Notes: This figure illustrates the average temperature between 1500 and 2000 in France, Great Britain, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, and Russia—as extracted from Luterbacher et al. (2004). One can observe secular fluctuations in temperature, sometimes common across European countries, and sometimes more country-specific. For visualization purposes, we report a 30-year moving average. This variation is used more systematically across countries and periods in Figure 10 of the paper.

Finally, Section 4.3 examines the emotional response conveyed through paintings to the secular but non-negligible climate variability affecting Europe from 1500 onward. Our approach relies on the harmonized temperature reconstructions from Luterbacher et al. (2004), which are based on multiple sources, including tree-ring data and ice cores. We report their temperature predictions for the main European "purveyors" of paintings in Figure B11. Two notably cold decades, around 1600 and 1700, are common across most European countries, though more pronounced in Germany, Denmark, and the Netherlands. Approximately 60% of the within-country temporal variation is explained by a "common" component—a share amplified by the anthropogenic nature of climate warming in the late twentieth century.

Indicators of economic development As part of a validation exercise (see Figure 9), we benchmark fluctuations in emotions against standard indicators of economic development. First, the Maddison Project provides estimates of output and population for a selected sample of countries from 1400 onward (Bolt and Van Zanden, 2014; Bolt et al., 2018); our analysis relies on their most recent revisions (Maddison Project Database; see Bolt and Van Zanden, 2025). We use these data to construct key macroeconomic variables, illustrated in Appendix Figure B12: (a) economic development, measured as (log) GDP per capita (Bolt and Van Zanden, 2025); and (b) economic uncertainty, proxied by

Figure B12. The variation in standard economic measures.

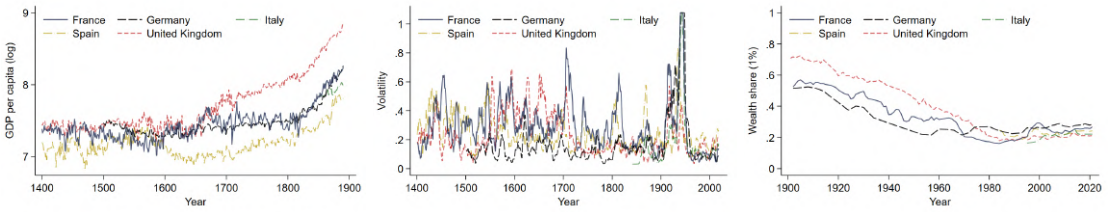

(a) GDP per capita (b) Uncertainty (c) Wealth inequality

Notes: These figures display the variation in France, Germany, Italy, Spain, and the United Kingdom for: (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% of individuals, mostly from the mid-nineteenth century onward (Alvaredo et al., 2020). This variation is used more systematically across countries and periods in Figure 9 of the paper.

six-year implied volatility in GDP per capita (Bloom, 2014).

Second, we incorporate a "valence index" derived from sentiment analysis of published books (Hills et al., 2019). Third, we use historical estimates of life expectancy at birth (Zijdeman and Ribeira da Silva, 2015). Finally, the World Inequality Database allows us to capture economic inequality through the concentration of wealth held by the top 1% or 10% of individuals, with particularly strong coverage for France (Alvaredo et al., 2020, see also Appendix Figure 9 and Appendix Figure B12).

Auction data We retrieve and clean entries from the sales catalogs digitized by the Getty Research Institute (part of the so-called Getty Provenance Index). These catalogs provide detailed records of public auctions, primarily from Belgium, France, Germany, Great Britain, and the Netherlands—each country being covered over a different time period between 1650 and 1945 (e.g., Great Britain in the eighteenth and early nineteenth centuries). After cleaning the entries from approximately 20,000 catalog—including variables such as price, clearly identified painter, and title of the piece—we are left with around 1,200,000 transactions, sometimes involving the same artwork. The dataset allows for the identification of artists (e.g., linking recurring artists such as Brueghel the Elder, Van Dyck, Poussin, Rembrandt, Rubens, Teniers the Younger across their different artworks and over time) as well as recurring sellers and buyers, typically painters, art collectors, aristocrats, or wealthy investors.

As discussed in Appendix B.4, we use the high-quality reporting of artist names to link entries in these catalogs to artists in our primary, harmonized collection of preserved paintings. This linkage enables us to better understand the scope and direction of attrition, commercialization, and selection into conservation.

C An algorithm to predict emotions

In Section 3, we describe our image classification algorithm. In this Appendix, we provide further insights into our model structure and its ability to identify emotions.

C.1 Image pre-processing

Image pre-processing While convolutional neural network architectures offer flexibility in how pixel values interact and contribute to the final prediction, they are more rigid in the required input format. Therefore, all images undergo pre-processing to convert them into a standardized $384 \times 384 \times 3$ array (height, width, color). Each pixel in this array contains Red-Green-Blue (RGB) information, which is stored as an integer between 0 and 255.

Figure C1. Image pre-processing.

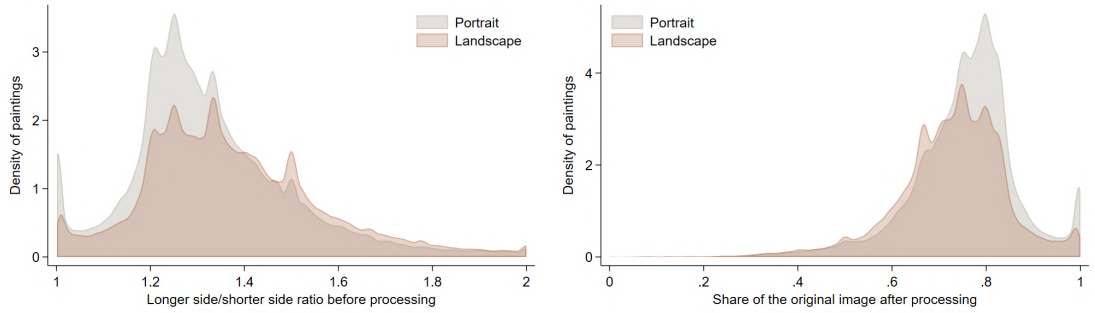

(a) Longer side/shorter side ratio

(b) Coverage

Notes: Panel (a) shows the distribution of the longer side/shorter side ratio across images in our harmonized collection; panel (b) displays the distribution of the average share of these initial images that is retained through our image pre-processing procedure. The areas of the two distributions reflect their contribution to the final sample of artwork. These distributions are calculated from artworks in *Google Arts and Culture*, *Wiki-Art*, and *Wiki-Data*.

This pre-processing step has two main implications: (i) the resulting image is slightly blurred compared to the original painting, and (ii) the final image may not retain all features of the original, as only the largest possible square with the same central point as the original image is preserved. Despite this, Figure 3 illustrates that reducing the original image to a 384×384 collection of pixels does not significantly alter our general perception of the painting. To further explore potential losses due to the exclusion of non-central pixels, we analyze the distribution of the longer side to shorter side ratio for paintings oriented as portraits or landscapes (see Figure C1). The results indicate that about 80% of centered images retain more than 70% of the original pixels. In theory, these preserved sections of the image should encompass the most important areas, typically around the focal point.

Figure C2. Structure of the model.

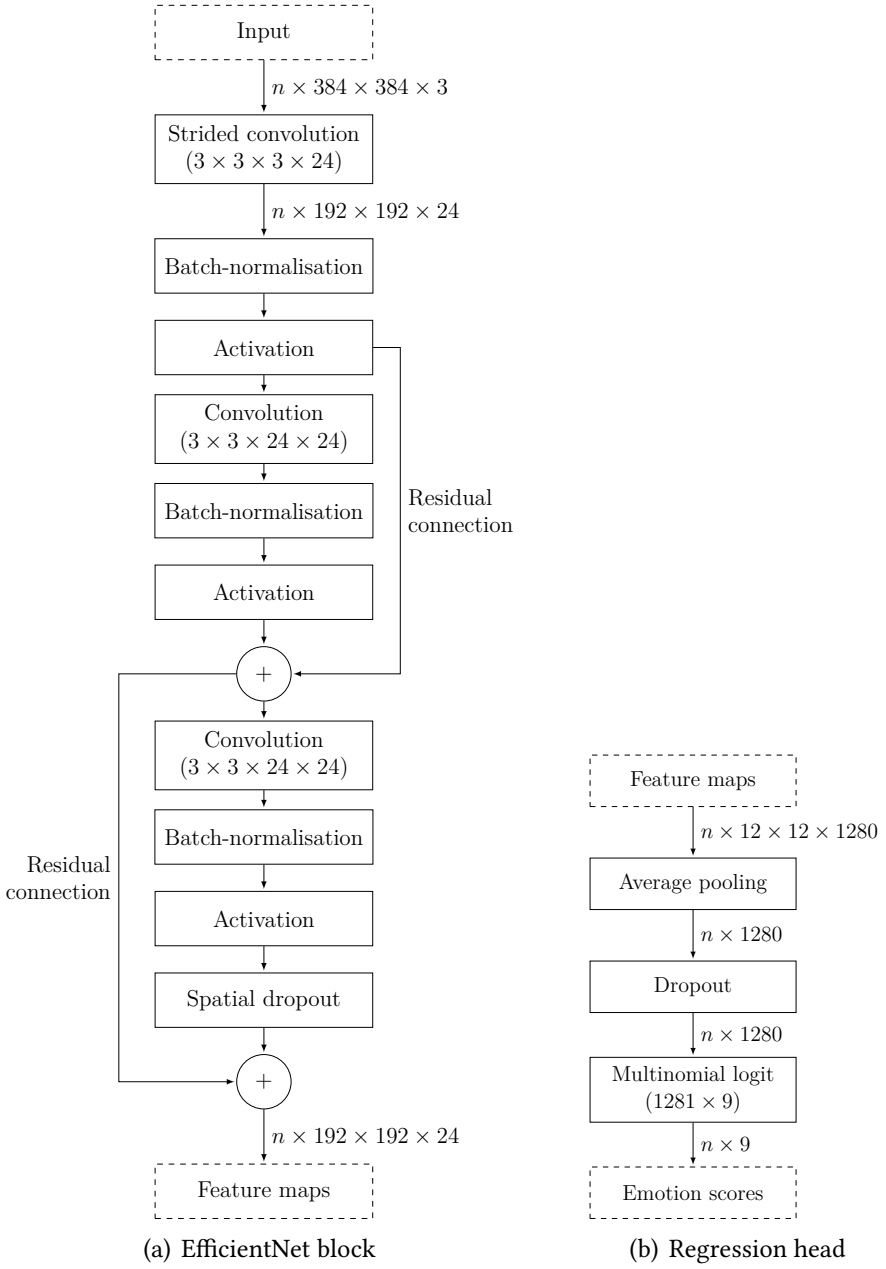

Notes: This figure illustrates the architecture used to generate emotion scores from input images. Panel (a) shows the first EfficientNet block, which processes input images of size $n \times 384 \times 384 \times 3$ and outputs features of dimension $n \times 192 \times 192 \times 24$. The block includes strided and residual convolutions, batch normalization, activation functions, and spatial dropout. Subsequent blocks contain attention mechanisms (i.e., "squeeze and excite" blocks) that enable the model to dynamically focus on specific features depending on global context information (i.e., global pooling and fully connected layers). Panel (b) displays the regression head, which takes as input feature maps of size $n \times 12 \times 12 \times 1280$ produced by the last EfficientNet block, applies global average pooling and dropout, and outputs a probabilistic vector of nine emotion scores through a multinomial logit layer.

C.2 A neural network structure

Structure, estimation, and optimization The prediction model integrates a pre-trained encoder model with a regression head. Intuitively, the encoder extracts a large

number of features from the images, while the regression head uses these features to predict emotion scores. The encoder is based on the convolutional layers of an EfficientNet V2-S model (Tan and Le, 2021), with the original output layer removed and replaced by our regression head. One advantage of using a pre-trained model is that it leverages extensive training—here, on 1,400,000 photos of objects across 1,000 different classes. However, the model was originally designed to predict object classes, not emotions, and was trained on real-life pictures rather than paintings. To address this, we attach a regression head to the encoder, mapping the extracted features to scores associated with each emotion. Letting 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.

Figure C2 provides a detailed description of both the encoder and the regression head. The regression head processes the extracted features through two densely connected hidden layers, each containing 64 units with ReLU activation, followed by batch-normalization (Ioffe and Szegedy, 2015) and dropout layers (Tompson et al., 2014). The final layer applies the previous multinomial logistic transformation, ensuring the emotion scores sum to 1 and allowing them to be interpreted as probabilities.

The model estimation minimizes the Kullback-Leibler divergence between the predicted and observed score distributions, better accounting for the probabilistic nature of our target and limiting the occurrence of false negatives—instances where true high emotion scores are predicted as low. The Kullback-Leibler divergence is defined as:

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

where q^e is the actual probability that one of the 15-25 labelers chooses emotion e as the dominant emotion, and p^e is the predicted probability. In practice, we estimate the model using maximum likelihood with an adaptive learning rate for each parameter, dynamically estimated learning rates, and sub-samples of 64 images at each optimization iteration (as in Kingma and Ba, 2015). The optimization process begins by freezing the encoder parameters, taking the pre-trained feature extraction as given and aligning the regression head to predict emotions. Once the regression head parameters are well initialized, we fine-tune the model by (i) unfreezing the encoder parameters (except for the batch-normalization layers, which are updated using an exponential moving average)

and (ii) training the model with a small learning rate to avoid over-fitting. The rationale is straightforward: the pre-trained encoder may not extract features optimally for our purposes, but its initial parametrization is valuable for initializing the regression head (justifying step i). Optimizing the entire model simultaneously would likely cause large gradient updates from the head, initialized with random values, potentially disrupting the encoder's ability to extract meaningful patterns. However, as the process of conveying emotions through paintings differs from real-life pictures—as evidenced by Figure 3 and Figure 4—the encoder model requires updating (justifying step ii).

Figure C3. Filters, composition and hierarchy.

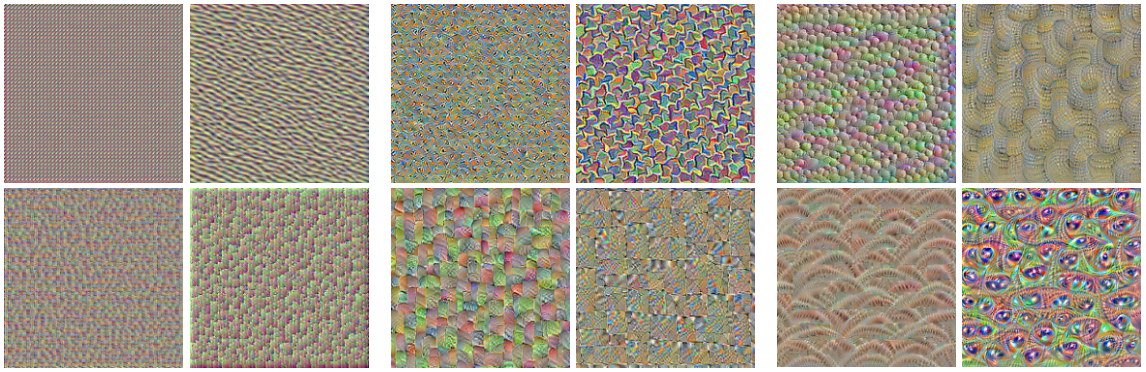

(a) Early layers (b) Intermediate layers (c) Deep layers

Notes: These filter visualizations are generated by optimizing a random input image to maximize the activation of selected filters at different depths in the network (Szegedy et al., 2015), providing insight into its internal representations. Two key aspects stand out: first, within-layer composition, where different filters specialize in detecting distinct patterns; and second, between-layer hierarchy, where deeper layers capture increasingly abstract features at larger spatial scales. Specifically, early layers (panel a) respond to fine-grained features such as colors, edges, and textures; middle layers (panel b) capture motifs and patterns at intermediate spatial scales; and deep layers (panel c) extract large-scale, abstract features, often giving rise to surreal or composite visual elements.

Figure C4. Feature projections.

(a) Input image (b) Early layer (c) Intermediate layer (d) Deep layer

Notes: This figure shows visualizations of activations from selected filters at different depths in the network, projected onto the original image of *A Sunday on La Grande Jatte* by Georges Pierre Seurat (1884), panel (a). Panel (b) displays activations from early layers, panel (c) from intermediate layers, and panel (d) from deep layers.

Convolutional layers and composition The use of multiple convolutional layers is fundamental to convolutional networks because it enables the network to model interactions hierarchically, capturing intensity patterns at different scales. Each layer builds on the previous one, creating increasingly abstract patterns and spatial relationships within the image. For example, the initial layers might capture simple features such as edges or textures; deeper layers tend to combine these to form more complex patterns, ultimately leading to high-level representations. This hierarchical approach allows convolutional networks to efficiently capture and represent the intricate structure within images, and within paintings in particular.

A way to grasp this functionality is to use filter visualization at different stages of the network (Zeiler and Fergus, 2014). In Figure C4, we consider the painting "A Sunday on La Grande Jatte" by Georges Pierre Seurat (1884) and show the activations of successive, hierarchical layers projected back onto the original image (Szegedy et al., 2015): these projections allow to better understand how convolutions and their differences throughout the network extract information about texture, contrast, or composition. For instance, panel (d)—the deep layer(s)—makes more salient: the symmetry in the foreground (between the two girls/women and the women and the dog); the general color contrasts; and the texture.

C.3 Validation

Fit and over-fitting In Section 3.2, we discuss the fit of the model across the test and training samples. The analysis then relies on the (targeted) Kullback-Leibler divergence.

To further evaluate the properties of the estimated model—specifically its ability to predict emotion scores on an out-of-sample set of paintings—we calculate the mean squared error (MSE), $E\left[\sum_{e\in\mathcal{E}}(q^e - p^e)^2\right]$, between the actual emotion ratings, $\{q^e\}_{e\in\mathcal{E}}$, and the predicted ratings, $\{p^e\}_{e\in\mathcal{E}}$, within the training sample (70% of the annotated dataset) and the test sample (15% of the annotated dataset, with the remaining 15% used for validation). The mean squared error provides an intuitive metric, allowing us to decompose the error by emotion. Figure C5 reveals several key insights: (i) the MSE is generally low (around 0.056 within the training sample and 0.067 within the test sample, indicating a minor prediction error, such as predicting 0.83 for *contentment* and 0.17 for *excitement* instead of 1 for *contentment*); (ii) the error is slightly lower for more recent paintings but the difference is not substantial (panel b); (iii) higher errors are observed for sketches, religious, and mythological paintings, likely due to their limited representation in the annotated sample (panel c); and (iv) the emotion-specific MSE shows larger errors for emotions with higher probabilities of being mentioned, but, relative to their

Figure C5. Model validation—mean squared error.

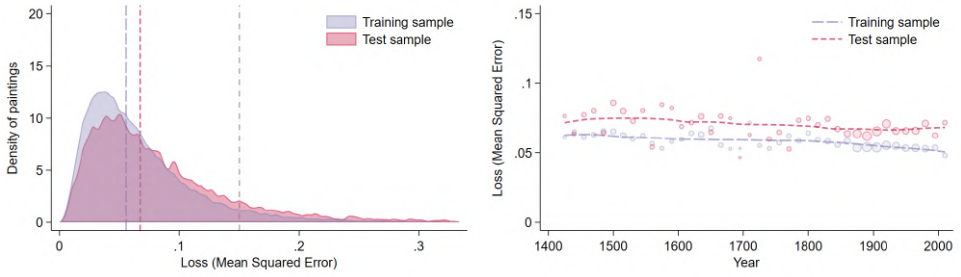

(a) Distribution (b) Period

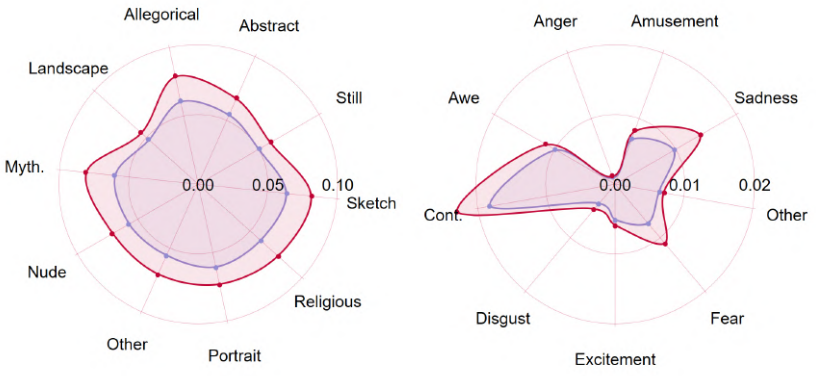

(c) Genre (d) Emotion

Notes: These figures provide statistics on the Mean Squared Error (MSE) metric between the emotion ratings and the predicted ratings within the training sample and the test sample. The metric is defined as $\sum_{e\in\mathcal{E}}(q^e - p^e)^2$, where $\{q^e\}_{e\in\mathcal{E}}$ is the actual probability distribution and $\{p^e\}_{e\in\mathcal{E}}$ is the predicted probability distribution. Panel (a) shows the distribution of this metric within the training sample (light blue) and the test sample (red); the median metric is displayed as dashed lines. For interpretation purposes, we report the median distance between predictions and randomly reshuffled labels to provide a fully uninformed benchmark. Panels (b) and (c) show the median distance metrics across sub-samples: (i) production periods, and (ii) genre or type of exercise—where *Abstract* stands for abstract, genre painting; *Landscape* stands for cityscape, landscape, marina; *Still* stands for animal painting, flower painting, still life; *Sketch* stands for sketch and study; *Nude* stands for nude painting; *Portrait* stands for portrait, self-portrait; *Religious* stands for religious painting; *Myth.* stands for mythological painting; and *Allegorical* stands for allegorical painting. Finally, panel (d) displays the emotion-specific error, i.e., the median of $\left[(q^e - p^e)^2\right]$ across all $e \in \mathcal{E}$. Relative to their importance, *sadness* and *fear* appear to have a slightly lower fit within the test sample.

importance, *sadness* and *fear* appear to have a slightly lower fit within the test sample (panel d).

In Figure C6, we provide a few illustrations of prediction quality by selecting eight artworks with a "dominant" emotion for which we benchmark the target dominant probability, q^e, against the predicted dominant probability, p^e (*co*: contentment, *am*: amusement, *ex*: excitement, *aw*: awe, *fe*: fear, *an*: anger, *sa*: sadness, and *di*: disgust).[20] In general, the model tends to accurately predict the most common emotions associated with paintings, i.e., *awe* (illustrated by opulent buildings), *contentment* (peaceful land-

[20]Anger is rarely mentioned by annotators such that it is rarely allocated a high probability, either in the annotated sample or in the prediction set.

Figure C6. Examples of predictions (*co*: contentment, *am*: amusement, *ex*: excitement, *aw*: awe, *fe*: fear, *an*: anger, *sa*: sadness, *di*: disgust).

(a) p^{co}: 0.70, q^{co}: 0.66 (b) p^{am}: 0.20, q^{am}: 0.21 (c) p^{ex}: 0.24, q^{ex}: 0.26 (d) p^{aw}: 0.65, q^{aw}: 0.55

(e) p^{fe}: 0.50, q^{fe}: 0.52 (f) p^{an}: 0.08, q^{an}: 0.08 (g) p^{sa}: 0.23, q^{sa}: 0.25 (h) p^{di}: 0.20, q^{di}: 0.22

Notes: This figure illustrates our model performance on selected paintings from the test sub-sample, where the latter are chosen as representative of each emotion. We display the dominant emotion for the target probability distribution, $\{q^e\}_{e\in\mathcal{E}}$, and the predicted probability distribution, $\{p^e\}_{e\in\mathcal{E}}$. Panel (a) shows *House with Red Roof* by Paul Cézanne (1890); the prediction and target for *contentment* are (p^{co}: 0.70, q^{co}: 0.66). Panel (b) displays *The Beauty and the Beast* by Leopold Survage (1939); the prediction and target for *amusement* are (p^{am}: 0.20, q^{am}: 0.21). Panel (c) shows *Street Decked with Flags* by Raoul Dufy (1906); the prediction and target for *excitement* are (p^{ex}: 0.24, q^{ex}: 0.26). Panel (d) displays *The Grand Canyon of the Yellowstone* by Thomas Moran (1872); the prediction and target for *awe* are (p^{aw}: 0.65, q^{aw}: 0.55). Panel (e) shows *Rough Sea* by Ivan Aivazovsky (1844); the prediction and target for *fear* are (p^{fe}: 0.50, q^{fe}: 0.52). Panel (f) displays *The Four Horsemen of the Apocalypse* by Viktor Vasnetsov (1887); the prediction and target for *anger* are (p^{an}: 0.08, q^{an}: 0.08). Panel (g) shows *During Haying* by Efim Volkov (1883); the prediction and target for *sadness* are (p^{sa}: 0.23, q^{sa}: 0.25). Panel (h) displays *Sculpture Drawing, SF64-571* by Sam Francis (1964); the prediction and target for *disgust* are (p^{di}: 0.20, q^{di}: 0.22).

scapes), *excitement* (through warm colors, lights and human activity), or *fear* (cold colors and indistinct texture evoking uncertainty).

Identification In Section 3.2, we discuss a convenient dichotomy to validate our general procedure, with two hypotheses: (H1) The model predicts the emotions sensed by labelers (which we discuss in the previous section); and (H2) the contemporary labels, and the resulting predictions, capture the artists' intentions. In Figure 7, we provide anecdotal support for hypothesis (H2) using an inversion procedure. More specifically, our neural network combines a large number of parameters into connected layers within regions of the image, then aggregates the output into averages across the entire image to finally produce a vector of emotional scores. A procedure called "Class Activation Mapping" allows to partially invert this process and measure the contribution of different parts of the image to explaining the variance in the prediction (Zhou et al., 2016).

Figure C7. The identification of emotion in paintings—additional evidence (1/2).

(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 isolate the parts of the image most used by the network to predict a latent emotional probability. We apply this methodology to a single painting—*The Promenade Woman with a Parasol* by Claude Monet, 1875—and display eight different mappings, each corresponding to one emotion (excluding "*other*"). The colors, from blue to green, yellow, and red, indicate the zones of activation for each dominant emotion.

In Section 3.2, we leverage this methodology to produce a mapping for each dominant emotion across the eight pieces displayed in Figure 3. In Figures C7 and C8, we rather focus on one artwork at a time and display eight different mappings corresponding to each emotion (but *other*). Figure C7 shows that the two characters in "The Promenade Woman with a Parasol" (Claude Monet, 1875) are predicted to suggest *contentment*, when the woman's robe (and its intricate folds) weakly convey *awe*. Figure C8 instead considers a negatively-toned painting: The Raft of the Medusa by Théodore Géricault (1818–1819), depicting a real event (the wreck of a French frigate). The central characters—crying for help—are predicted to convey *fear*, while lying (dead?) bodies strongly suggest *sadness*. Note, also and importantly, that zones are not exclusive in their emotional content: there is some overlap between the heatmaps for *fear* and *sadness*; and, reciprocally, many areas of the image are not suggestive of any emotional content.

The previous two examples provide "visual" support for hypothesis (H2). Testing hypothesis (H2) more formally is challenging: we have limited, observable information about (i) the reasoning behind the annotators' choices and (ii) the intentions and mindsets of artists at the time of production. In Figure C9, we show that the textual descriptions provided by annotators better match their allocated painting than a ran-

Figure C8. The identification of emotion in paintings—additional evidence (2/2).

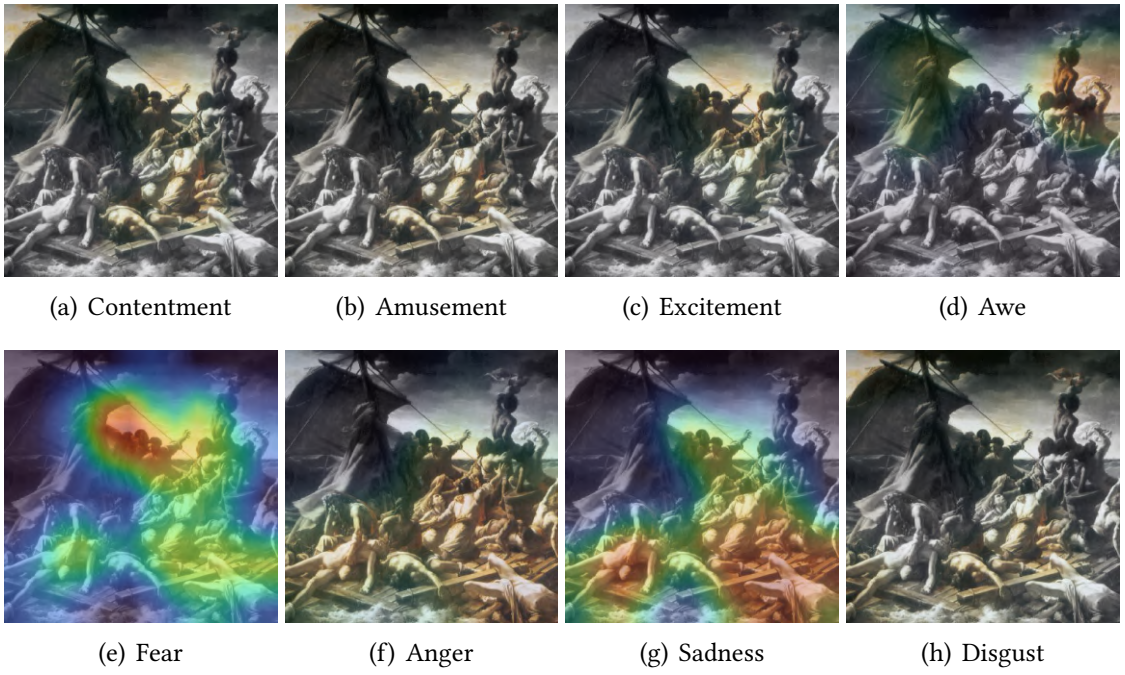

(a) Contentment (b) Amusement (c) Excitement (d) Awe

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

Notes: This figure displays heatmaps generated using Class Activation Mapping (Zhou et al., 2016)—which isolate the parts of the image most used by the network to predict a latent emotional probability. We apply this methodology to a single painting—*The Raft of the Medusa* by Théodore Géricault, 1818–1819—and display eight different mappings corresponding to each emotion (excluding "*other*"). The colors, from blue to green, yellow, and red, indicate the zones of activation for each dominant emotion.

Figure C9. The internal and external validity of annotated and predicted emotions.

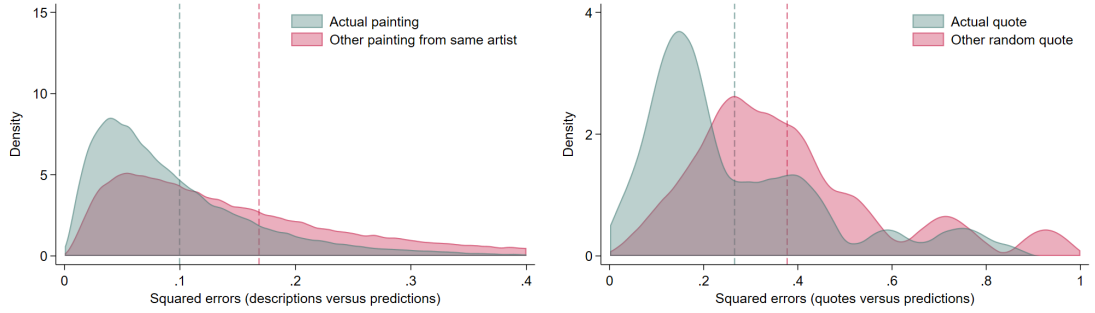

(a) Descriptions against labels (b) Predictions against quotes

Notes: Panel (a) displays the distribution of squared errors, $\sum_e (q^e - \vartheta^e)^2$, between the probabilistic vectors of emotions, as averaged across the emotions provided by annotators, $\{\mathbf{q}\}_i$, and the probabilistic vectors of emotions associated with their justifications, $\{\vartheta\}_i$. The red distribution provides a benchmark in which these justifications are randomly allocated across paintings within the same artist. Panel (b) displays the distribution of squared errors, $\sum_{e \in \mathcal{E}} (q^e - \rho^e)^2$, between the probabilistic vectors of emotions, as averaged across the emotions provided by annotators, $\{q^e\}$, and the probabilistic vectors of emotions associated with quotes from the artists, $\{\rho^e\}$. The red distribution provides a benchmark in which these quotes are randomly allocated across paintings. The textual classification relies on a transformer model (DistilBERT).

dom painting from the same artist (panel a), and we quantify the mapping between our predictions, $\{\mathbf{q}\}$, against the probabilistic vectors of emotions associated with selected

quotes from artists about their artwork, $\{\rho\}$.

D The dynamics of emotions over time

This section provides complements to Section 4.

D.1 Empirical strategy

Section 4.1 outlines our empirical strategy, which isolates variation in conveyed emotions within an artist's production cycle. In this Appendix, we first examine this residual variation and then describe a Principal Component Analysis (PCA) with two objectives: (i) to construct simplified indices associated with the probabilistic emotion vector, and (ii) to explore the latent structure underlying the probabilistic emotion vector $\{p^e\}_{e\in\mathcal{E}}$.

Figure D1. Explanatory power of artist and context fixed-effects.

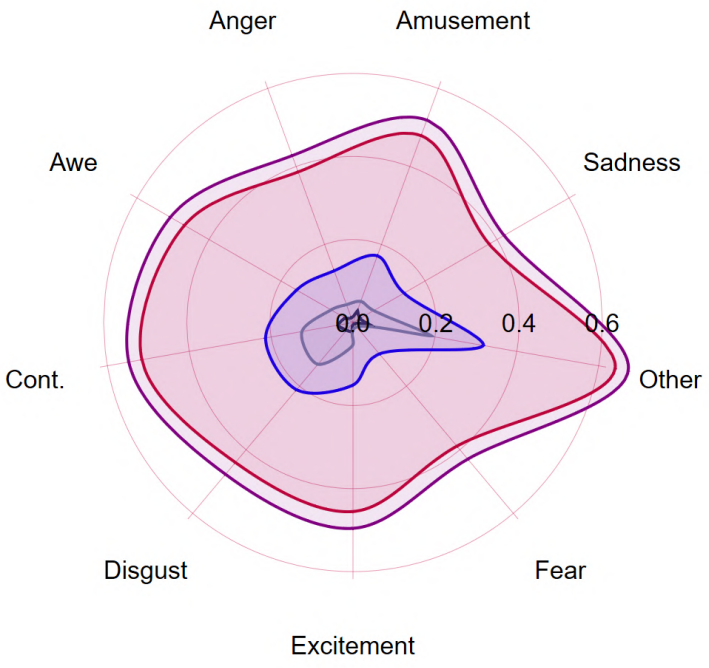

Notes: This figure reports a variance decomposition of our probabilistic emotion scores, quantifying the respective explanatory power of artist and context fixed effects. For each regression, the unit of observation is a painting j, and the left-hand side variable is the predicted probability to evoke emotion e, p_j^e. The inner, black line corresponds to the adjusted R-squared of linear regressions with 100 country fixed effects; the green line corresponds to the adjusted R-squared of linear regressions with 600 year fixed effects; the blue line corresponds to the adjusted R-squared of linear regressions with 14,384 context fixed effects; and the red line corresponds to the adjusted R-squared of linear regressions with 29,000 artist fixed effects. The purple, outer line conditions on the two previous sets of fixed effects.

Explanatory power of artist and context fixed-effects Figure D1 quantifies the respective explanatory power of artist and context fixed-effects. Let j index an image produced by artist $a(j)$ in location $\ell(j)$ and year $t(j)$, and let $p^e_{j,\ell,t}$ denote the predicted

probability of evoking emotion e. We consider specifications of the form,

$$p^e_{j,\ell,t} = \psi_{\ell,t} + \alpha_{a(j)} + \varepsilon_{j,\ell,t}$$

where $\psi_{\ell,t}$ are context fixed effects and $\alpha_{a(j)}$ represents unobserved artist effects. The inner black line in Figure D1 shows the adjusted R-squared from a linear regression with 100 country fixed effects; the green line corresponds to a specification with 600 year fixed effects; the blue line reflects a model with 14,384 context fixed effects; and the red line shows results from a regression with 29,000 artist fixed effects. The outer purple line conditions on both artist and context fixed effects.

The figure reveals that country differences explain less than 2-3% of the overall variance in depicted emotions, and common yearly dynamics across countries explain less than 8% for most emotions. In contrast, unobserved artist-level factors account for 40-50% of the variance. Conditioning additionally on context fixed effects (ℓ, t) increases the explained variance by a further 3-4 percentage points across our emotion indicators.

Table D1. A reduction of dimensionality.

	p^{c_1}	p^{c_2}	p^{c_3}	p^{c_4}
Eigenvalue	*2.56*	*1.91*	*1.35*	*1.10*
Difference	*0.65*	*0.56*	*0.25*	*0.48*
Awe	-0.234	-0.114	**0.613**	**0.444**
Contentment	**-0.503**	-0.174	-0.366	0.025
Amusement	-0.167	**0.502**	-0.211	-0.382
Excitement	-0.190	0.351	**0.421**	**-0.412**
Fear	**0.424**	-0.204	0.346	-0.221
Anger	**0.456**	0.143	0.033	-0.027
Sadness	0.340	**-0.418**	-0.320	-0.140
Disgust	0.323	**0.448**	-0.058	0.109
Other	0.146	0.381	-0.205	**0.637**

Notes: This table reports the loadings of the different emotion scores from the Principal Component Analysis (PCA). We present the eigenvalue associated with each eigenvector of the variance-covariance matrix in italics, along with the difference from the next lower eigenvalue.

A reduction of dimensionality In Table D1, we instead explore the *cross-emotion* variation and identify latent factors underlying our probabilistic emotion vector $\{p^e\}_{e\in\mathcal{E}}$ through a Principal Component Analysis. The first principal component ($p^{c_1}_j$) loads positively on negative emotions (*fear*, *anger*, *sadness*, and *disgust*) and is associated with a very high eigenvalue. As suggested by Figure 8, which focuses on *fear*, this principal component captures meaningful fluctuations around salient, often geopolitical, events.

The second component ($p^{c_2}_j$) also loads on non-neutral emotions, including *amusement*, *excitement*, as well as *disgust* and *other*, and is associated with a reasonably high

eigenvalue. As shown in Appendix Figure D6, its fluctuations only partly reflect secular trends in the depiction of positive emotions, coinciding with rising living standards across developed economies.

The third component ($p_j^{c_3}$) loads predominantly on *awe*, while the fourth component ($p_j^{c_4}$) is primarily associated with the residual category *other*. An interesting feature of the data is that *contentment*, the most frequently expressed emotion, is not positively loaded on any of these principal components: it is negatively loaded on the first component, suggesting that artists may substitute *contentment* with emotions such as *fear*, *anger*, and *sadness* during turbulent times.

An illustration of the empirical strategy based on well-known artists In this section, we discuss our empirical strategy through the lens of the following 30 well-known artists, representative of their respective eras (in parentheses):

- Jan van Eyck (1400–1450);
- Hieronymus Bosch (1450–1500);
- Giovanni Bellini (1450–1500);
- Sandro Botticelli (1450–1500);
- Albrecht Dürer (1500–1550);
- Leonardo Da Vinci (1500–1550);
- Michelangelo Buonarroti, known as Michelangelo (1500–1550);
- Raphael Sanzio da Urbino, known as Raphael (1500–1550);
- Doménikos Theotokópoulos, known as El Greco (1550–1600);
- Pieter Bruegel The Elder (1550–1600);
- Tiziano Vecellio, known as Titian (1550–1600);
- Michelangelo Merisi da Caravaggio (1600–1650);
- Diego Velazquez (1600–1650);
- Peter Paul Rubens (1600–1650);
- Johannes Vermeer (1650–1700);

- Rembrandt Harmenszoon van Rijn, known as Rembrandt (1650–1700);
- Giovanni Battista Tiepolo (1700–1750);
- Jean-Antoine Watteau (1700–1750);
- Jacques-Louis David (1750–1800);
- Francisco Goya (1750–1800);
- Eugene Delacroix (1800–1850);
- Joseph Mallord William Turner (1800–1850);
- Claude Monet (1850–1900);
- Edouard Manet (1850–1900);
- Vincent Van Gogh (1850–1900);
- Henri Matisse (1900–1950);
- Pablo Picasso (1900–1950);
- Wassily Kandinsky (1900–1950);
- Andy Warhol (1950–2000);
- Francis Bacon (1950–2000);
- Jackson Pollock (1950–2000)

Analyzing the emotional content of these artists' work yields the following insights. First, they differ markedly in their emotional expression. Table D2 presents a systematic selection of the most expressive artists by emotion.

Appendix Figure D2 then illustrates these differences for three prominent artists, Wassily Kandinsky, Jan van Eyck, and Jackson Pollock. Panel (a) shows the average emotion scores for each artist across nine emotion categories, capturing distinct emotional profiles: Kandinsky expresses relatively high levels of *amusement* and *excitement*, van Eyck emphasizes *awe* and *contentment*, while Pollock features more *anger*, *fear*, and *disgust*. Panel (b) highlights the substantial variation *within* each artist's body of work by plotting the distribution of the *positivity* index across their respective paintings. The spread of each distribution indicates that, even within a single artist's oeuvre, emotional tone can vary widely from one painting to another.

Table D2. Most expressive artists by emotion.

Emotion	Artists
Contentment	Monet, Vermeer
Amusement	Picasso, Kandinsky
Excitement	Kandinsky, Warhol
Awe	Tiepolo, van Eyck, Rubens
Fear	Bruegel the Elder, Bosch, Bacon, Pollock
Anger	Pollock
Sadness	Caravaggio, Bellini, Rembrandt, El Greco
Disgust	Pollock
Other	Pollock

Notes: This table presents a systematic selection of the most expressive artists by our nine emotion categories: *contentment*, *amusement*, *excitement*, *awe*, *fear*, *anger*, *sadness*, *disgust*, and *other*.

Figure D2. Emotional expressions of different artists (Kandinsky, van Eyck, and Pollock).

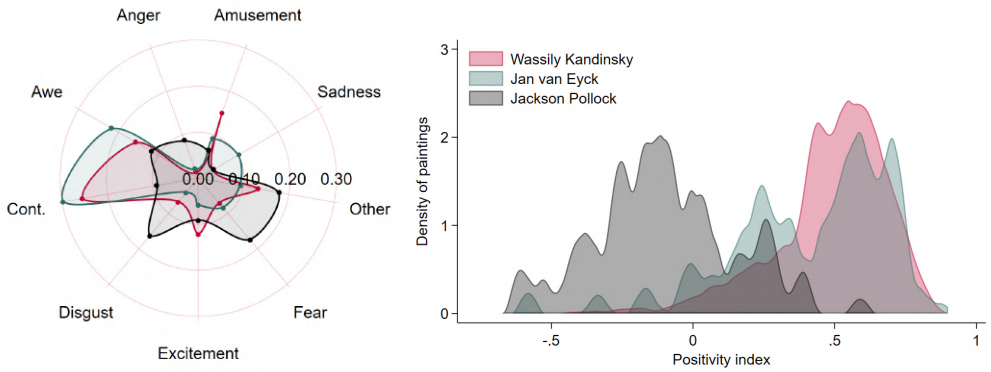

(a) Emotions (b) Within-variation

Notes: Panel (a) displays the distribution of predicted emotion scores for 1,025 paintings produced by Wassily Kandinsky (red), Jan van Eyck (green), and Jackson Pollock (black). Panel (b) displays the distribution of the *positivity* index, I_i, across the same set of 1,025 paintings.

Second, there is substantial variation in emotional expression within each artist's production cycle. Figure D3 extends the insight from the previous Figure D2 to the full sample of 30 selected artists. Panel (a) places each artist in "emotion space," using the average *positivity* index to summarize overall emotional tone and a within-artist disagreement index to capture the extent of emotional variability across works. Panel (b) shows the temporal coverage of the sample by century, revealing distinct clusters in artistic production over time.

The results highlight notable differences in average emotional tone across artists. Some, like Jan van Eyck and Johannes Vermeer, exhibit consistently high positivity, while others, such as Jackson Pollock and Caravaggio, display much lower average scores.

Figure D3. Emotional expressions of different artists (30 selected artists).

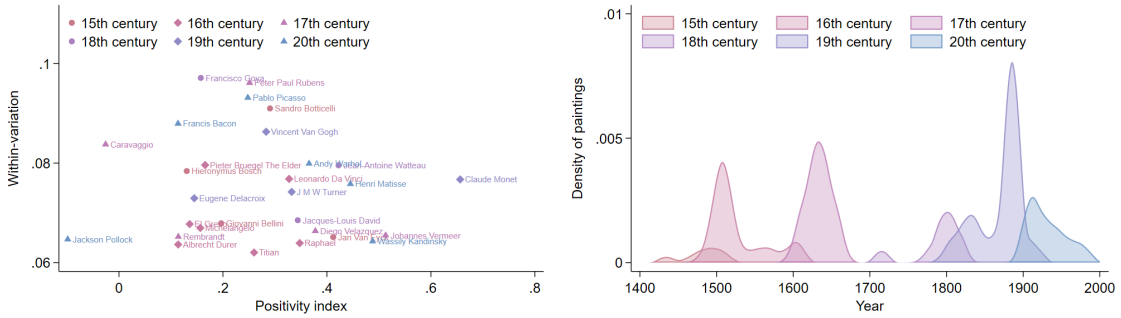

(a) Emotions (b) Coverage

Notes: Panel (a) displays the degree of disagreement within artist against the average positivity index, I_i—the figure is based on 27,387 paintings produced by 30 well-known artists: Dürer, Warhol, Caravaggio, Monet, Velazquez, Manet, El Greco, Delacroix, Bacon, Goya, Tiepolo, Bellini, Matisse, Bosch, Turner, Pollock, David, van Eyck, Watteau, Vermeer, Da Vinci, Michelangelo, Picasso, Rubens, Bruegel the Elder, Raphael, Rembrandt, Botticelli, Titian, van Gogh, Kandinsky. The disagreement measure is the average of the sum (across emotions) of absolute deviations between each painting's score and the artist's average score. Note that the disagreement measure would be distributed as follows if emotion scores were randomly assigned: 0.035, 0.043, 0.064. Panel (b) displays the distribution over time of the paintings covered in the previous exercise.

The within-artist disagreement index further reveals meaningful disparities in emotional range across artists. While painters like van Eyck, Vermeer, and Kandinsky demonstrate relatively focused emotional expression, others, including Francisco Goya and Pablo Picasso, exhibit wide variation. In the case of Goya, this range reflects the contrast between his early, decorative works and his later, darker period marked by political violence and psychological unrest. For Picasso, the emotional diversity corresponds to his successive stylistic periods, from the Blue Period to Cubism and beyond.

Additionally, certain individual paintings diverge sharply from their creators' typical emotional profiles. For example, Claude Monet's *Still Life with Pheasants and Plovers* (1879) is unusually somber, reflecting the personal grief following the death of Camille Doncieux. Similarly, Sandro Botticelli's *Lamentation over the Dead Christ with the Saints Girolamo, Pietro, and Paolo* carries the weight of Counter-Reformation austerity, likely influenced by the spiritual severity of Girolamo Savonarola's rule in Florence.

Third, variation in emotional expression sometimes reflects broader secular changes in the artists' environments. Sandro Botticelli, Caravaggio, Titian, Pieter Bruegel the Elder, and El Greco all experienced the chastening effects of the Counter-Reformation in Italy or the rule of Philip II in the Low Countries and Spain. To quantify the impact of these contextual "shocks" on conveyed emotions, we embed them within a two-way fixed effects event study, controlling for artist, time, and country fixed effects (see Figure D4). We find that fear increases by approximately 2 percentage points, contentment decreases by 3 percentage points, and the positivity index declines by around 0.10.

Figure D4. The dynamics of emotions after counter-reformations (Sandro Botticelli, Pieter Bruegel the Elder, Titian, El Greco, Caravaggio).

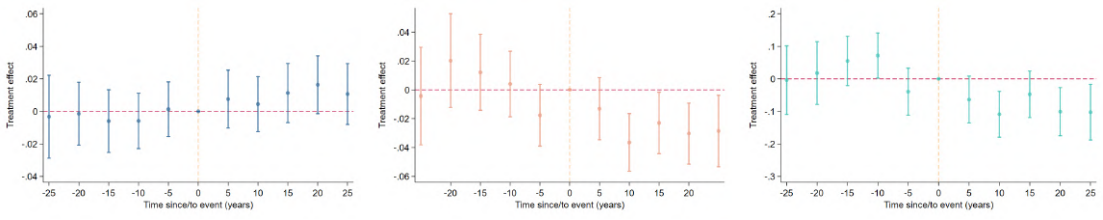

(a) Fear (b) Contentment (c) Positivity

Notes: This figure presents a collapsed two-way fixed effects event study, controlling for artist fixed effects, time fixed effects, and country fixed effects. The specification considers 5 events: Sandro Botticelli, religious extremism, 1490; Pieter Bruegel the Elder, religious persecution, 1560; Titian, Counter-Reformation, 1545; El Greco, Spanish Inquisition, 1582; Caravaggio, Counter-Reformation, 1595. For visualization purposes, we group years into the closest five-year intervals to identify the estimates β_τ.

D.2 The dynamics of emotions across countries

Negative emotions over time Figure D5 depicts the evolution of negative emotions within the borders of 12 contemporary countries, each contributing more than 10,000 artworks: France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. The selected negative emotions are *fear* (blue, plain), *sadness* (teal, long dash), and *anger* (black, shorter dash). As discussed in Section 4.2, the dynamics of *fear* display considerable short-run volatility, which appears to be context-specific—except during major global conflicts, such as the World Wars. A few countries, e.g., as France, Germany, Austria, Spain, and Belgium, exhibit more impermanent emotional patterns, in contrast to the more stable trajectories observed in the United States or the United Kingdom.

For example, Austria experienced large fluctuations in fear corresponding to the numerous conflicts affecting its territory prior to 1800—a strong correlation between the depiction of *fear* and the outbreak of war that appears across many societies. The evolution of *sadness* sometimes mirrors or follows that of fear, but not consistently. The end of the nineteenth century and the turn of the twentieth century are "sad" periods for most countries—a pattern echoed in the rise of nineteenth-century European realism—while the remainder of the twentieth century is not. One possible explanation is that the sustained rise in living standards, which began around 1800, was only broadly felt by the average individual in our sample locations after 1900, coinciding with the expansion of leisure, labor protections, social welfare, and a general decline in inequality.

Finally, *anger* emerges as a more residual emotion, with limited variation prior to the second half of the twentieth century.

Figure D5. The dynamics of negative emotions (*fear*, *sadness*, and *anger*) over time.

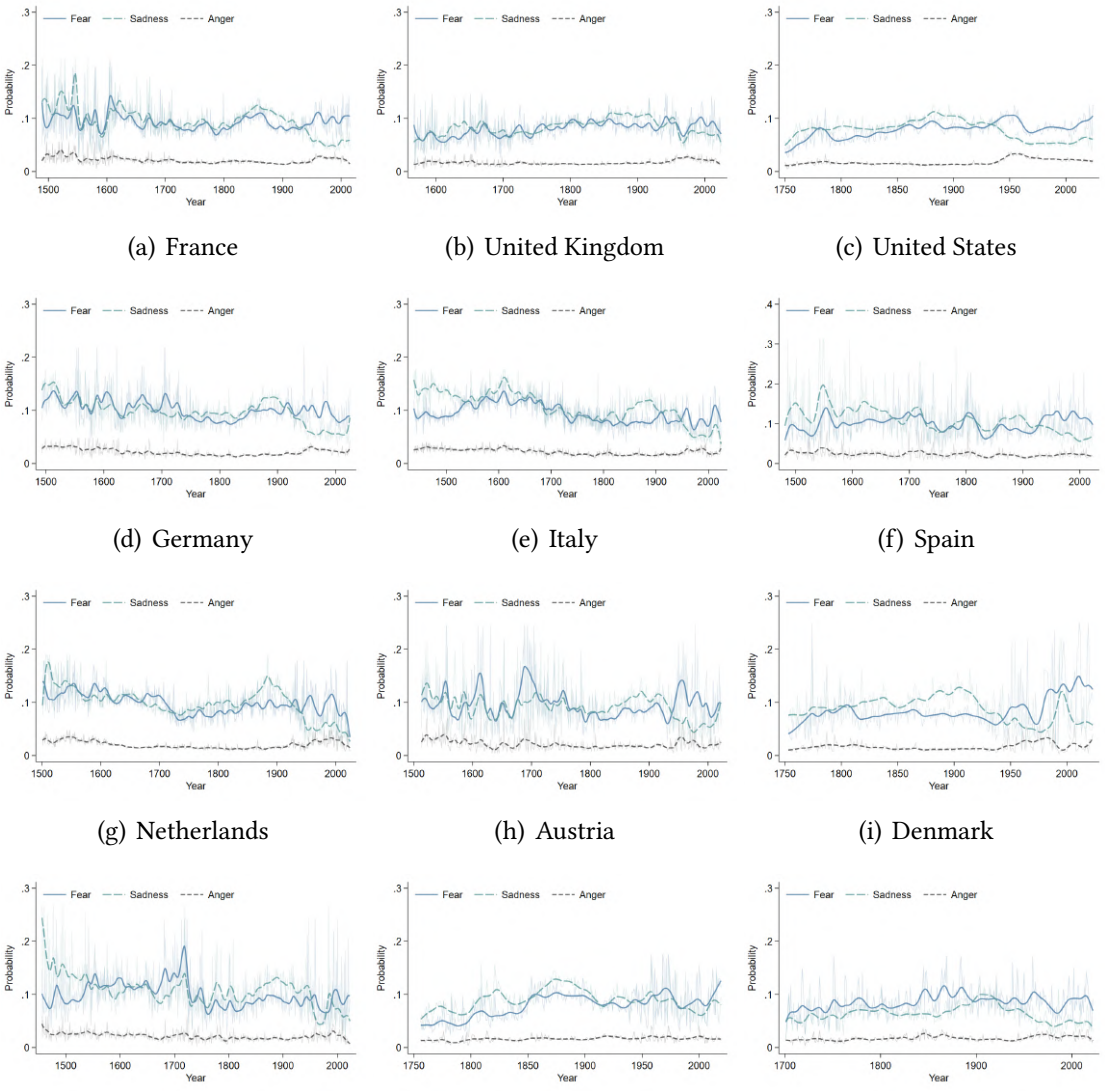

(a) France (b) United Kingdom (c) United States

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

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

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

Notes: This figure illustrates the dynamics of the predicted probability to evoke emotion e, $p^e_{\ell,t}$, for negative emotions in France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. Negative emotions are: *fear* (blue, plain), *sadness* (teal, long dash), and *anger* (black, shorter dash). 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 coefficient of 100—as a thicker line. Note that the respective starting dates are set to exclude the first percentile of artworks in terms of production time, and each observation within a year is weighted to account for duplicates and the overall production of each artist.

Positive emotions over time Figure D6 depicts the evolution of positive emotions within the borders of twelve contemporary countries. The selected positive emotions are *excitement* (orange, plain), *amusement* (pink, long dash), and *contentment* (salmon, shorter dash). *Contentment* and *amusement* display non-negligible volatility, the former also showing significant secular trends: predicted *contentment* steadily rises (or remains stable) from 1500–1900 and sharply decreases afterward before stabilizing from 1950–

Figure D6. The dynamics of positive emotions (*excitement*, *amusement*, and *contentment*) over time.

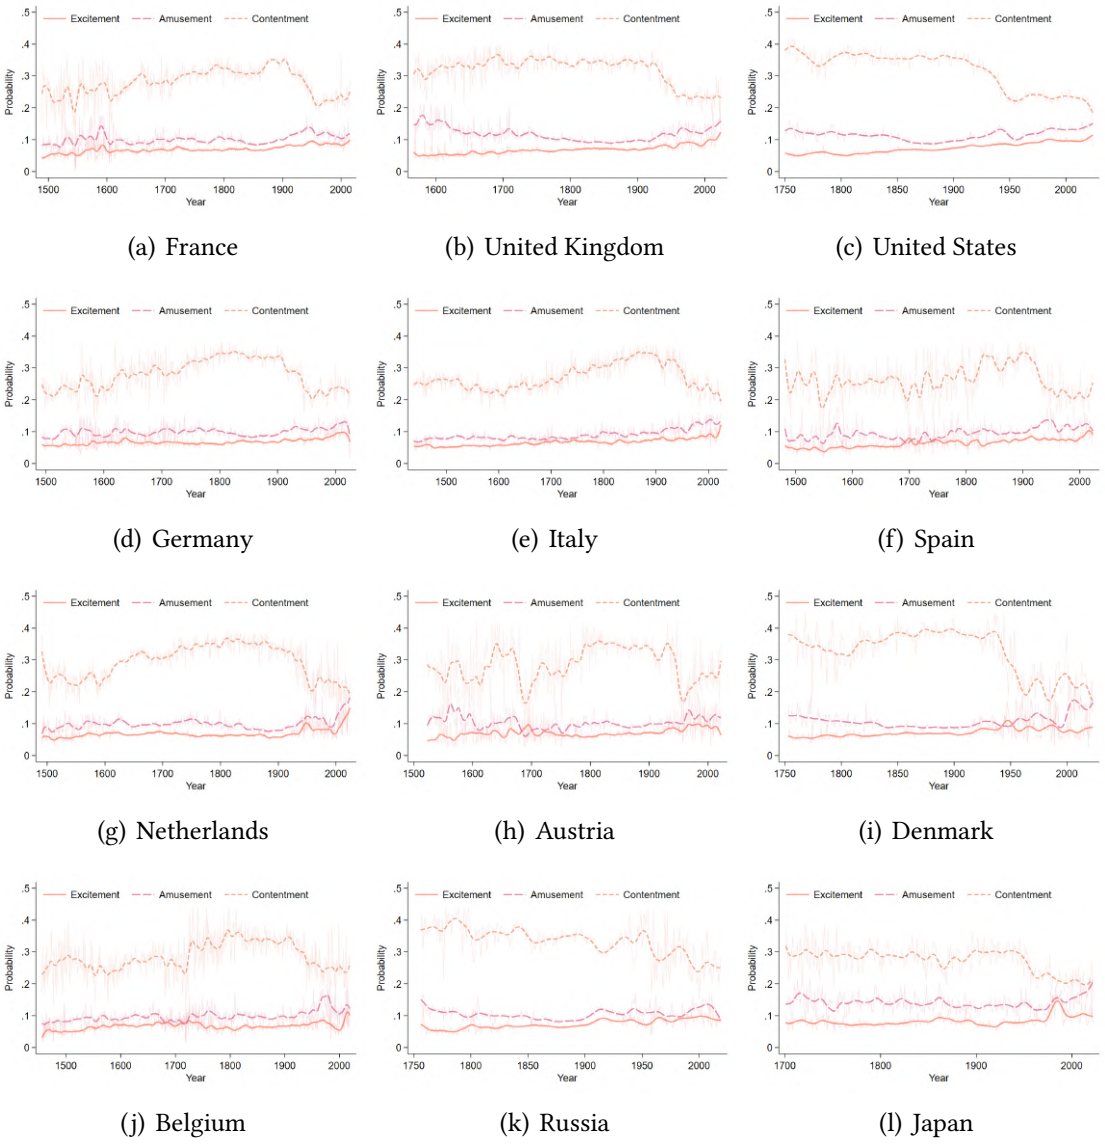

(a) France (b) United Kingdom (c) United States

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

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

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

Notes: This figure illustrates the dynamics of the predicted probability to evoke emotion e, $p^e_{\ell,t}$, for positive emotions in France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. Positive emotions are: *excitement* (orange, plain), *amusement* (pink, long dash), and *contentment* (salmon, shorter dash). 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 coefficient of 100—as a thicker line. Note that the respective starting dates are set to exclude the first percentile of artworks in terms of production time, and each observation within a year is weighted to account for duplicates and the overall production of each artist.

2020 (yearly variations across countries explain about 22% of the overall variance in depicted *contentment*, $p^c_{\ell,t}$). By contrast, *excitement* is less often evoked and apparently less volatile.

Other emotions over time and a positivity index Figure D7 focuses on the three remaining emotions: *awe* (gold, plain), *disgust* (green, long dash), and *other* (gray, shorter

Figure D7. The dynamics of other emotions (*awe*, *disgust*, and *other*) over time.

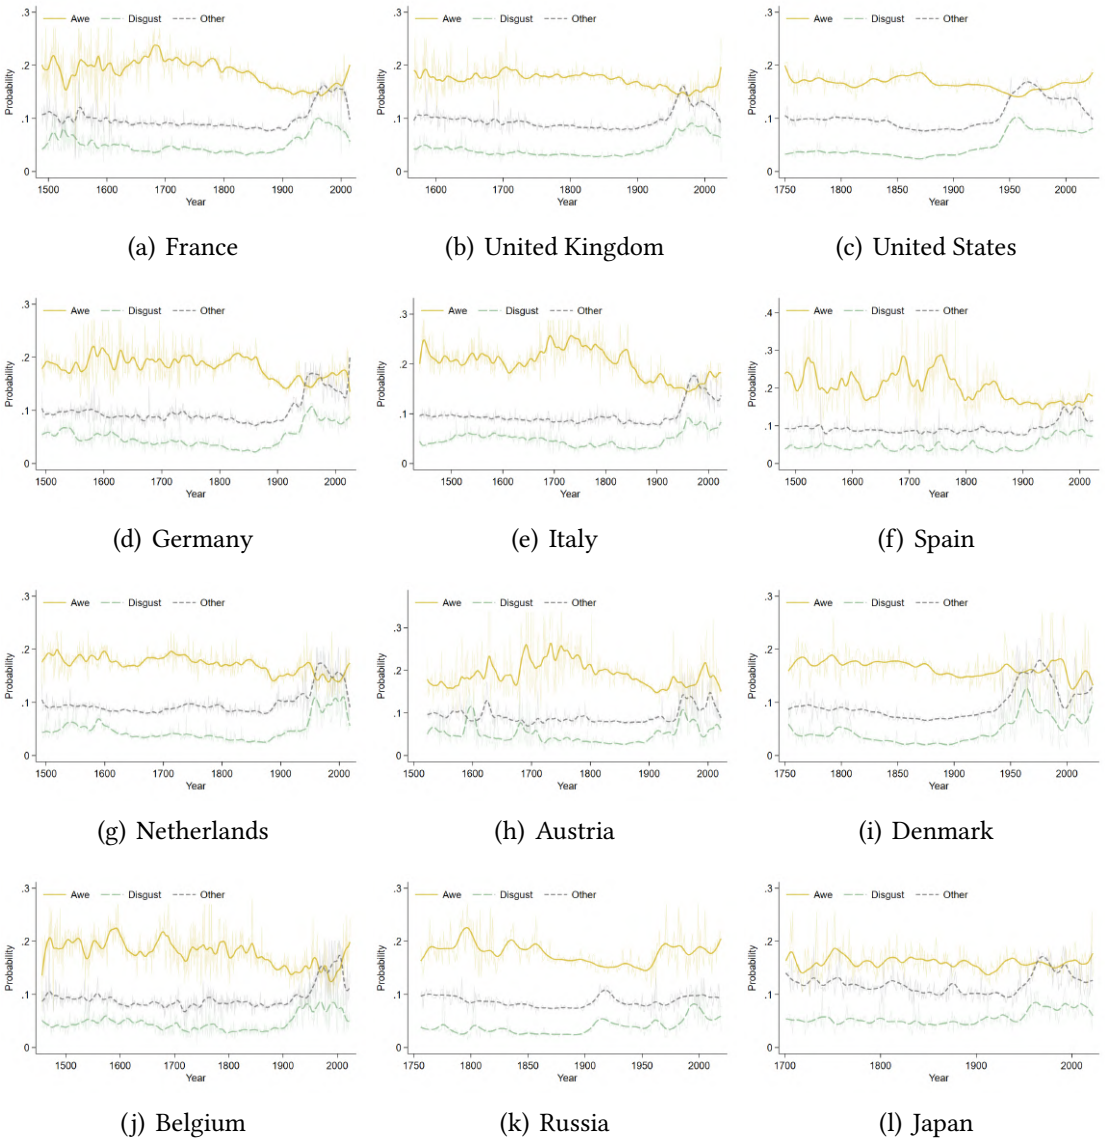

(a) France (b) United Kingdom (c) United States

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

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

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

Notes: This figure illustrates the dynamics of the predicted probability to evoke emotion e, $p^e_{\ell,t}$, for other emotions in France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. Other emotions are: *awe* (gold, plain), *disgust* (green, long dash), and *other* (gray, shorter dash). 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 coefficient of 100—as a thicker line. Note that the respective starting dates are set to exclude the first percentile of artworks in terms of production time, and each observation within a year is weighted to account for duplicates and the overall production of each artist.

dash). The last two emotions follow similar dynamics across all societies: they are flat and with limited volatility between 1500 and 1900 before they rise swiftly during the twentieth century. *Awe*, in a similar way as *fear* or *contentment*, is more volatile and with fluctuations that seem less correlated across economies. For instance, it reaches a peak during the reign of Louis XIV in France and a low point before the French Wars of Religion; it varies markedly in Spain, with local maxima around 1520 (under the rule of

Charles V), 1660 (before the end of the Habsburg dynasty), and 1750 (under the Bourbon dynasty); and it increases significantly after the Second World War in the (then) Soviet Union, possibly reflecting propaganda efforts and the regime's use of cultural production for political support.

Figure D8. The dynamics of *positivity* ($I_{\ell,t}$) over time.

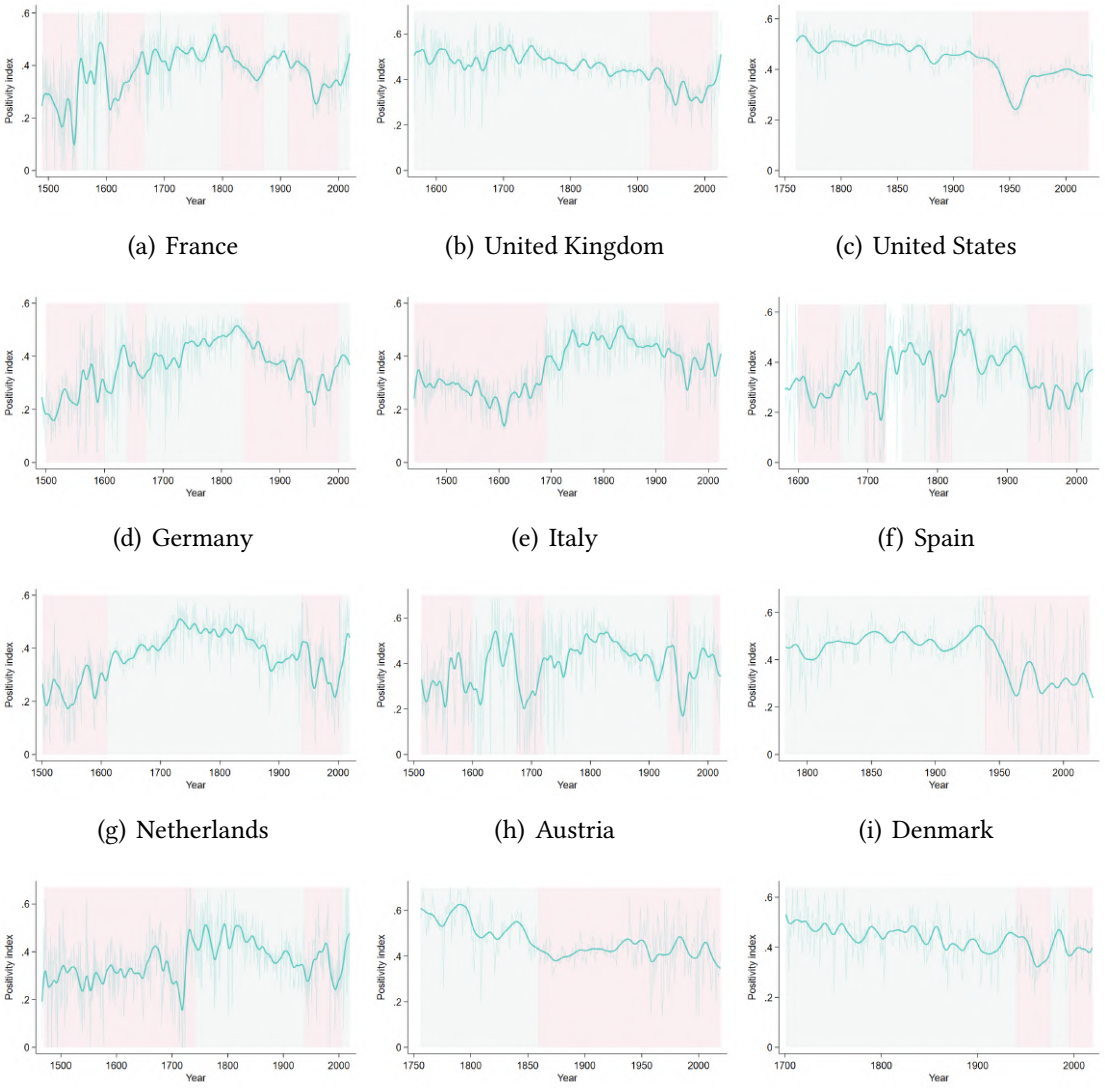

(a) France (b) United Kingdom (c) United States

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

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

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

Notes: This figure illustrates the average evolution of the *positivity* index ($I_{\ell,t}$) 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 coefficient of 100—as a thicker line. Note that the respective starting dates are set to exclude the first percentile of artworks in terms of production time, and each observation within a year is weighted to account for duplicates and the overall production of each artist. The red/green areas indicate general periods of negativity/positivity.

Finally, we collapse all emotion indices into a *positivity* index ($I_{\ell,t}$) in Figure D8 and show that: (i) there are differences in levels, with Great Britain, the United States, and

Denmark being persistently more positive in earlier periods, while Italy, Spain, Germany, and the Low Countries are persistently more negative; (ii) there are differences in volatility, with France, Germany, Spain, and Austria displaying the most variation over time; (iii) many continental European countries experience a prolonged negative episode, possibly coinciding with the Little Ice Age; (iv) the eighteenth and nineteenth centuries are marked by a sustained period of *positivity*; and (v) the geopolitical context of 1914–1963 appears to correspond with a widespread decline in *positivity* across all countries.

Interestingly, *positivity* reaches its highest point around 1850—coinciding with the German revolutions of 1848–1849—within the borders of contemporary Germany, followed by a long decline that only reverses after World War II. Additional spikes in positivity are also observable: in France during la Belle Époque, in the early twentieth century or after the liberation from Ottoman rule in Austria, during the Dutch Golden Age in the Netherlands, and under the reign of Nicholas I in Russia.

The disparity in emotions within a context The previous evidence sheds light on the "commonality" of different artworks produced in a given context. In Figure D9, we illustrate the "disparity" in expressed emotions within a context, $\delta_{\ell,t}$, as captured by the distance to the average depiction within a context, $1/|\mathcal{E}| \sum_{e \in \mathcal{E}} |p_j^e - \bar{p}_{\ell,t}^e|$, and its average across all paintings produced in any given year t and location ℓ, and by a similar index but based on emotion scores residualized by the identity of artists, their age, the harmonized movements, and the harmonized exercises.

One striking shift in the disparity of displayed emotions over time appears in the second half of the sixteenth century in France, corresponding to the French Wars of Religion—during which many artists fled to Great Britain, Saxony, Prussia, the Low Countries, and Switzerland.

Alternative indicators of economic development In this Appendix, we also contrast our own measures of welfare to alternative, sometimes more standard, measures of economic development in economic history, where and when such data are available (using data from Hills et al., 2019; Zijdeman and Ribeira da Silva, 2015). While our baseline approach focuses on measures of GDP per capita and wealth concentration (see Figure 9), we correlate the predicted probability of evoking each emotion $e \in \mathcal{E}$, p_j^e, the *positivity* index, I_j, and the *disagreement* index, d_j, with: a sentiment analysis of books written in a few countries (Hills et al., 2019); and a measure of life expectancy at birth (Zijdeman and Ribeira da Silva, 2015).

In Figure D10, one can see that *sadness* and the aggregate *positivity* index correlate strongly with these alternative measures of living standards, in much the same way as

Figure D9. The dynamics of *disagreement* ($d_{\ell,t}$) across countries and over time.

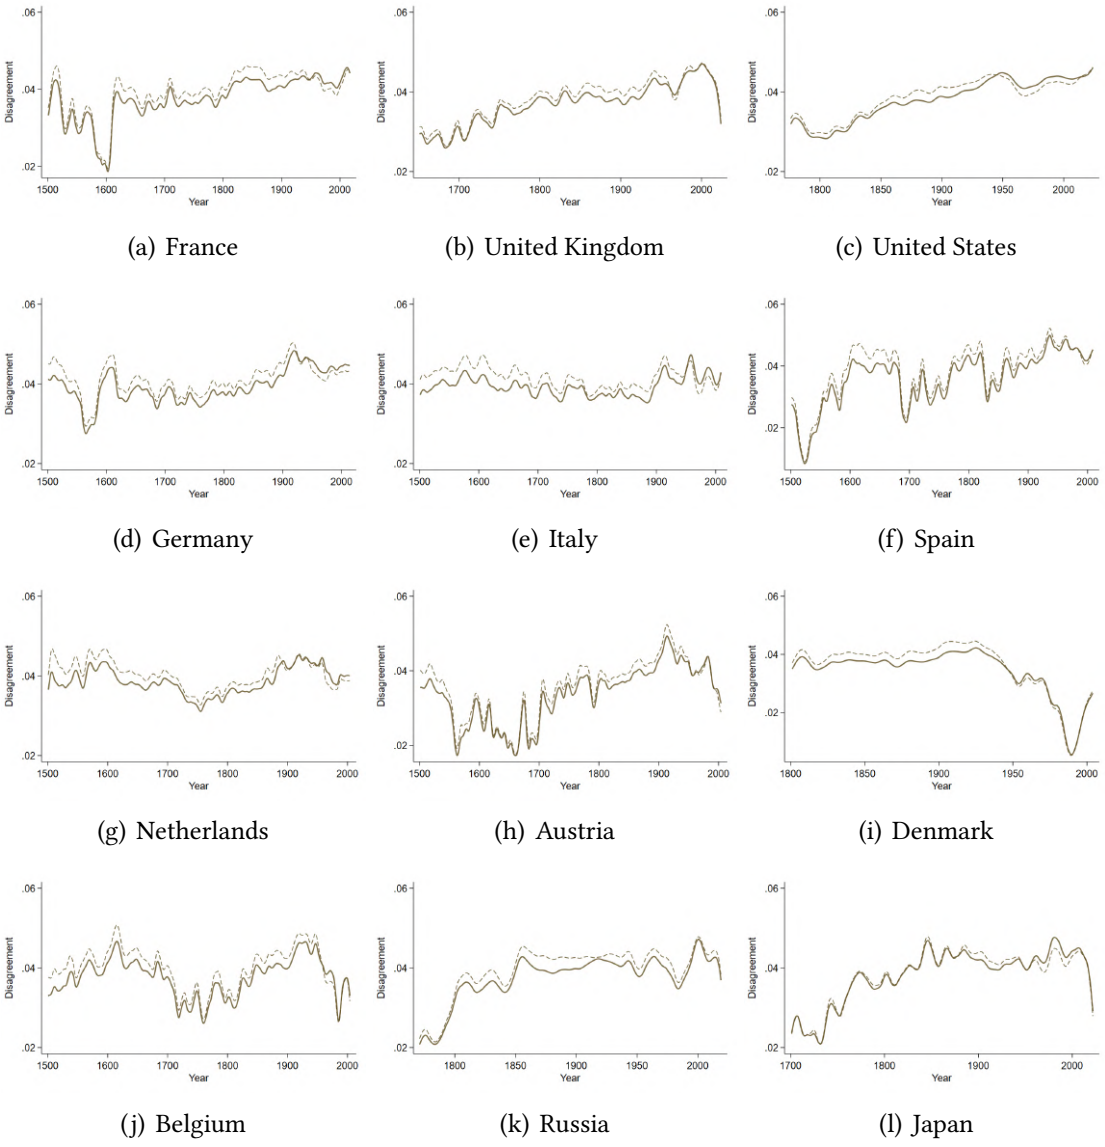

(a) France (b) United Kingdom (c) United States

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

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

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

Notes: This figure displays the dynamics of *disagreement*, $d_{\ell,t}$, in France, the United Kingdom, the United States, Germany, Italy, Spain, the Netherlands, Austria, Denmark, Belgium, Russia, and Japan. For illustration purposes, the time series are smoothed using a Hodrick-Prescott filter with a coefficient of 100, and the respective starting dates are set to include the bulk of artworks in terms of production time. The context-level disagreement index is computed as the average level of disagreement across all paintings produced in a given year t and location ℓ, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}}|p_j^e - \bar{p}_{\ell,t}^e|$. The dashed line is based on measures that are residualized by genre, movement, age, and artist fixed effects.

with the baseline measures. However, compared to Figure 9, these alternative indicators appear to give more prominence to awe than to other positive emotions, possibly because their variation is concentrated in the nineteenth century.

Figure D10. Alternative indicators of economic development and the emotions conveyed through paintings.

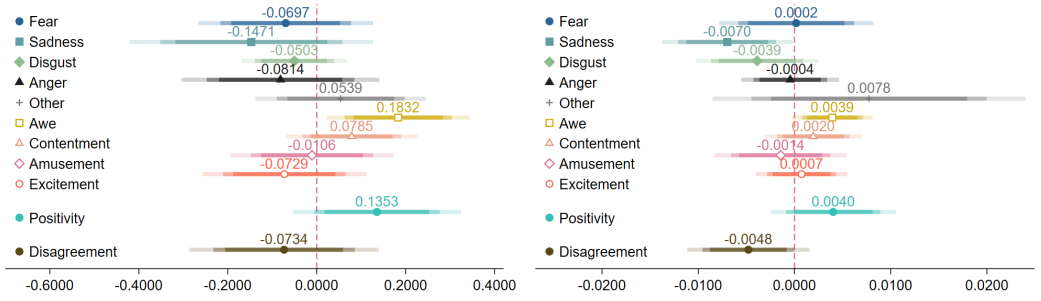

(a) A valence index (b) Life expectancy

Notes: These Figures display the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to alternative indicators of economic development. 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 "valence index" using sentiment analysis on published books (Hills et al., 2019); and (b) a measurement of life expectancy at birth (Zijdeman and Ribeira da Silva, 2015). 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, $\frac{1}{|\mathcal{E}|}\sum_{e\in\mathcal{E}} |p_j^e - \tilde{p}_{\ell,t}^e|$. 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.

D.3 The perception of large socio-economic changes

This section provides complements to Section 4.3.

Basic subsistence costs In panel (a) of Figure D11, we provide a complement to Figure 10: climatic fluctuations have very marked effect on the emotional expression of artists, possibly due to their impact on basic subsistence costs. To support this hypothesis, we regress our emotion indicators on a residual measure of bread prices in Figure D11 (a): A 50% increase in the price of bread raises the probability of evoking *anger* and *sadness* by 0.023–0.033 standard deviations, at the expense of *awe* and *contentment*—mirroring the effects of temperatures.

Luxury trade before the first wave of globalization Panel (b) of Figure D11 complements Figure 11: before the first wave of globalization, trade in luxury goods did not benefit a broad share of the population, and we observe different effects compared to the large trade expansion that followed. Specifically, an increase in the price of luxury goods is associated with heightened depictions of *anger* and *excitement*.

Figure D11. Basic subsistence costs, and luxury trade.

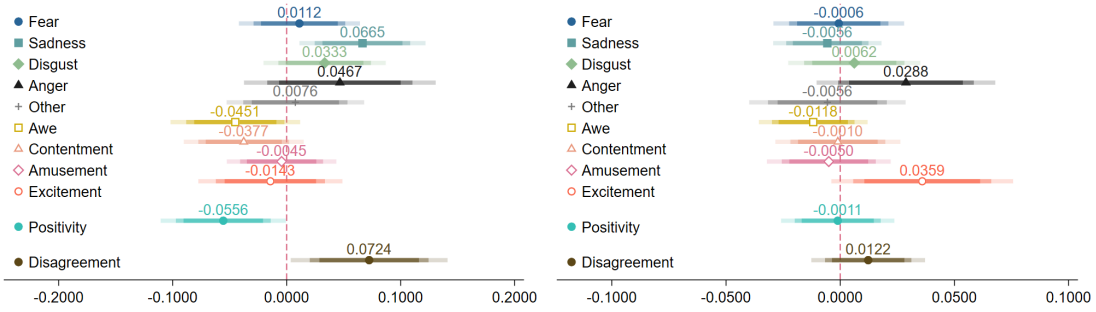

(a) Price of bread (b) Price of luxury

Notes: These figures display the estimates of regressions relating the predicted emotions of a given painting, $\{p^e\}_{e\in\mathcal{E}}$, to local prices. 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) the residualized (log) price of bread as inferred from the market prices reported in the Allen-Unger Global Commodity Prices Database (Allen and Unger, 2019, see Appendix B.4); and (b) the price of luxury goods. 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 country/decade.

Universities, the decline of religiosity, and the scientific revolution The evidence reported in Section 4.3 focuses on academic openness—measured by the (log) number of foreign scholars—and the share of scientific research in total academic production. Figure D12 (panels a and b) further illustrates the large 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, based on two additional characteristics: research "quality"—measured by the (log) number of influential scholars—and religiosity—the share of religious research in total academic production. Across these economies, academic output increases steadily over time, though less markedly in Italy and more prominently in Great Britain. Spain is a notable exception, where academic output stagnates from 1550 onward, following the reign of Charles V, Holy Roman Emperor.

Appendix Figure D12 (c) shows that academic output is associated with a more positive emotional tone of contemporary artistic output across all countries and periods. An increase in academic output equivalent to the secular trend in the United Kingdom between 1400–1800 would be associated with -0.10 standard deviations in the depiction of *fear* and increases of 0.10 and 0.08 standard deviations in the depiction of *contentment* and *amusement*, respectively. Appendix Figure D12 (d) shows that the relative dominance of religiosity against scientific research in academic production yields negative

Figure D12. Universities, the decline of religiosity, and the scientific revolution.

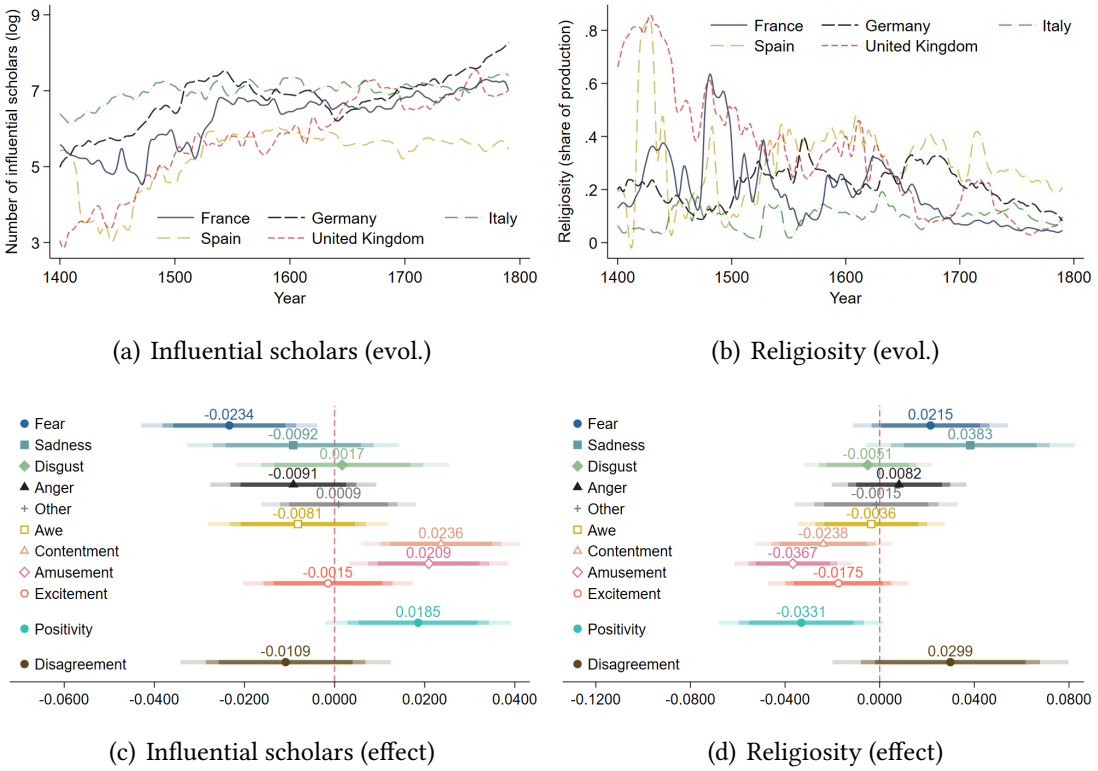

(a) Influential scholars (evol.)

(b) Religiosity (evol.)

(c) Influential scholars (effect)

(d) Religiosity (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 "quality" or influence, as measured by the (log) number of scholars with more than 10 Wikipedia pages (under different languages); and panel (b) reports the weighted share of knowledge production in theology. 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: (a) the (log) number of scholars with more than 10 Wikipedia pages and (b) the weighted share of knowledge production in theology. 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 country/decade.

emotions—particularly *sadness* and *fear*.

Labor strikes, emancipation, and revolutions Section D.3 of the paper focuses on objective indicators of regime characteristics (see, e.g., Figure 14). In Figure D13, we turn to a list of key events and define indicators equal to 1 in the ten years following a major reform of labor laws and slavery (panel a) or a revolution/transition to a democratic regime (panel b). The insights from Figure D13 suggest that reforms reducing labor coercion and political revolutions significantly influence the emotions conveyed in artworks—and in a generally positive direction. Following labor emancipations or

Figure D13. Labor strikes, emancipation, revolutions, and the emotions conveyed through paintings.

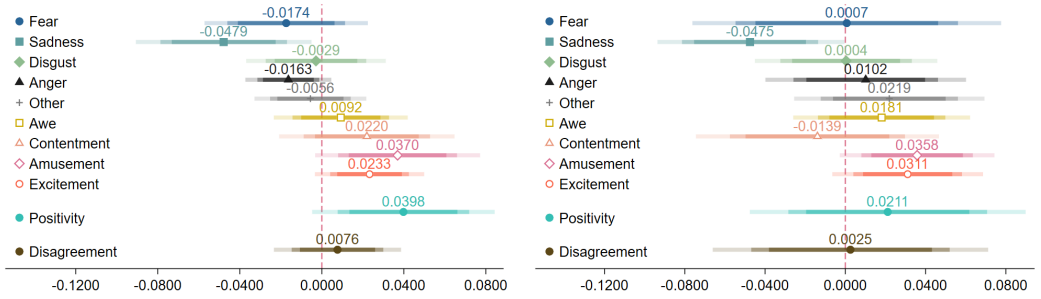

(a) Labor unions/emancipation (b) Revolutions

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 reforms and revolutions. 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) an indicator equal to 1 in the ten years following a major reform of labor laws and slavery (e.g., in France, the Waldeck-Rousseau law recognizing the right to organize labor unions in France, May 1968, or the abolition of slavery in its colonies); and (b) an indicator equal to 1 in the ten years following a transition to democracy. 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 country/decade.

revolutions, the *positivity* index increases by 0.04 standard deviations, primarily driven by a more prevalent depiction of *amusement* at the expense of *sadness*.

Figure D14. The dynamics of emotions during conflicts.

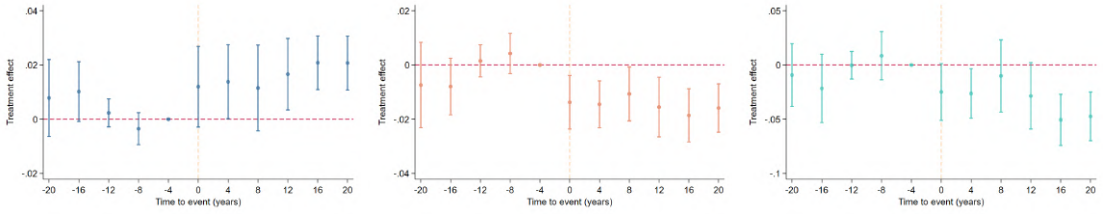

(a) Fear (b) Contentment (c) Positivity

Notes: This figure presents a collapsed two-way fixed effects event study, controlling for artist fixed effects, time fixed effects, and country fixed effects. Standard errors are clustered at the country level. The specification considers 24 events affecting France, Germany, Italy, the Netherlands, Spain, the United Kingdom, the United States, China, Austria, Japan, Portugal, Russia, Sweden, Switzerland, and Turkey (e.g., the French Wars of Religion, the Thirty Years' War, the War of the Austrian Succession, the Anglo-Dutch Wars, the Russo-Japanese Wars, the Russo-Turkish War (1768–1774), the English Civil War, the American Civil War, and the Sino-Japanese War). For visualization purposes, we group years into the closest five-year intervals to identify the estimates β_τ.

Event studies While our baseline approach in Equation (1) considers correlations without addressing potential dynamic effects, this appendix presents results from event studies that control for artist, time, and country fixed effects (as in Equation 1). To do so,

we focus on abrupt events with clearly defined start dates and analyze 62 event-country treatments associated with 24 distinct wars. The analysis includes the following countries or current political entities: France, Germany, Italy, the Netherlands, Spain, the United Kingdom, the United States, China, Japan, Portugal, Russia, Sweden, Switzerland, and Turkey.

We use this exercise as a proof of concept: can we detect a large negative shock through shifts in the emotions conveyed by artists before and after an event? Figure D14 reports the results for three main indices: *fear*, *contentment*, and *positivity*. We find that *fear* increases by approximately 2 percentage points, *contentment* moves in the opposite direction, and *positivity* declines by about 0.05. On average, the effect is long-lasting: emotion indices do not return to baseline levels even 20 years after the event—despite a median conflict duration of around 10 years in our sample.