Image(s) *

Hans-Joachim Voth † David Yanagizawa-Drott ‡

September 20, 2026

Abstract

From clothes and hairstyles to fashion accessories, humans use a wide range of style choices to express themselves, impress others, stand out, or signal group membership. Despite the importance of style for social interaction and self-expression, systematic evidence on style choice over time and space remains scarce. Using 54 million individual portraits of graduating US seniors from more than 360,000 yearbooks between 1930 and 2010, we introduce granular measures of culture based on the similarity of style choices among peers, and across generations. Image-based measures of norm deviation and style novelty display rapid cultural change starting in the mid-1960s. This shift generates convergence between norm deviations by male and female students. Regional patterns follow similar trajectories until the mid-1970s. Thereafter, polarization emerges: the American South reverts to more homogeneous behavior and traditional styles, whereas the West sustains high levels of norm deviation and a faster pace of stylistic innovation. These regional cleavages persist. Style-based norm deviations and style novelty predict cohort outcomes later in life. These include countercultural attitudes, differential marriage patterns, fertility choices, labor market behaviors, and patenting. Our results illustrate the usefulness of images as a source for cultural economics.

Keywords: Culture, norms, cultural change, similarity, innovation.

JEL Classification: A12, Z13, N34

*We gratefully acknowledge financial support from the Swiss National Science Foundation (grant 100018-208092). We thank Anjali Adukia, Elliott Ash, Ran Abramitzky, Marianne Bertrand, Leonardo Bursztyn, Emir Kamenica, Joel Mokyr, Alexander Todorov, Sendhil Mullainathan, Vicky Fouka, and Katia Zhuravskaja for helpful comments, seminar audiences at Chicago-Booth, BGSE, Harvard, the ETH-Zurich AI workshop, Mannheim, Stanford, Nottingham, BU, the 2024 UBS Advisory Board workshop in Zurich, and the NBER cultural economics meetings for feedback, and Emilia Bullano, Etienne Compérat, Olaf de Rohan Willner, Ada Gianassi, James Godlonton, Daniel Göttlich, Dominik Loibner, Ewan Rawcliffe, Simone Ricci, Jérôme Romano Blanc and Vanessa Sticher for outstanding research assistance. We are particularly grateful to Martin Fiszbein for sage advice.

†University of Zurich. Email: voth@econ.uzh.ch

‡University of Zurich. Email: david.yanagizawa-drott@econ.uzh.ch

1 Introduction

'Culture' is both "the way of life, especially the general customs and beliefs, of a particular group of people at a particular time" and "the attitudes, behaviour, opinions, etc. of a particular group of people within society", according to the Cambridge Dictionary. Empirical studies of culture have largely focused on the second aspect. Culture in the broader sense transcends attitudes and beliefs, and encompasses 'a way of life'. Recent work in economics has begun to analyze additional dimensions of culture, exploiting anthropological surveys, naming patterns, consumption baskets, and folklore traditions (Atkin 2016; Bazzi et al. 2020; Bertrand and Kamenica 2023; Michalopoulos and Xue 2021; Michalopoulos and Rauh 2024).

In this paper, we turn to a new source—photographic portraits. By analyzing photographs, we contribute to a growing literature in economics using images (Adukia et al. 2023; Ash et al. 2022; Caprini 2023; Ludwig and Mullainathan 2024; Shofia 2022). The use of portrait pictures for social science research originated in the 19th century. Galton (1878) infamously sought to identify the typical facial characteristics of criminals and the mentally ill by superimposing multiple portrait photographs. Here, instead of faces, we study people's *style choices*. Such choices are part of everyday culture: While many are unconscious, the clothes we wear, the haircut we sport, and the glasses we pick often reflect personal preferences. Style choices depend on social context and are linked to norms (Hancock et al. 2013):[1] Is it permissible to deviate from what others wear? Does a woman need to cover her hair? Can men wear wigs or tights? The idea of style choices as an indicator of status, prestige, and social change has a long lineage in the social sciences (Veblen 1899; Simmel 1957). Many religions closely circumscribe what is considered acceptable dress; sumptuary laws routinely determined which groups in society were allowed to wear what (Riello and Rublack 2019).

We study cultural change through the lens of style choices, as documented in portrait pictures. We use more than 54 million senior portrait pictures from approximately 363,000 yearbooks, taken of graduating students in thousands of US high schools during the period 1930 to 2010. To construct meaningful measures of culture from images, we take several steps.

First, images represent high-dimensional objects. We extract low-dimensional representations of the information they contain. Deep neural networks can identify key style char-

[1] Men's hair length has repeatedly been fought over in the courts (Graham 2004). The 1960s cult movie "Easy Rider" depicts the murder of two freedom-seeking young men at the hands of rural rednecks, for whom the heroes' hair is the main trigger for visceral hate.

acteristics, from hair length and hairstyle to the use of necklaces, the wearing of ties, or the depth of necklines. In this way, for each portrait, we obtain a sparse vector reflecting individual style choices. Each combination of individual features we define as a "style". Styles do not diffuse like color TVs or washing machines; they rarely follow an S-curve. While some new features like low necklines for females or earrings for men start from a low base, grow rapidly, and then stay high, others show a cyclical pattern of explosive growth, followed by sharp falls of adoption.[2]

Second, we use two measures to track aggregate trends: *norm deviation* and *style novelty*. Horizontal norm deviation captures (low) similarity of style choices in each class-gender-year cell, using 1-average cosine similarity of style vectors. It shows whether clothes, accessories, and hairstyles of students differ from those of their peers. High horizontal norm deviations, i.e. an absence of conformity, can reflect differences of underlying preferences, or a low level of "norm tightness" (Gelfand et al. 2011). Horizontal norm deviations are known to vary markedly in the cross-section of individuals, across cultures, and over time (Bazzi et al. 2020; Hofstede 2011; MacFarlane 1978). In contrast, vertical norm deviation captures how different style choices in a given cohort are from those of seniors at the same high school a generation earlier, by gender. Persistent local culture implies high similarity over time: senior-year students in the year in question look similar, on average, to those graduating twenty years earlier. Under rapid cultural change, vertical norm deviation would be high. Since our image data start in 1930, we begin measuring vertical norm deviation in 1950. Both horizontal and vertical norm deviation reflect how much individuals deviate from others in the same location. We take the average of the two and label it *norm deviation*. This variable starts in 1950 and is available until 2010.[3]

Third, we measure the arrival of new styles. Individuals can deviate not only from previous generations in the same locality (vertical norm deviation) or from their peers (horizontal norm deviation). They can engage in new behaviors that were previously not observed in the same high school, state, or country. *Style novelty* captures style combinations that either did not exist before, or were very rare. For each student's style, we examine the full history in the country as a whole and ask how rare it has been and is (up to and including

[2]We can also show that adoption of styles appears to occur in school-specific "waves". We compare the actual prevalence of common full-style combinations in graduating high school cohorts with a binomial benchmark based on independent draws from the gender-specific national style distribution. Relative to this benchmark, adoption has excess mass at both low and high take-up rates. Hence, while not everyone switches to a new style, the spread of style choices clusters heavily at the high school-cohort level.

[3]We also show in the online appendix that the two variables are highly correlated, although far from perfectly collinear.

the present year). If a student's style combination falls in the bottom one percent of the cumulative style-frequency distribution for that gender-year, we call it novel. This metric is distinct from horizontal and vertical norm deviations. Someone may show high horizontal norm deviation (wearing a set of clothes, accessories, and a hairstyle different from their classmates) and high vertical norm deviation (different from their parent's generation in the same high school), but not wear a novel style because it had been adopted at scale in other high schools previously. We aggregate our data at the commuting zone level.[4]

Several new facts emerge: Until the middle of the 1960s, culture as reflected in high school portraits was remarkably stable over time: norm deviations for men were low. For women, norm deviations were substantially higher. Both the horizontal and vertical measures show the same pattern. Style novelty was rare for both genders. Style choices were relatively homogeneous across the country. Males overwhelmingly had short hair and wore dark suits with ties, shirts with collars, and sported no facial hair. Compared with their parents' generation, young men looked almost indistinguishable. Among females, we document greater variation, both over time and across space. Style novelty was more common among females and stable.[5]

Starting around 1965, rapid cultural change became visible. Men increasingly deviated from the looks of their parents. Importantly, there was not a uniform shift in styles, which would just signal a new manifestation of conformist behavior. Instead, similarity among peers declines precipitously, and new styles begin to arrive at greater frequency. These trends accelerated further in the late 1960s, following the 'Summer of Love' and the Woodstock music festival. Male norm deviation rose, and style novelty surged.

By the mid-1970s, counterculture peaked in the form of higher norm deviation and rapid style innovation. Over the same period, gender differences narrowed: women started from higher levels of norm deviation, but these measures declined over time, while female style novelty increased only gradually. By the 1990s, we observe convergence across genders: men and women showed near-identical levels of norm deviation, although differences in style novelty persisted. During the second half of the 20th century, as equality between the sexes became the norm before the law and in popular culture, the heterogeneity of style choices converged.

[4]We also construct a state-level panel dataset, for merging with other sources of data unavailable at more granular levels.

[5]Given that we use a fixed vector of style features, the universe of potential novel styles mechanically shrinks over time, by construction. However, there are many novel combinatorial possibilities left in the data, several orders of magnitude more than the styles actually observed in the data.

Aggregate trends mask a dramatic spatial rupture. Before the mid-1960s, the country was stylistically homogeneous: a graduating senior in Alabama was as likely to innovate or deviate from the style of his peers as one in New York or Texas. Since the counterculture peak of the mid-1970s, that uniformity has given way to persistent polarization. After a peak in innovation and norm deviations, trends in the American South reversed. Vertical norm deviations declined. In contrast, western and northern parts of the U.S. drove innovation in styles while showing high levels of norm deviation. Standard socioeconomic correlates are not the main driving force behind this fracture. Across the urban-rural divide, the Democratic vote margin, and income, levels differ little and trends move in step. footnoteOur evidence is broadly consistent but not highly correlated with documented political polarization (Gentzkow et al. 2019). Major shocks leave equally faint traces: neither the "China Shock" nor Vietnam War casualties bend style choices in any clear direction.

How meaningful are these style-based measures of culture? Do they correlate with or predict other, more standard, measures? We leverage data from the General Social Survey (GSS), available from 1972 and onwards. A key limitation is its size, with few observations per state-year. We aggregate our data to the state-by-cohort level, and match survey data of respondents aged 18-25. Norm-deviation and style-novelty measures correlate with attitudes related to counter-culture, like pre-marital sex, homosexuality, drugs, and church attendance. Other attitudes reflected in the GSS, like those towards women's rights and anti-authoritarianism, show weaker, null, or opposite patterns. Pictures and the GSS agree where one would expect a clear relationship—the 1968 generation emphasized rule-breaking self-expression in many dimensions, including personal style. The mixed pattern suggests that portraits capture a specific dimension of countercultural visual style rather than a broad index of all progressive or anti-authoritarian attitudes. Null-results may also reflect a lack of granularity and coverage in the GSS.

We then turn to later-life behaviors and choices. Do they differ systematically according to the style choices made by graduating high school seniors? We collect data on family form from the Current Population Survey (CPS), available from 1962, studying individual-level outcomes at ages 25-40. Young men and women from cohorts that defied their parents' style and were more likely to stand out in front of their peers – showing high norm deviation – chose less traditional family and work patterns: marriage rates are lower, women are more likely to work full-time, men are less likely to work full-time, and a composite index of traditional lifestyle choices declines.

Style innovation, norm deviations, and technological innovation are also closely linked

in our data. Cohorts from commuting zones with more norm deviation and more style innovators file significantly more patent applications and receive more grants—a pattern that survives commuting zone and cohort fixed effects as well as socioeconomic controls. A one-standard-deviation increase in norm deviations is associated with approximately a 0.14 standard deviation rise in patent applications and a 0.085 standard deviation rise in grantees; for style novelty, the associated increases are 0.04 and 0.03 std. This is not a causal claim. Instead, local environments tolerant of style originality are more likely to be ones novel ideas are more likely—cultural openness cuts across multiple domains simultaneously.

Finally, we examine naming patterns. They have been used extensively in the literature to measure individualism (Bazzi et al. 2020). How does this indicator of parental choice compare with style expression—and in particular, horizontal norm deviations—in our senior portraits? We link the names and pictures of *individual* high school seniors, finding a positive correlation between the two. At the same time, the correlation is low, and the share of variation explained in a simple bivariate regression is never more than 5%, for either males or females. This suggests that names capture a related but distinct dimension of cultural expression.

We make several contributions to the existing literature. First, we build on the rich literature on cultural persistence and change. When fundamental economic and environmental factors change, norms of behavior can shift and increase vertical norm deviation (Giuliano and Nunn 2021). Beyond material factors, the literature has pointed to additional mechanisms of identity, social image incentives and beliefs about what is socially acceptable. Notable examples include Kuran (1987, 1989) and Benabou and Tirole (2011), who point to the existence of potentially fragile equilibria when people care about the actions or attitudes of others. These can be disrupted when a few individuals, or new information, initiate a wave of change. Experimental evidence has been able to pinpoint the key role of perceptions of what is socially acceptable, including what is deemed as "cool" in educational settings (Bursztyn et al. 2019), and how local equilibria may be particularly fragile when those perceptions are misaligned (Bursztyn et al. 2020, 2023). More broadly, various interdisciplinary perspectives exist on the dynamics of norms (Gelfand et al. 2024), including models characterizing how local communities can reach different stable, bi-modal, equilibria as they cross certain tipping points (e.g., Gavrilets (2020)). We contribute to this literature by studying style choices as a cultural phenomenon. This approach generates granular measures of changes in local outcomes. Their availability enables us to analyze how local fundamental economic factors such as income and rural-urban cleavages correlate with cultural change.

Our polarization result echoes this idea of tipping points: an initially uni-modal distribution across localities passes through a phase of rapid cultural change and settles into a stable bimodal one. Our paper demonstrates the potential of using portrait pictures to speak to key questions about culture, innovation and economic growth (Gorodnichenko and Roland 2011).

Our paper also relates to research on the information content of images in the social sciences, and on style choice and fashion as economic and social phenomena. Theoretical analysis of fashion goes back to the classic work by Veblen (1899) and Leibenstein (1950) who differentiated between "snob" and "bandwagon" effects, with adoption becoming either more or less likely as a function of others' adoption decision and type. Becker and Murphy (1993) derive a micro-founded model of such consumption externalities. Matsuyama (1993) explicitly models groups of conformists and non-conformists and shows that demand for goods can fluctuate cyclically ("fashion cycles"), and Pesendorfer (1995) models firms' introduction of new fashion items. Imitation of first movers in consumption is also prominent in Banerjee (1992), while Karni and Schmeidler (1990) present a model of social stratification in consumption. While several theoretical papers in economics analyze the diffusion of fashion as a social phenomenon, there is little empirical work on the origins, spread and similarities of style and fashion over time and space. The yearbook-image data help fill this gap by providing high-frequency, geographically granular measures of style choices over several decades.[6]

We introduce several types of analysis, using portrait pictures to examine style choice as an indicator of cultural change, evaluating local culture with high spatial and temporal resolution. In this way, we also relate to a recent literature in economics using image analysis, with researchers evaluating the effect of facial features on judges' bail decisions or on lending behavior (Athey et al. 2022; Ludwig and Mullainathan 2024), of faces and perceived trustworthiness (Todorov et al. 2009), or the distribution of protagonists' skin color in children's books (Adukia et al. 2023).

The paper is structured as follows. In Section 2, we provide a brief background on the history of portraits and photography, as well as some broad patterns of cultural change in the post-war United States. In Section 3, we present the data sources and in Section 4 we explain our methodological approach. Section 5 contains the main analysis. Section 6

[6] A related strand of literature in computer science has documented fashion trends in recent years. Hidayati et al. (2014) analyze fashion trends during New York's fashion week, based on the couture shown. Along similar lines, Kiapour et al. (2014) create five main clothing style categories from human coding and roll out a categorization scheme based on models trained on these classifications.

concludes.

2 History and Background

In this section, we provide some brief historical background of portraits and pictures, as well as of the period on which we focus: the American post-war years.

2.1 Portraits and Photography

Clothing and jewelry are as old as mankind. Early pre-historic art depicts skirts and animal skins, dresses, and different hairstyles (Bigelow and Kushino 1979). Accessories have been unearthed by archaeologists for periods as far back as the 16th century BC (Nosch et al. 2014). While styles came and went in many locations and periods, rapid changes in preferred or acceptable clothing and hairstyle—fashion—is probably a relatively recent phenomenon. Braudel (1975) famously argued that rapid changes in dress originated in Europe, among the upper classes during the late Middle Ages, as a way to distinguish themselves from the lower orders. This view is controversial and undoubtedly Euro-centric. Fashion, and periodic—if not necessarily rapid—changes in dress are probably human universals (Welters and Lillethun 2018).

Industrialization coincided with the spread of textile manufacturing, especially of cotton (Crafts 1985). As productivity surged and the cost of new cloth fell, fashion items became more widely accessible. Some historians have located a "consumer revolution" in 18th century English society, centered on new fashions. McKendrick et al. (1982) argue that even servants could afford several fashion accessories every year, making it easier to follow new trends—and creating a greater need among the upper classes to distinguish themselves.

Dress and style choices have served as a way for individuals to present themselves to others and posterity since antiquity. Kings and emperors had their faces depicted on coins and on marble statues. Popes, kings, and officials down to early modern burghers commissioned portrait paintings, presenting themselves to the world as warriors or in the stark simplicity of black robes, in front of their worldly possessions or with family, friends, and favorite pets. Many famous artists painted self-portraits (Carbon 2017), from Dürer to Picasso, depicting themselves in every style from darkly realistic images to idealized versions (Beyer and Lindberg 2003). The arrival of photography changed the extent to which images could be embellished. At the same time, it created new scope for highlighting one's preferences and

horizontal norm deviation, from the style of photograph chosen to the manner in which one dressed and presented oneself to the world. The very first portrait photographs date back to 1839; by the 1840s, daguerreotypes had become common. From the 1930s, roll film allowed a quantum leap in the mobile use of cameras, and brought costs down; soon, family outings and celebrations were not complete without an—often staged—picture commemorating the occasion (Prodger 2021).

High-school yearbooks became popular in the US from the 1930s; by the 1940s, many high schools compiled annual overviews depicting every student, ordered by class. The yearbooks often describe sporting events and teams, depict high school clubs and associations, portray teachers, and show ads by local firms.[7]

2.2 Post-War America and Counterculture

As the country emerged triumphant from World War II, American culture exerted a strong influence around the world. Hollywood, American TV shows, American universities and music combined into a powerful and seductive form of "soft power". Youth rebellion against established norms became a dominant and recurring theme in fashion and an important form of self-stylization.

Growing economic prosperity and rapid demographic expansion were accompanied by a cultural revolution, particularly among young people, who began to challenge traditional social norms and values. In the 1950s, teenagers began to embrace rock n' roll music and a more rebellious teenage culture. This led to the rise of a youth counterculture in the 1960s, which was marked by a rejection of traditional values and the embrace of a more liberal lifestyle. The youth rebellion of the 1950s and 1960s had a profound impact on American culture. It created frictions in civic society, and between old and young, progressives and conservatives. While seeking to build a more tolerant and diverse society, it also furthered the rise of consumerism (Heath and Potter 2004).

The "counterculture" of the late 1960s and early 1970s centered around three main themes: opposition to the Vietnam War, rejection of traditional social and sexual mores, and the use of psychedelic drugs (Issitt 2009). While in some ways similar to the earlier Beatniks, the counterculture of the sixties is a distinct cultural phenomenon. Hippies and anarchists, or the Hells Angels, made their rejection of traditional society clear in many dimensions, but their physical appearance(s) was often what shocked older observers the most.

[7] Yearbook pictures are also explored by (Ginosar et al. 2015), but their exercise is limited by sample size – 30,000 images – and a focus on facial expressions.

Men would wear their hair long; facial hair made a comeback; many hippies cultivated a deliberately casual look, some even refusing to wear shoes. Hand-printed shirts and skirts in psychedelic colors were common, as were long flowing dresses for women.

Hippie style and culture largely originated with middle-class youth. They diffused widely in society, possibly because its torch-bearers were ethnically, culturally and in terms of social status, close to the mainstream (Davis 2013). By the late 1970s, many stylistic elements of the counter-culture had become "normal" (Kopkind 1979). To provoke required something new, like the mohawk hairstyle, combined with leather and spike accessories and piercings of punks.

New waves of youth culture emerged. Some subcultures and styles—from street style to ghetto chic, and from goths to punks and 1980s power dressers—reached wider audiences through the popular cable television channel Music Television (MTV). Our data, covering yearbooks until 2010, document style changes and innovation within each graduating cohort and location, over time.

3 Data

American high school yearbooks have a long lineage. From the early 20th century onward, student associations began to publish annual yearbooks containing a range of information on clubs and societies, events and sporting competitions. Initially focused on collecting memorable utterances of seniors for the enlightenment of juniors, these quickly evolved into a collection of student portraits. By the 1930s, a high share of American children attended high schools. Many schools published yearbooks containing portraits.

Figure A.1 shows examples of yearbook pages from different schools, states, and years. Most images are relatively small, and portray only the head and upper torso. Black-and-white pictures give way to color from the late fifties onwards. Most pictures are frontal or ¾ frontal portraits; pictures in profile are rare.

While a range of different sources exists, the commercial website *www.classmates.com* has by far the most comprehensive collection. Yearbooks are widely available from the 1930s onwards, and the raw collection is most extensive for the 1980s; coverage in our retained senior sample, however, peaks around 1970 (Figure A.2). The company acquires physical yearbooks and digitizes them, but its selection process is not disclosed. The data are unlikely to be fully representative of the population at large, whether at the national, state, or commuting zone level (for details on construction, see Appendix B).

Table B.3 breaks down yearbooks collected by state in our sample, from year 1930 to 2010. In total, the final data contains approximately 363,000 yearbooks. Figure A.2 shows the number of high schools covered by the data as well as the total number of available images over time. We use some 54 million images in this national dataset. The sample is not randomly drawn, but our data provide broad coverage across states, commuting zones, high schools, and cohorts.[8] In each yearbook retained for analysis, we use the entire detected set of senior portraits.

Figure B.1 lays out the full pipeline. A portrait-recognition algorithm scans each page for faces and rectangular frames; a symbolic rule then decides where the senior section begins. We keep only runs of consecutive pages whose images share a uniform size and color profile and—because senior portraits are more often in color and larger than the rest—use size and color to single them out. Human audits confirm our pipeline classifies senior pages with about 99% accuracy, recovering 95.4% of true senior pages. Downstream denominators differ by measure.[9] A hand-labeled subset of these images is then used to train classifiers (Google Vertex AI) for various style attributes such as hair length, clothing, and accessories, with each classifier predicting the likelihood of each attribute through a continuous score; the trained models are then applied to the full corpus. Gender is classified separately (subsection B.5). Model predictions from Google Vertex AI are discretized to become dummy variables using a threshold of 0.3 for a prediction to be categorized as 1. High school locations in the dataset are geocoded via OpenStreetMaps API tool where possible, with the remaining share manually coded. Appendix B offers an in depth explanation on data construction.

The main dataset is at the image-level for more than 54 million senior portraits. The key variables are style attributes (model predictions), horizontal norm deviation, vertical norm deviation and style novelty for each individual. We describe how these are constructed below.

We also construct commuting zone by cohort and state by cohort datasets. These facilitate matching style measures from image data with auxiliary datasets on patenting behavior and innovation from Bell et al. (2019), family formation and work-related outcomes from

[8]The correlation coefficient between the (log) number of portraits we use per commuting-zone-year cell and the (log) population is 0.817.

[9]Horizontal norm deviation requires same-gender classmates in the school-year cell. Vertical norm deviation additionally requires a comparable same-school same-gender cohort twenty years earlier, so it is missing when that lagged cohort is unavailable. Style novelty is defined only for eligible full-style combinations. Linked-outcome analyses then impose further sample restrictions from the external data: the GSS and CPS matches are available at state-by-cohort cells, while the patenting analysis is limited to the cohorts and commuting zones covered by the Bell et al. (2019) patenting data and the required socioeconomic controls.

the CPS, and attitude and belief data from the GSS. Table 1 shows summary statistics for these two panels, related to style attributes and style measures (horizontal norm deviation, vertical norm deviation, style novelty), and reports the effective number of observations for each measure. Table A.1 reports summary statistics for the three auxiliary datasets used in the analysis [10] —the commuting-zone patenting panel from Bell et al. (2019), the CPS, and the GSS—each with its outcomes, standardized style measures, and shared controls.[11]

Note that few of the new styles diffuse in an S-curve style pattern. Figure A.5 looks at nine novel styles that diffused nationally. Bowl-cuts for women, long hair for men, and Afros for both genders grew at an accelerating pace first, but then collapsed. Low necklines for women, spiky hair for men, earrings for men, and mustaches spread at a steady rate before stabilizing, in some case, at high levels.

4 Method

We present simple methods to compare style choices over time and space, and demonstrate their reliability, using a human audit.

4.1 Object Identification and Image Vision

While the human eye processes images rapidly (Ginosar et al. 2015), we require a systematic approach to evaluate image similarity (Wang et al. 2004). While easily implemented, it is affected by changing technology (color film, lens characteristics, printing technology) and photographic style (depth of field, contour or frontal lighting, etc.). The second approach is to create a vector of attributes: clothing, accessory, or hair characteristics of each image. We focus on the analysis of such style vectors. These capture combinations of identifiable and easy-to-interpret features, reflecting *style choices*. Even sparse vectors can differentiate between thousands of different "styles" (i.e., combinations of style elements). One limitation

[10] Beyond these three, we match our style measures to several further datasets for the correlational and robustness exercises in Section 5 and the appendices—for example, income-mobility, county-level presidential-vote, and economic- and political-shock data—each described where it is first used.

[11] As control variables in regression and correlation analyses we make use of *Historical, Demographic, Economic, and Social Data*, Inter-university Consortium for Political and Social Research (ICPSR) Study 2896. These data contain socioeconomic variables recorded every ten years at the county level, which we aggregate to commuting-zone averages for each cohort. We then linearly interpolate annual values between adjacent decennial census observations. The commuting-zone-by-year dataset of our main measures (horizontal norm deviation, vertical norm deviation, and style novelty) is freely available at https://huggingbeard.github.io/images/.

of our source is that portraits only depict seniors' heads and upper torsos; fashion change below the waistline (changes from long skirts to miniskirts, say) is not captured.[12]

To populate the vector, we need to identify objects. We first convert all images to grayscale to avoid confusion from the diffusion of color pictures from the 1960s onwards. To identify individual style elements in our pictures, we define a vector of 25 style elements, ranging from necklines and the wearing of ties to jewelry, hairstyle, and the type of collar worn, for both men and women. For style classification, we asked 155 Amazon Mechanical Turk workers to make 18,878 image-category decisions, using a set of 6,540 randomly chosen images from our high school corpus for style classification.[13] For the hair style classifier, we used 2,127 images.

We then train convolutional neural networks, using Google Vertex AI, to extract style features, using two multi-label classifiers.[14] We evaluate the predictive performance of our algorithms. For gender classification, we adopt a two-stage process combining two off-the-shelf models (subsection B.5 gives the details). After training the classifiers for up to 16 hours each on Vertex, we obtain a classification accuracy for the 656 images in our validation sample of 93.4% to 99.3%. We then hand-coded an additional dataset (not part of the training and validation data), asking Prolific workers to classify style characteristics of up to 2,500 additional images. In subsection D.1 we discuss the details of this out-of-sample exercise. Table D.1 gives an overview of the training and prediction results. We obtain an accuracy of 98.5% for mustache identification, marking the best result in our dataset; at the opposite end of the spectrum, our algorithm agrees 67.4% of the time with the (median) human coding of curly hair. Humans agree somewhat more with classification decisions by other humans than they agree with the algorithm. Across categories, we average 79% (clothes) to 98% (facial hair) out-of-sample accuracy. Table C.2 reports AUC and other performance metrics of our classifier.

4.2 Basic patterns

Figure A.3 gives an overview of style attributes in our data over time. Men with ties are common in the 1930s, and their share stays above 70% of our sample into the late 1960s. It then falls precipitously, declining to roughly a quarter of the sample by 2010. The share of

[12] Note that this may affect the assessment of female fashion trends more than that of males.

[13] Scoring was carried out in 2022, before the widespread availability of VLMs.

[14] In addition, we use a single-label algorithm for the detection of long hair (which we run separately from the hair style classification, as styles are mutually exclusive but long versus short hair can come in many different styles). Further details of our procedure are in Appendix C.

men with suits follows a similar downward trend. Long hair is exceedingly rare throughout the 1950s and 1960s, and then spikes for a short time in the early 1970s, before falling sharply and then partially rebounding in later decades. Clean-shaven portraits are dominant in the 1950s and 60s, and their share falls in the 1970s, before bouncing back from the 1980s onwards. The style attributes for women also show significant changes over time, but long-term trends are less pronounced. Long hair is en vogue twice in our sample, the 1940s and early 50s, and then again from the 1970s onwards; jewelry is rare in the 1940s and then becomes common from the 50s onwards.

Styles are not independent of each other. Figure A.4 shows the correlation pattern of style attributes. Suits are positively correlated with ties and shirts with a collar, as one would expect for the case of men. Long hair is only weakly (and negatively) correlated with indicators of formal wear, such as ties and suits. For women, correlations are generally lower, but jewelry, such as necklaces and earrings, correlate with low necklines, and long hair is also correlated with necklaces.

Next, we classify unique combinations of attributes as a 'style' and examine their evolution over time. In Figure 2, we show the composition of styles in our sample for men and women. In the 1940s, around 85% of men were in the same full-style combination—suit, tie, shirt with a collar, no glasses, short hair, no jewelry, and clean-shaven (style "N"). By the 1960s, the broader share wearing both a suit and tie had declined to around 75%; another 8-10% had short hair, suit, bow tie, collared shirt, and no facial hair (style "Q"). Until the late 1960s, four styles for men alone account for nearly 90% of the sample. By the early seventies, the influence of the more traditional styles ("N" and "Q") falls sharply, from a share of around 80% in 1966 to less than 40% by the mid-seventies. Suit and tie, combined with long hair, enjoys a brief return to popularity in the mid-1970s, with a share of more than 10%.

It is striking that the residual "other" category is one of the most common types in our sample by the 1970s already, and by far the single largest category by 2010—a sharp change from the 1940s and 1950s, when 5% or less of high school seniors were in this category. In other words, fewer students fit the leading named style combinations over time. The share outside the top categories rose from a small minority of each class to one of the largest groups, a pattern consistent with increasing style fragmentation.

Women's styles tell a different story. Already in the 1930s no single combination dominates: the residual "other" category accounts for more than half of the total and—fluctuations aside—shows no clear trend across the following decades. The one clear movement is the

rise of low necklines from the 1980s onward (highlighted in blue in Figure 2).

4.3 Measures

We calculate three statistics to measure the evolution of style change. Two of these are direct measures of norm deviations—within each class of high school students (horizontal), and across generations (vertical), comparing the current cohort with graduates 20 years earlier. The third looks at the upper tail of style innovation.

4.3.1 Horizontal Norm Deviations

We calculate *horizontal norm deviations* in a school s, of gender g, in year y, as:

$$h_i = 1 - \frac{1}{|C_i|} \sum_{j \in C_i} \frac{v_i \cdot v_j}{\|v_i\| \|v_j\|},$$
$$H_{s,g,y} = \frac{1}{N_{s,g,y}} \sum_{i \in I_{s,g,y}} h_i, \tag{1}$$

where v_i is the style vector for individual i; $I_{s,g,y}$ is the set of individuals of gender g in school s and year y; $N_{s,g,y} = |I_{s,g,y}|$; and, for each $i \in I_{s,g,y}$, $C_i = I_{s,g,y} \setminus \{i\}$ is the set of all other same-gender individuals in that school-year cohort. The first line defines individual i's horizontal norm deviation as one minus the average cosine similarity to same-gender classmates. The second line averages these individual scores to form the school-year-gender horizontal norm deviation index. The statistic is defined only when $N_{s,g,y} \geq 2$; singleton same-gender cells are treated as missing for this measure.

Figure 3 illustrates the approach using senior portraits from Preston High School in Idaho. We compare each image (in the illustrated case, the first one) with all others in the same high school senior cohort of the same gender. We first average the focal image's style vector cosine similarity with all others in this cohort and subtract that average from one to obtain the focal image's horizontal norm deviation score. We then repeat this calculation for all images and average the individual scores. Cosine similarity takes value 1 if two vectors are identical, and value 0 if they are orthogonal. In the displayed 1964 cohort, the resulting horizontal norm deviation is 0.03, indicating that style choices (tie, suit, hair length) are mostly identical among men in the same cohort.

4.3.2 Vertical Norm Deviations

We define *vertical norm deviations* analogously, as one minus the average cosine similarity, but now use the set of style vectors of same-gender individuals from the same institution from twenty years earlier as the comparison group. The vertical measure is therefore defined only when the focal cohort and earlier same-gender comparison cohort are observed; otherwise the corresponding comparison is missing. We use twenty years as a rough measure of changes across generations but will show results where we iterate across many time horizons, from one year all the way to twenty years. Aggregation to the high school cohort level proceeds as before. As we do not have a balanced panel of high school yearbooks, some missing observations are present in our data.

Figure 3 illustrates the vertical norm deviation calculation using male graduates from Preston High School in Idaho in 1964 and 1984. The 1964 portraits are visually homogeneous, whereas the 1984 portraits show greater within-cohort stylistic dispersion. Comparing the 1984 style vectors with those from the 1964 cohort yields an average vertical norm deviation score of 0.54, indicating substantial change in male style choices across generations within the same high school.

4.3.3 Style Novelty

To calculate our third measure, style novelty, we compare the style choices of each high school senior with those observed among graduates of the same gender who came before them plus the current year (and not just those in the same high school). We want to know how rare the style chosen by an individual is, compared to previous and current cohorts' style choices. We first restrict attention to full-style combinations observed at least 30 times in the full sample. For each remaining style, we then calculate the cumulative count of high school senior images that use the same full-style combination.[15] Formally, the cumulative style count of each style is defined as:

$$C_{s,g,t}(t_0, t) = \sum_{y=t_0}^{t} T_{s,g,y}, \tag{2}$$

where $T_{s,g,y}$ is the number of portraits of gender g in year y assigned to full-style com-

[15]The arrival of new styles can be viewed through the lens of the familiar experimentation vs. exploitation framework (Murdock et al. 2017).

bination s. We begin in 1930 (t_0). Let $\mathcal{S}_{g,t}^{30}$ denote the set of eligible style combinations that are observed for gender g in year t. We then determine whether a style belongs to the rare-style tail:

$$I_{s,g,t} = \begin{cases} 1 & \text{if } s \in R_{g,t} \\ 0 & \text{otherwise} \end{cases}, \tag{3}$$

where $R_{g,t}$ is constructed by sorting the eligible same-gender styles in year t from lowest to highest cumulative count and selecting styles until their cumulative counts reach 1% of all eligible portraits of gender g observed through year t. We include all style combinations tied at the cutoff, so the realized image-level share classified as novel can differ from exactly one percent. Once a style drops out of the rare-style tail, we allow it to re-enter only after ten years without being classified as novel. Figure 4 illustrates our method. The left-most image is one of the first observed examples of the displayed style combination in our data. The same style appears again in later years and remains in the rare-style tail in the middle image; by the final image, after broader diffusion, it no longer qualifies as novel.

Not all choices that affect senior portraits are necessarily made at the level of the individual. Consider Figure 3. Some style choices in senior portraits may not be independent within a cohort, and could even reflect school policy. However, the continuation or collapse of a local social norm of this kind is exactly what we are interested in.[16] We examine reappearing trends, and the effect of style innovation measured only locally, in the appendix (Figure E.12 and Figure E.13)

4.4 Human Audit

We use a sparse vector of image characteristics to robustly capture the style of high school seniors in our sample. The vector should be long enough to span the overall set of possible forms of stylistic expression; it should be sparse enough to be robust and allow for reliable identification of style attributes. To test the validity of our approach, we use human audit samples. How important are the features we use, in the eyes of people? To this end, we recruited a total of 60 workers on Prolific who examined 4,980 pictures.

Participants were presented with three images—a reference image, and two comparison

[16] subsubsection E.5.1 shows that using or not using data from schools with a de facto "uniform" does not affect our aggregate results.

images. We then asked respondents to rate the similarity of the comparison images with the reference image. We randomly selected reference images and comparison images (with the same gender and the same decade as the reference image). A simple but powerful test of the validity of our approach is to analyse how often the human agrees with the prediction of the closer match, based on cosine similarity scores.

Humans looking at images react instinctively and quickly to many attributes including facial expressions, emotions, skin color, facial symmetry, and other markers of attractiveness, in addition to style choices. While we prompted respondents to focus on style,[17] all other factors are known to affect behavioral responses. Despite these distractions, the likelihood that humans agree with the algorithm, based on our style vector, is twice as high as disagreement.

Figure D.2 illustrates agreement between humans, and between humans and the algorithm. Remarkably, humans disagree with each other almost as much as they disagree with the algorithm. Agreement rates are not constant across images, and we summarize them two ways. First, against the modal human choice: humans agree with that consensus 73% of the time, the algorithm 66%—a gap of just seven percentage points (Table D.3). Second, comparing humans one-on-one: pairwise inter-human agreement ranges from 50 to 96.6% across images, and human-machine agreement is similarly variable, so the algorithm sits squarely inside the band of human disagreement.

[17] *"In this survey, we will present you with various images extracted from yearbooks. Your task will be to compare a reference image with two other images and select the one that is most similar to the reference image. It is important that you focus on style characteristics rather than facial features."*

5 Trends, Validation, and Applications

In this section, we apply our method to our new source, measuring style choices. We demonstrate the potential of our approach by mapping the geography, speed, and scale of cultural change. We also show that our measure of cultural change, style choice, is associated with consequential outcomes over the life cycle.

5.1 Patterns of Cultural Norm Deviation and Change

We first present broad patterns of change over time and space, using the dataset of high schools over the period 1930-2010. We examine how much each student differs from other individuals in his or her class ("horizontal norm deviation"). We also compare each student in year t with people graduating from the same high school 20 years earlier ("vertical norm deviation"). Figure 5 plots the norm-deviation scores for both types of analysis over time (Panel A), as well as our measure of style novelty (Panel B). Appendix Figure A.9 reports the correlations among these style measures.

Women start out with high levels of horizontal norm deviation in the 1950s, while men register very low values. Within each high school, the style choices of men were near-identical, while those of women varied. Rapid changes do not necessarily occur simultaneously for males and females. Instead we see waves of change occurring intermittently. There is no clear evidence that one gender is leading and the other is lagging. Over the second half of the 20th century, horizontal norm deviation scores converge. Men start to show much more horizontal norm deviation from the mid-sixties onwards, with further gains in the nineties, whereas the female score declines gradually. By the late 2000s, men and women are at near-identical levels.

These developments are mirrored in terms of vertical norm deviation trends. Males start out at very low levels—each generation looked largely the same as the one before, almost without variation, as indicated by cosine similarities close to 1. In other words, it is only when the first cohorts affected by the '1968' movement are compared with earlier cohorts that we find evidence of renewed, rising vertical norm deviation in style choices. In contrast, women showed relatively high vertical norm deviation for much of the post-war era. From the late 1960s and 1970s, male values rise sharply while female values decline, and the two series converge by the 1990s. By 2010, male vertical norm deviation is higher than female vertical norm deviation, indicating a greater rate of style change from one generation to the next.

Style novelty starts out lower for men than for women in the 1950s and 60s, but then surges in the 1970s. After a dip in the 1980s, it recovers and rises, with big swings around a rising trend. Women start off with higher innovation rates, but these remain steady for most of our sample period.

In Figure A.7, we examine vertical norm deviation over time more closely. Instead of the (arbitrary) 20-year horizon used in Figure 5, we plot average cosine dissimilarity at different horizons. Darker colors show greater vertical norm deviation; lighter colors indicate lower vertical norm deviation. On average, longer horizons are associated with higher vertical norm deviation. Until the counterculture era, vertical norm deviation remained low across horizons, indicating that the style of seniors graduating from high school barely changed over time. Moreover, our data allows us to examine shorter time horizons such as ten years, where we see vertical norm deviation rise around 1970 and the alternative horizons show that the timing of this break shifts with the comparison window. More broadly, there are visible diagonal patterns in the heatmap as we move 'northeast' in the figure, tracking longer time horizons and later cohorts simultaneously. For example, the 12-year horizon value is high around 1982 and the 15-year horizon value is high around 1985, both comparisons to cohorts around 1970. These comparisons show that the counterculture-era break remains visible across alternative horizon definitions. Thus, our measure of vertical norm deviation can capture fashion cycles as they come and go; a topic studied theoretically in the literature (Matsuyama 1993; Pesendorfer 1995) but for which appropriate data systematically measured across time and space has so far not been available.

What about variability over space? Figures 6 and 7 show our style measures by decade and by U.S. Census region. The 1950s and 1960s are characterized by relative uniformity across space in terms of norm deviation and style novelty. From the early 1970s, spatial variation increases. In the regional time series, norm deviation rises in all regions around the counterculture period, but the South peaks in the 1970s and then falls back relative to the Midwest and West. Style novelty also rises sharply in the 1970s, with a pronounced Southern spike, before the Midwest and West sustain the highest rates from the late 1980s onward.

These regional patterns are robust to sample composition and school dress codes: restricting to long-run schools (observed in at least 30 graduation years), excluding likely-uniform school-years (the modal apparel combination exceeding 80 percent of students in both the male and female cohorts), and recomputing the measures on the small and single-gender cohorts excluded by our minimum-size rule all leave the post-1970 regional ordering and

timing essentially unchanged (Appendix Figures E.9, E.10, and E.11).

In Figure A.8, we examine distributions by year and commuting zone, exploiting the granularity of our data. The relative uniformity of the 1950s and early 1960s is reflected in most commuting zones having very similar levels overall—there were few differences between the norm-deviation scores across high school students, and compared with earlier decades. By the late 1960s, the sharp ridges of high conformity and low vertical norm deviation start to spread out, with some high schools changing much more than others. By the 1990s there are clear signs of an emerging bimodal distribution, with a set of high and a set of low horizontal norm deviation areas, mirrored by high- and low-vertical-norm-deviation areas. The bimodality then appears stable. This polarization result, where a stable bimodal distribution across localities emerges, after initial unimodality followed by rapid cultural change, appears broadly consistent with Gavrilets (2020). The uniformly low style novelty gives way to a wide pattern of dispersion, with some schools remaining at the 1950s levels, while others score high, but there is no evidence of bimodality in this dimension.

Figure A.8 splits the data by gender. Again, we see substantial heterogeneity. The bimodality in horizontal and vertical norm deviations appears to be driven primarily by males. The dispersion in style novelty across the country in recent decades is also driven by males. We refrain from speculating about the determinants that give rise to these gender differences, beyond noting the fact that they are accompanied by changes in society, such as greater female labor force participation, reduced gaps in educational outcomes and beyond.

A caveat concerns the historical portability of our measures. Style variation and style novelty quantify the statistical rarity of an individual's appearance relative to classroom peers. The measures are therefore norm-relative by construction: they identify departures from the locally prevailing standard, whatever that standard is. What the measures cannot hold fixed is the *meaning* of a given departure. The symbolic content and social cost of deviating from appearance norms depend on the norm environment: an unconventional hairstyle in 1958 constituted a costly signal of nonconformity in a way that the same statistical rarity does not in a period of style pluralism. The mapping from observed stylistic deviation to latent cultural traits is thus context-dependent and need not be stable across periods. Our results should accordingly be read as internally valid for the transformation of the 1960s and 1970s — an episode in which style demonstrably became a salient margin of cultural expression — rather than as a general law linking stylistic and cultural change. Applying these measures to other periods would require re-establishing, for that context, that style carried comparable expressive content.

5.2 Drivers of Norm Deviation

What stylistic changes drove the collapse of the conformity equilibrium in the late 1960s, and the triumph of individualistic expression and style diversity in the early 1970s? To describe how individual style features relate to these patterns, we use simple OLS regressions within each decade–gender cell. We use horizontal norm deviation, vertical norm deviation, and style novelty as dependent variables, and style features as independent variables.

The resulting coefficients vary substantially across variables and decades. Appendix Figure A.10 plots coefficients from OLS regressions relating individual style characteristics to horizontal norm deviation, vertical norm deviation, and style novelty, for men and women separately. Take the horizontal norm deviation results, for example. In the 1950s and 1960s, wearing a shirt without a collar is associated with a higher horizontal norm deviation score for young men, corresponding to lower similarity with other, randomly chosen classmates. In later decades, the coefficient diminishes, consistent with the same attribute becoming less distinctive as it becomes more common. Bow ties have some of the most negative coefficients for horizontal norm deviation – and increasingly so in recent decades. In contrast, selected coefficients for sweaters, long hair, glasses, and facial hair are generally positive. For style novelty, coefficients for bowlcuts in female high-school senior portraits also vary by decade, illustrating that selected associations differ across both outcomes and periods.

Similar changes in coefficient size and sign occur for vertical norm deviation, where the sign must be read as an association with 1 minus cosine similarity to the earlier cohort. Positive coefficients therefore indicate greater deviation from the comparison cohort twenty years earlier. For example, crew cuts for men have negative coefficients for vertical norm deviation in the 1950s–1970s; by the 1990s, the coefficients are positive. Curtained hair also changes sign across decades, with positive coefficients in some earlier decades and negative coefficients later on. For women, low-neckline coefficients are positive in some earlier decades and negative in later decades. The changing coefficient patterns derived from our OLS exercise highlight one important advantage of our analytical approach. The association between individual elements of style choices and each measure is conditional on other people's choices; we can use them as time-varying, context-sensitive indicators of the social signal embedded in style choices.

5.3 Socioeconomic Correlates

Style choices suggest a dramatic fracture across regions from the 1970s onwards. Are socioeconomic forces or educational cleavages responsible? For example, greater norm-deviations in clothes among the less affluent could reflect that they have few other options to express their individuality – create a false impression of a class gradient.[18] To shed light on this question, we first examine how style choices (horizontal norm deviation, vertical norm deviation and style novelty) correlate with local socioeconomic and political characteristics. Since no such data exists at the high school level, we use county-level data from the decennial censuses, aggregate them to commuting zones, and linearly interpolate annual values between adjacent census years. The political split uses the Democratic two-party presidential vote margin, computed from the county-level election returns of Algara and Amlani (2021) (in turn compiled from CQ Press and ICPSR historical election returns) and aggregated to commuting zones.

Figure 8 shows the trend in norm deviation over time for four sample splits—urban vs. rural, rich vs. poor, Republican- vs. Democratic-leaning, and low vs. high share of nonwhites. The urban and income splits compare the bottom and top quartiles of commuting-zone urban population share and median family income, while the political split compares the corresponding quartiles of election-based Democratic vote margin. All show small differences and highly correlated trends. For men, rural areas show greater norm deviation from the 1970s onwards, but differences are small and both groups of commuting zones show parallel trends. Differences are even smaller between Republican- and Democratic-leaning areas, and the sign of differences reverses several times. Parallel trends also dominate for rich vs. poor commuting zones, with poorer areas slightly higher over much of the sample, though among men richer commuting zones overtake poorer ones after roughly 1990.[19] While in line with the hypothesis that style choices may be a favored way of expressing individual preferences, income-stratified differences are dwarfed by aggregate trends. Overall, our findings echo those of (Desmet and Wacziarg 2021), who also find that pre-defined categories like race do not have high predictive power for cultural markers.

Figure 8 reveals an interesting fact. The series for commuting zones in the bottom quartile of the nonwhite population share diverges over time from the series for commuting zones in the top quartile. This pattern could lead one to worry that the regional differences

[18] We thank Andrei Shleifer and several referees for emphasizing this point.

[19] Appendix Figure A.6 shows the corresponding style-diffusion patterns by commuting-zone income quartile.

in Figure 7 partly reflect differences in local racial composition rather than regional variation in style dynamics. Appendix Figure A.11 reports the same regional trends after adjusting the style measures for commuting-zone nonwhite population share within graduation year. The main patterns remain stable; the regional evolution of norm deviation and style novelty is not driven mechanically by differences in local racial composition.

5.4 Correlated Adoption and the Effect of Shocks

How did new styles diffuse in our high schools? We already documented the collapse of conformity in the late 1960s and growing regional divergence. But is adoption mainly driven by patterns at the high school level—or do aggregate trends dominate?

If high school seniors adopted styles as independent draws from a common national frequency distribution, then cross-school dispersion around each style's national take-up rate would follow a binomial benchmark (Figure 9). We show the binomial benchmark—the average of 500 random draws from a binomial distribution for the top 10 full-style combinations (the dashed blue line)—with adoption measured relative to the overall gender-specific take-up rate. The black line shows the actual school-level adoption frequencies, with the shaded band giving its 95% confidence interval. For both men and women, school-level adoption has excess mass in both tails, with fewer observations in the middle than the binomial benchmark predicts. Adoption is therefore not well described by independent individual draws from a common national style distribution; it is clustered at the school-year-gender level. This suggests either spatial sorting or possible externalities in adoption, with a few pioneers convincing larger numbers of classmates to follow suit, or correlated adoption decisions.

Two further checks reinforce this picture (Appendix E.2): within schools, transitions are orderly—the within-school split on a transitioning style rises temporarily and recedes as the school converges to a new norm—and adoption propensity is a persistent local trait, with commuting zones that closely track national trends for one style doing so for others as well.

Do style choices bear the marks of major economic and political shocks? We examine two—the "China shock" and the wave of casualties from the Vietnam war. The former is concentrated in time and space; the latter, less so. Neither produces a clear footprint in our data. Areas affected by the China shock did not become less norm-defiant, for example, and exposure to Vietnam casualties did not affect norm deviation in the 1960s and 1970s. We report additional descriptive evidence on these shocks in Appendix E.1.

5.5 Validations and Applications

How do our new measures of norm deviations and style novelty compare with other, alternative indicators? And are they correlated with substantive, high-stakes outcomes? Here, we examine similarities with four measures—attitudes from the GSS, and first names, gender roles in the home and workplace, and patenting.

5.5.1 GSS Attitudes

The GSS is conducted every two years, and provides disaggregated data at the state level. We compare our style measures for graduating high school seniors with the responses to a set of questions plausibly related to conformity, horizontal norm deviation, and norms. We do so for three broad sets of themes—counter-culture, gender roles, and anti-authoritarianism. The questions in the GSS related to counter-culture cover pre-marital sex, homosexuality, legalization of marijuana, political views (with liberalism entering positively), and church attendance (reverse-coded, i.e., higher frequency = lower value).[20] We use the Kling–Liebman–Katz standardized composite index to construct an indicator of counter-culture values from the responses. We use responses of GSS respondents aged 18–25, and match these cohorts to the style choices of graduating seniors in the same state. Panel C, Table A.1 summarizes our data.[21]

Table 2 gives the results. Panel A uses standardized norm deviation; Panel B repeats the exercise for standardized style novelty. In Panel A, we find a positive and significant coefficient for countercultural values (0.059, $p < 0.01$ pooled)—cohorts in states with more norm-defying styles support more counter-culture ideas. The corresponding full-specification R^2 is 0.096, including fixed effects. Effects are larger and statistically significant for men (0.070, $p < 0.05$) and smaller and not significant for women (0.036). The style-novelty estimates in Panel B show a similar positive pattern for countercultural values, especially

[20]The countercultural values index aggregates 5 GSS items: PREMARSX, HOMOSEX, GRASS, POLVIEWS (reversed), and ATTEND (reversed). The women's autonomy index aggregates 5 GSS items: FEFAM, FECHLD (reversed), FEPRESCH, FEPOL, and ABANY. The anti-authoritarianism index aggregates 6 GSS items: SPKATH, SPKRAC, SPKCOM, COLATH, LIBATH, and CAPPUN (reversed). Before standardization, all items are oriented so that higher values indicate more of the named construct. For example, higher ATTEND is recoded as lower countercultural values, FECHLD is recoded so disagreement that working mothers harm children raises women's autonomy, and CAPPUN is recoded so opposition to capital punishment raises anti-authoritarianism.

[21]We cannot link individual images with individual survey responses in GSS (or below, the CPS), or identify which high school survey respondents went to. Thus, the associations in the data should not be misinterpreted as individual-level correlations, but rather more aggregate relationships under some measurement error.

in the pooled and male specifications. The indicator for support of women's autonomy is constructed similarly. We aggregate questions on women's work vs family roles, the effect of working mothers on children, the impact on preschool children, the suitability of women for political office, and the permissibility of abortions. The pooled and female coefficients are negative and significant—more norm deviation is associated with *less* support for women's autonomy. The pooled coefficient is -0.077 ($p < 0.01$), with a larger estimate for women (-0.104, $p < 0.01$); the coefficient for men is negative but not significant (-0.064). For anti-authoritarianism, we use questions on the death penalty, free speech for atheists, racists, and communists, whether atheists should be allowed to teach, and whether anti-religious books should be kept in libraries. The coefficients are not statistically significant. We conclude that while style choice is significantly associated with some of the attitudes measured by the GSS, this is limited to a subset of questions closely related to 1960s counter-culture. The women's autonomy results are puzzling: more norm deviation predicts less, not more, support for women's independence, particularly among women. For anti-authoritarianism, there is no clear-cut pattern of association. Given the lack of granular data in the GSS, the absence of strong correlations for all indicators may well reflect a lack of variation.

5.5.2 Mobility

We also examine how style measures are associated with income and geographical mobility. While effects are not identified, the broad associations are interesting.

Using Opportunity Insights data on income mobility, we find higher chances of moving up the income distribution relative to parents for children who grew up in more norm-defiant, style-innovating commuting zones (subsection E.4).

Interstate movement – as documented in the CPS – is more frequent for persons from the CPS, aged 25-40, who graduated from high schools with greater norm deviations and more style innovations. The pattern is positive for men and women, but more pronounced — and significant — for the former.

5.5.3 Gender Roles and Life Cycle Outcomes

Are style choices in high school indicative of later-life outcomes? When counter-culture hit in the 1960s, it questioned traditional gender roles and family definitions. Here, we use data from the CPS to examine whether norm deviation among high school seniors predicts 'rebellion' against traditional gender roles.

The CPS contains individual-level microdata. We use respondents aged 25–40 surveyed up to 2010, and impute their high school graduation year (survey year minus age plus 17). For high school styles, we aggregate the norm deviation index to state × graduation-year cell means. We then merge CPS respondents to the norm deviation index of their state of residence at the time of survey and graduation cohort, along with census-based demographic controls (log population, urban share, white share, and population density) at the same state × cohort-year level. Panel B, Table A.1 gives an overview. We interpret the CPS estimates as descriptive state-cohort associations rather than individual high-school exposure estimates.

Table 3 shows that norm deviations in high school are negatively associated with being married in later life—consistent with a rejection of traditional family structures. The coefficient on the norm deviation index is significant for both genders: a one-standard-deviation increase reduces the probability of marriage by 0.7 percentage points for women and 0.8 percentage points for men (cols. 1–2). Parenthood is also lower for women and not statistically different from zero for men (cols. 3–4). For work outcomes, the associations run in opposite directions by gender: women are significantly more likely to work full-time and participate in the labor force, while men are significantly less likely to work full-time and show no significant change in labor-force participation (cols. 5–8). For both genders, we see a positive association with college education, too. The style-novelty estimates in Panel B show a similar, though smaller, pattern.

Table 4 aggregates these choices into a summary measure using the Kling–Liebman–Katz (2007) Traditional Lifestyle Index—a standardized composite of being married, having children, following a traditional work pattern (full-time for men, not working full-time for women), and not having college education.[22] Both women and men in state cohorts with more norm-defying high-school styles score lower on this traditional-lifestyle index.

5.5.4 Style Choices and Patenting

"Irreverence is a key to progress."

— Joel Mokyr, *A Culture of Growth*

Are style choices associated with economically important outcomes? Or are they only related to fashion? We examine the relationship between style-based, norm-defying behaviors and patenting using data on commuting-zone-level innovation from Bell et al. (2019). They

[22]The work component captures full-time work status, not household breadwinner status or whether a respondent is out of the labor force.

assemble data on both the number of inventors per commuting zone and birth cohort in the US, as well as patent applications and grants. For each year of age, their data trace out outcomes long after the high school portraits are taken. Panel A, Table A.1 gives an overview of the patenting data.[23]

To fix ideas, consider Figure 10. It shows the high school senior portrait of Steve Jobs, who graduated from Cupertino Homestead High in 1972. He is wearing a tuxedo, bow tie, and long hair, no facial hair, no glasses. The style combination appeared before, but rarely so—placing him in the rare-style tail for male students in 1972, with 0.0390% of the individuals in 1972 in our sample sharing his exact style. The style takes off in a big way soon after Jobs adopted it. Jobs also went on to apply for 1,114 patents during his lifetime, of which 960 were granted.

Next, we examine the pattern more generally. Figure A.12 shows patenting rates by years since high school graduation for commuting-zone cohorts with and without style innovators. We group commuting-zone-by-cohort cells by whether they contain at least one senior classified as a style innovator. Differences are substantial: cells with at least one style innovator show higher patenting rates than cells with none, for both male and female cohorts. The gap is considerable early on and then grows quickly; it stays high, reaching a peak of an additional 50% for men, and a comparative boost of more than 100% for women, roughly a decade after graduation, before declining gradually. [24] Higher patenting rates are associated with style innovation in all time periods.

In Table 5 we go beyond the visual evidence and exploit continuously varying measures of norm deviation and style innovation.[25] Applicant share (col. 1), grantee share (col. 2), and the number of grants awarded (col. 3) are all positively associated with both the norm deviation average (across horizontal and vertical measures), and the style novelty measure. Effects are substantial—a one standard deviation rise in norm deviations is associated with 0.144 standard deviations higher patent applications. This corresponds to moving the

[23]Some people move out of the commuting zone after graduating high school. Others may move into it. This creates measurement error—the cohort of patentees in any given commuting zone is only partly populated by people who graduated high school there. This measurement error will likely grow over time. This could cause a form of attenuation bias, leading to smaller estimates. For the most recent cohorts we only observe patenting for a few years after graduation. Magnitudes should be interpreted with caution; not the entire life cycle is captured for all the cohorts.

[24]If we abstract from the age dimension and collapse style innovation and patenting rates at the commuting zone/year level. We do this separately by gender, and plot the pattern by decade in Figure A.13.

[25]Table 5 requires availability of three data sources, i.e., the image data, the patenting data from (Bell et al. 2019) and all socioeconomic control variables from ICPSR 2896. This limits data to cohorts 1950-1989 from 485 commuting zones.

male Martinsville, Virginia, cohort in 1960 from about the 37th to the 53rd percentile of patent-application rates. Style novelty results are somewhat smaller, but highly significant throughout. Areas may be more fertile ground for technological innovation when youth culture permits or even encourages *more creative and original behavior*, whether between peers (horizontal norm deviation), over time (vertical norm deviation), or globally (style novelty).[26]

The relationships in Table 5 are not causal—putting an earring on a young man in 1956, or a tie on a young female, will not turn them into inventors. They are, however, compatible with a broader literature linking cultural openness, tolerance for nonconformity, and innovation (Gorodnichenko and Roland 2011). School policies or other local forces may jointly shape style innovation and technological innovation. In this sense, the style choices documented in high school yearbook images reflect broader social and cultural forces that are associated with innovation.

5.5.5 Norm Deviations from First Names

Important recent work has used unusual first names as an indicator of horizontal norm deviation (Bazzi et al. 2020; Raz 2025). Is our indicator of horizontal norm deviation capturing the same variation? Names are chosen by parents; clothes, subject to social constraints, by students themselves. We collect the first names of 752,444 graduating seniors from a random sample of up to 250 high-school yearbooks in each sampled year, every other year from 1950 to 2010, representing 7,484 yearbooks in total; the validation figures and regressions use the 746,744-observation complete-case sample. We extract the name of each graduating senior and the picture to which it corresponds. Name frequencies by state, sex, and cohort are taken from U.S. Social Security Administration's state-level baby-name files (Social Security Administration 2025), which report given-name counts by state of birth, sex, and year of birth. Let y denote the senior year and $b = y - 18$ the matched SSA cohort.[27] Because SSA reports name counts by state of birth, we benchmark within state-gender-cohort cells and read the exercise as validating a state-cohort analogue of our within-school measure.

Name horizontal norm deviation is defined as one minus the ratio of log count of frequency of name i relative to the most common name in the matched SSA state-gender cohort

[26] We do not know whether *youth* culture specifically fosters innovation later in life; similar patterns may emerge with style choices of adults at a later age.

[27] The CPS and patenting analyses match cohorts using a 17-year offset (the typical senior age); here we use 18 to align the SSA birth cohort with the graduating class.

(s, g, b): $\text{HND}^i_{s,g,b} = 1 - \frac{\log(c^i_{s,g,b})}{\log\left(\max_{j \in G_{s,g,b}} c^j_{s,g,b}\right)}$, where $c^i_{s,g,b}$ denotes the count of name i in that cohort.[28] Higher values correspond to more distinctive names. Name vertical norm deviation is constructed analogously, but compares the name to the matched SSA cohort 20 years earlier: $\text{VND}^i_{s,g,b} = 1 - \frac{\log(c^i_{s,g,b-20})}{\log\left(\max_{j \in G_{s,g,b-20}} c^j_{s,g,b-20}\right)}$, with higher values indicating greater cross-generational drift.

Figure E.4 shows scatterplots with Lowess smoothers of name horizontal norm deviation against style horizontal norm deviation (Panel A) and the analogous vertical norm deviation comparison (Panel B); for comparability with the name measure, the style-side variable here is benchmarked at the state-gender-cohort level rather than within school. The two share a common tendency — as style nonconformity rises, name horizontal norm deviation increases slightly over the range of the distribution. While the association is statistically significant (Table E.1), only a small share of variation in names is explained by style nonconformity. Within state, senior-year, and gender fixed-effect cells, we obtain a within-R^2 of 0.0004, suggesting that over 99.9% of variation in name horizontal norm deviation is unexplained by style nonconformity. A one standard-deviation increase in style horizontal norm deviations (see Figure E.5) is associated with only a 0.006 increase in name horizontal norm deviations. In the 1986 Kentucky male senior cohort, such an increase corresponds to moving from "Jerry" to "Terry" as a first name (the 28th and 29th most frequent names in that cohort out of 419 names observed). Overall, the evidence suggests that style and naming choices reflect related but distinct aspects of individuality.[29]

6 Conclusion

Imagine walking the halls of a typical American high school in the 1950s. For the boys, a sea of uniformity —clean-shaven faces, short, neatly trimmed hair, and, in their senior portraits, the same suit and tie, each looking earnestly into the camera. The girls vary more: most dress modestly and conservatively, but a few already experiment with new hairstyles, jewelry, or lower necklines. Fast forward to the late 1960s, and the yearbook pages show a dramatic shift. Suddenly, long hair, beards, and mustaches become visible markers of counterculture for men. Some men still dress conservatively, in the style of their parents; women on the

[28] The use of name frequencies at the state or national level to compile measures of *local* norm deviation is common (Raz 2025).

[29] We also derive name individuality from high school frequencies, with similar results.

other hand, while showing lower necklines than a generation before, and unkempt, long, undulating hippie-style hair, are more likely to look like each other.

The counterculture movement reshaped gender relations, political attitudes, and the meaning of work across the United States. Yet how can we measure such cultural transformation systematically? Visual self-representation through style and fashion is a central dimension of culture, but one that has so far largely resisted quantification.

We analyze millions of senior portraits spanning decades, extracting key style attributes. The length of a young man's hair or the depth of a young woman's neckline are more than just a personal choice; they become data points in a broader story of societal transformation. We use sparse feature vectors capturing key attributes of style. We train algorithms to identify style features in images. The resulting vector representations facilitate standard similarity measures, which we use to map two key dimensions of culture: norm deviation—both horizontal, within each graduating class, and vertical, between students and their parents' generation—and style novelty. This allows us to create a data-driven narrative of cultural evolution. Our approach contributes to a growing literature in economics that uses image data systematically (Adukia et al. 2023; Ludwig and Mullainathan 2024).

This setting has three key advantages. First, high school senior portraits cover a large share of American adolescents over decades. Second, they are standardized in format—head-and-shoulder shots taken under comparable conditions—which makes style variation meaningful rather than an artifact of photographic convention. Third, the data span a period of extraordinary cultural upheaval, offering the kind of variation needed to study the dynamics of conformity and rebellion.

There is marked convergence in style choice between men and women. In the 1950s and early 1960s, men's senior portraits were close to uniform, while women displayed substantially more variation. The late 1960s and early 1970s mark a clear break: male conformity within schools collapsed, continuity with the previous generation weakened, and the gender gap in norm deviation narrowed sharply. By the 1990s, male and female norm-deviation levels had largely converged, but the path was not symmetric. Men's norm deviation rose from very low levels, while women's initially high norm deviation declined and then stabilized at lower levels.

The counterculture-era break appears across U.S. Census regions, but the post-1970 evolution is uneven. From the 1970s onwards, polarization increases, with the South showing lower norm deviation than the rest of the country. Style novelty follows a related but distinct pattern: the South shows an especially sharp 1970s spike, whereas the Midwest and West

sustain higher novelty rates from the late 1980s onward. Thus, the collapse of the old conformity equilibrium was national in timing, but its persistence and later intensity varied substantially across regions.

Our measure of norm-defying behavior predicts other outcomes over the life cycle. In particular, it contains information about future gender roles—in areas where men and women showed more norm-breaking style choices in their high school senior portraits, traditional gender roles of male breadwinners and female home-makers declined more rapidly than in the US as a whole. Norm-defying high school graduates were also more likely to have rare and unusual names, but the pattern is weak. There is also some overlap between the attitudes captured by the GSS and style choices in high school yearbooks—with counter-culture attitudes favoring 'sex, drugs, and rock'n'roll' more common for cohorts that made bolder, norm-defying choices with their hair and clothes.

Finally, style choices of high school seniors correlate with an economically important activity—patenting. In commuting zones where more high school seniors adopted novel, unusual styles, dared to stand out in front of their peers, and deviated from the style of their parents, patenting rates are higher. We show that this finding holds at the cohort-commuting-zone level. While not causally identified, the association points to important complementarities between broad culture and high-end human creativity.

The philosopher Georg Simmel (1957) defined fashion as "a form of imitation. . . ,[which] paradoxically, in changing . . . differentiates one time from another." How imitation can simultaneously produce rapid change and long periods of uniformity has proved difficult to document empirically. Our methods—applied to uniquely granular, high-frequency data—allow us to trace how major style shifts emerge and the degree to which they break with prior patterns. We can also link sartorial choices with consequential outcomes: Graduates from high schools where norm-defying style was more common went on to patent more and to adopt less traditional gender roles. Daring to dress differently, it turns out, is associated with thinking — and living — differently.

References

Adukia, Anjali, Alex Eble, Emileigh Harrison, Hakizumwami Birali Runesha, and Teodora Szasz. 2023. "What We Teach About Race and Gender: Representation in Images and Text of Children's Books." *The Quarterly Journal of Economics* 138 (4): 2225–2285.

Algara, Carlos, and Sharif Amlani. 2021. "Partisanship & Nationalization in American Elections: Evidence from Presidential, Senatorial, & Gubernatorial Elections in the U.S. Counties, 1872–2020." *Electoral Studies* 73 102387.

Ash, Elliot, Ruben Durante, Maria Grebenshchikova, and Carlo Schwarz. 2022. "Visual Representation and Stereotypes in News Media." CEPR Discussion Paper No. 16624.

Athey, Susan, Dean Karlan, Emil Palikot, and Yuan Yuan. 2022. "Smiles in Profiles: Improving Fairness and Efficiency Using Estimates of User Preferences in Online Marketplaces." Working Paper 30633, National Bureau of Economic Research.

Atkin, David. 2016. "The Caloric Costs of Culture: Evidence from Indian Migrants." *American Economic Review* 106 (4): 1144–1181.

Autor, David H., David Dorn, and Gordon H. Hanson. 2013. "The China Syndrome: Local Labor Market Effects of Import Competition in the United States." *American Economic Review* 103 (6): 2121–2168.

Banerjee, Abhijit V. 1992. "A Simple Model of Herd Behavior." *The Quarterly Journal of Economics* 107 (3): 797–817.

Bazzi, Samuel, Martin Fiszbein, and Mesay Gebresilasse. 2020. "Frontier Culture: The Roots and Persistence of "Rugged Individualism" in the United States." *Econometrica* 88 (6): 2329–2368.

Becker, Gary S, and Kevin M Murphy. 1993. "A Simple Theory of Advertising as a Good or Bad." *The Quarterly Journal of Economics* 108 (4): 941–964.

Bell, Alex, Raj Chetty, Xavier Jaravel, Neviana Petkova, and John Van Reenen. 2019. "Who Becomes an Inventor in America? The Importance of Exposure to Innovation." *The Quarterly Journal of Economics* 134 (2): 647–713.

Benabou, Roland, and Jean Tirole. 2011. "Laws and Norms." Working Paper 17579, National Bureau of Economic Research.

Bertrand, Marianne, and Emir Kamenica. 2023. "Coming Apart? Cultural Distances in the United States over Time." *American Economic Journal: Applied Economics* 15 (4): 100–141.

Beyer, Andreas, and Steven Lindberg. 2003. *Portraits: A History*. H.N. Abrams.

Bigelow, Marybelle S, and Kay Kushino. 1979. *Fashion in History: Western Dress, Prehistoric to Present*. Class: Fashion Design, Burgess Publishing Company.

Braudel, Fernand. 1975. *Capitalism and Material Life, 1400-1800*. Harper Colophon.

Bursztyn, Leonardo, Alexander W Cappelen, Bertil Tungodden, Alessandra Voena, and David Yanagizawa-Drott. 2023. "How Are Gender Norms Perceived?" Working Paper 31049, National Bureau of Economic Research.

Bursztyn, Leonardo, Georgy Egorov, and Robert Jensen. 2019. "Cool to be Smart or Smart to be Cool? Understanding Peer Pressure in Education." *The Review of Economic Studies* 86 (4): 1487–1526.

Bursztyn, Leonardo, Alessandra L González, and David Yanagizawa-Drott. 2020. "Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia." *American Economic Review* 110 (10): 2997–3029.

Caprini, Giulia. 2023. "Visual Bias." Economics Series Working Papers 1016, University of Oxford, Department of Economics.

Carbon, Claus-Christian. 2017. "Universal Principles of Depicting Oneself across the Centuries: From Renaissance Self-Portraits to Selfie-Photographs." *Frontiers in Psychology* 8 246418.

Chetty, Raj, Will Dobbie, Benjamin Goldman, Sonya R. Porter, and Crystal S. Yang. 2026. "Changing Opportunity: Sociological Mechanisms Underlying Growing Class Gaps and Shrinking Race Gaps in Economic Mobility." *The Quarterly Journal of Economics* 141 (2): 1137–1210. 10.1093/qje/qjaf057.

Crafts, Nicholas F R. 1985. *British Economic Growth During the Industrial Revolution*. Oxford: Oxford University Press.

Davis, Fred. 2013. *Fashion, Culture, and Identity*. University of Chicago Press.

Desmet, Klaus, and Romain Wacziarg. 2021. "The cultural divide." *The Economic Journal* 131 (637): 2058–2088.

Galton, Francis J. 1878. "Composite Portraits." *Nature* 18 97–100.

Gavrilets, Sergey. 2020. "The Dynamics of Injunctive Social Norms." *Evolutionary Human Sciences* 2 e60.

Gelfand, Michele J, Sergey Gavrilets, and Nathan Nunn. 2024. "Norm Dynamics: Interdisciplinary Perspectives on Social Norm Emergence, Persistence, and Change." *Annual Review of Psychology* 75 341–378.

Gelfand, Michele J, Jana L Raver, Lisa Nishii et al. 2011. "Differences Between Tight and Loose cultures: A 33-Nation Study." *Science* 332 (6033): 1100–1104.

Gentzkow, Matthew, Jesse M Shapiro, and Matt Taddy. 2019. "Measuring Group Differences in High-Dimensional Choices: Method and Application to Congressional Speech." *Econometrica* 87 (4): 1307–1340.

Ginosar, Shiry, Kate Rakelly, Sarah Sachs, Brian Yin, and Alexei A Efros. 2015. "A Century of Portraits: A Visual Historical Record of American High School Yearbooks." In *Proceedings of the IEEE International Conference on Computer Vision Workshops*, 1–7.

Giuliano, Paola, and Nathan Nunn. 2021. "Understanding Cultural Persistence and Change." *The Review of Economic Studies* 88 (4): 1541–1581.

Gorodnichenko, Yuriy, and Gerard Roland. 2011. "Which Dimensions of Culture Matter for Long-Run Growth?" *American Economic Review* 101 (3): 492–498.

Graham, Gael. 2004. "Flaunting the Freak Flag: Karr v. Schmidt and the Great Hair Debate in American High Schools, 1965–1975." *The Journal of American History* 91 (2): 522–543.

Hancock, Joseph, Toni Johnson-Woods, and Vicki Karaminas. 2013. *Fashion in Popular Culture: Literature, Media and Contemporary Studies*. Intellect Books.

Heath, Joseph, and Andrew Potter. 2004. *Nation of Rebels: Why Counterculture Became Consumer Culture*. New York: HarperBusiness.

Hidayati, Shintami C, Kai-Lung Hua, Wen-Huang Cheng, and Shih-Wei Sun. 2014. "What are the Fashion Trends in New York?" In *Proceedings of the 22nd ACM International Conference on Multimedia*, 197–200.

Hofstede, Geert. 2011. "Dimensionalizing Cultures: The Hofstede Model in Context." *Online Readings in Psychology and Culture* 2 (1): 8.

Issitt, Micah. 2009. *Hippies: A Guide to an American Subculture*. ABC-CLIO.

Karni, Edi, and David Schmeidler. 1990. "Fixed Preferences and Changing Tastes." *The American Economic Review* 80 (2): 262–267.

Kiapour, M Hadi, Kota Yamaguchi, Alexander C Berg, and Tamara L Berg. 2014. "Hipster Wars: Discovering Elements of Fashion Styles." In *Computer Vision–ECCV 2014: 13th European Conference, Zurich, Switzerland, September 6-12, 2014, Proceedings, Part I 13*, 472–488, Springer.

Kopkind, A. 1979. "Dressing Up: Mysteries of Fashion Revealed." *The Village Voice* 24 (1): 33–38.

Kuran, Timur. 1987. "Preference Falsification, Policy Continuity and Collective Conservatism." *The Economic Journal* 97 (387): 642–665.

Kuran, Timur. 1989. "Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution." *Public Choice* 61 (1): 41–74.

Leibenstein, Harvey. 1950. "Bandwagon, Snob, and Veblen Effects in the Theory of Consumers' Demand." *The Quarterly Journal of Economics* 64 (2): 183–207.

Ludwig, Jens, and Sendhil Mullainathan. 2024. "Machine Learning as a Tool for Hypothesis Generation." *The Quarterly Journal of Economics* 139 (2): 751–827.

MacFarlane, Alan. 1978. "The Origins of English Individualism: Some Surprises." *Theory and Society* 6 (2): 255–277.

Matsuyama, Kiminori. 1993. "Custom Versus Fashion: Path-Dependence and Limit Cycles in a Random Matching Game." Discussion Paper / Institute for Empirical Macroeconomics 82, Federal Reserve Bank of Minneapolis.

McKendrick, Neil, John Brewer, and John Harold Plumb. 1982. *The Birth of a Consumer Society: The Commercialization of Eighteenth-Century England*. Routledge.

Meta Llama., "Llama-4-Scout-17B-16E." https://huggingface.co/meta-llama/Llama-4-Scout-17B-16E, Hugging Face model page, accessed April 20, 2026.

Michalopoulos, Stelios, and Christopher Rauh. 2024. "Movies." Working Paper 32220, National Bureau of Economic Research.

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

Murdock, Jaimie, Colin Allen, and Simon DeDeo. 2017. "Exploration and exploitation of Victorian science in Darwin's reading notebooks." *Cognition* 159 117–126.

National Archives and Records Administration., "Combat Area Casualties Current File (CACCF)." Electronic Records Division, U.S. National Archives and Records Administration, Records of U.S. military casualties during the Vietnam War.

Nosch, Marie-Louise, Cécile Michel, and Mary Harlow. 2014. *Prehistoric, Ancient Near Eastern & Aegean Textiles and Dress*. Oxbow Books.

Pesendorfer, Wolfgang. 1995. "Design Innovation and Fashion Cycles." *The American Economic Review* 771–792.

prithivMLmods., "Realistic-Gender-Classification." https://huggingface.co/prithivMLmods/Realistic-Gender-Classification, Hugging Face model page, accessed April 20, 2026.

Prodger, Phillip. 2021. *Face Time: A History of the Photographic Portrait*. WW Norton.

Raw, Nate. 2023. "vit-age-classifier (Revision 461a4c4)." 10.57967/hf/1259.

Raz, Itzchak Tzachi. 2025. "Soil Heterogeneity, Social Learning, and the Formation of Close-Knit Communities." *Journal of Political Economy* 133 (8): 2643–2691.

Riello, Giorgio, and Ulinka Rublack. 2019. *The Right to Dress: Sumptuary Laws in a Global Perspective, c. 1200–1800*. Cambridge University Press.

Shofia, Naila. 2022. "Why veil? Religious headscarves and the public role of women." *Religious Headscarves and The Public Role of Women (July 21, 2022)*.

Simmel, Georg. 1957. "Fashion." *American Journal of Sociology* 62 (6): 541–558.

Social Security Administration. 2025. "State-Specific Data on the Relative Frequency of Given Names in the Population of U.S. Births Where the Individual Has a Social Security Number." Data file: `namesbystate.zip`, Tabulated from Social Security records as of March 2, 2025.

Todorov, Alexander, Manish Pakrashi, and Nikolaas N Oosterhof. 2009. "Evaluating Faces on Trustworthiness After Minimal Time Exposure." *Social Cognition* 27 (6): 813–833.

Veblen, Thorstein. 1899. *The Theory of the Leisure Class*. New York: Macmillan.

Wang, Zhou, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. 2004. "Image Quality Assessment: From Error Visibility to Structural Similarity." *IEEE Transactions on Image Processing* 13 (4): 600–612.

Welters, Linda, and Abby Lillethun. 2018. *Fashion History: A Global View*. Bloomsbury Publishing.

Figures

Figure 1: Sample Images and Style Attributes from High School Yearbooks

Style Feature	Model Dummies	
	Image A	Image B
long hair	1	1
curtained hair	0	0
curly hair	1	1
spiky hair	0	0
bowlcut hair	0	0
crew-cut hair	0	0
afro	0	0
beard	0	0
moustache	0	0
earrings	0	1
necklace	1	1
glasses	0	0
low neckline	0	1
sweater	0	0
suit	0	0
shirt with collar	1	0
shirt without collar	0	0
bowtie	0	0
tie	0	0



Image A



Image B

Style Feature	Model Dummies	
	Image C	Image D
long hair	0	1
curtained hair	0	1
curly hair	0	0
spiky hair	0	0
bowlcut hair	0	0
crew-cut hair	1	0
afro	0	0
beard	0	0
moustache	0	0
earrings	0	0
necklace	0	0
glasses	0	1
low neckline	0	0
sweater	0	0
suit	1	1
shirt with collar	1	1
shirt without collar	0	0
bowtie	0	1
tie	1	0



Image C

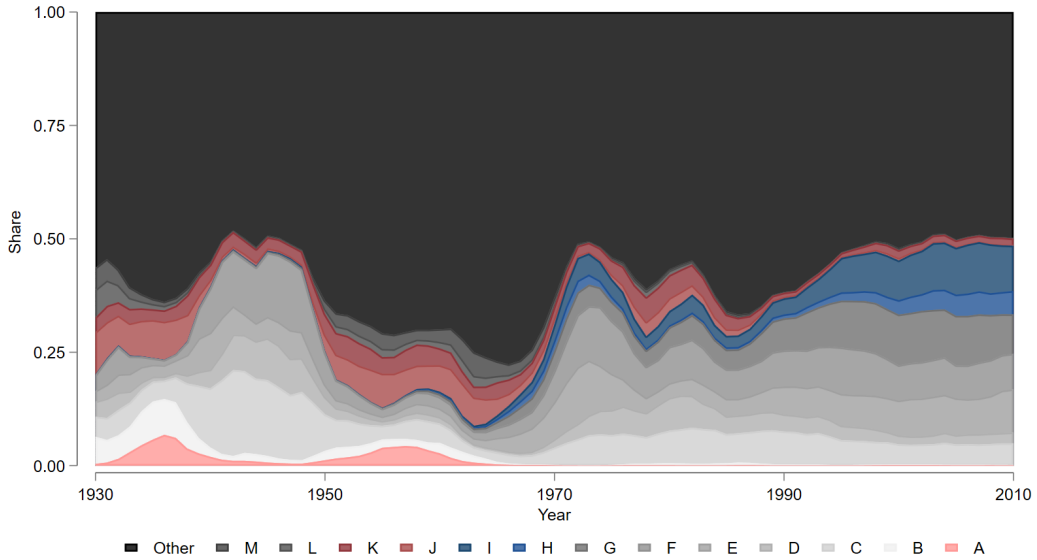


Image D

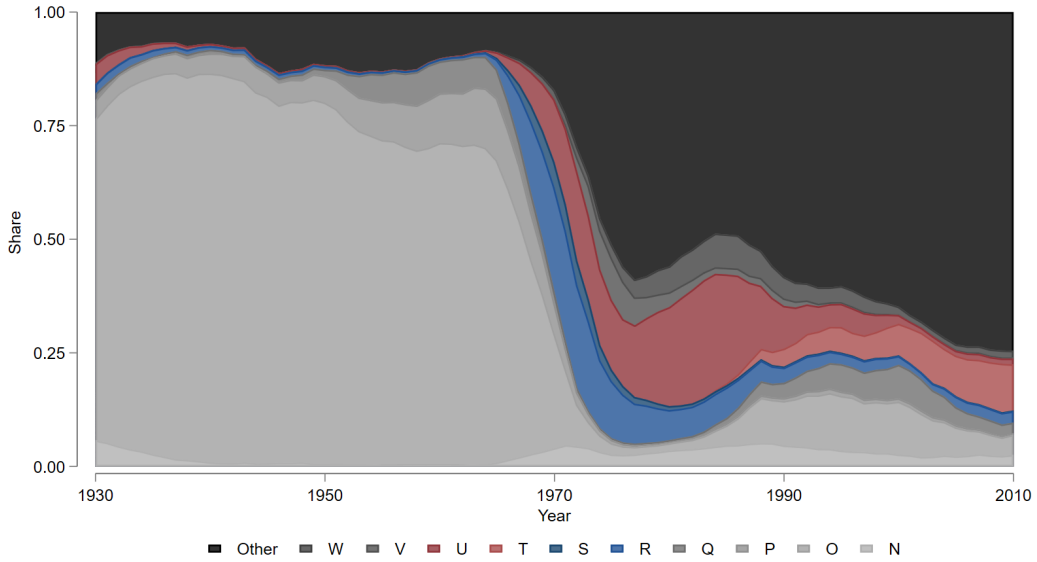
Notes: The figure illustrates the construction of individual style vectors from high school senior portraits. For each image, we report binary indicators for visible hairstyle, clothing, facial hair, and accessory features. A value of 1 indicates that the feature is present, while 0 indicates that it is absent. The indicators are obtained from Google Vertex AI model predictions after converting images to grayscale. Continuous prediction scores are discretized using a threshold of 0.30. See Appendix C for details on model training, validation, and the construction of style vectors.

Figure 2: Styles in High School Senior Yearbook Pictures

(a) Women

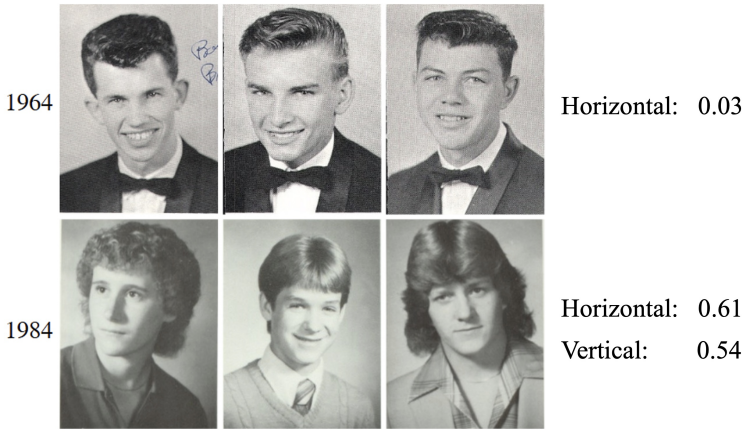


(b) Men



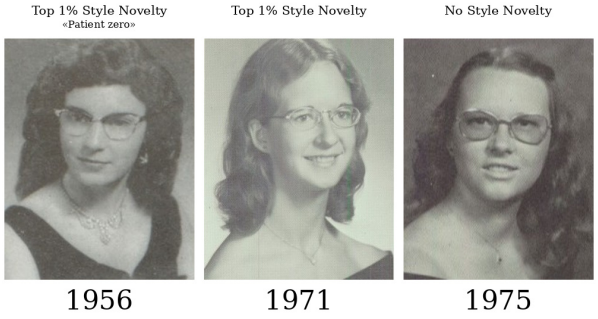
Notes: This graph depicts the share over time of discrete combinations of style features, split by gender. Each image is assigned a style according to the features available. As there are many thousand unique combinations, for illustrative purposes here, we assign the residual category "Other" to all unique style combinations that comprise less than 5% of the sample in every year. In our analysis of style novelty, we use all unique style combinations as they are defined, and not "Other".

Figure 3: Norm Deviation, Example



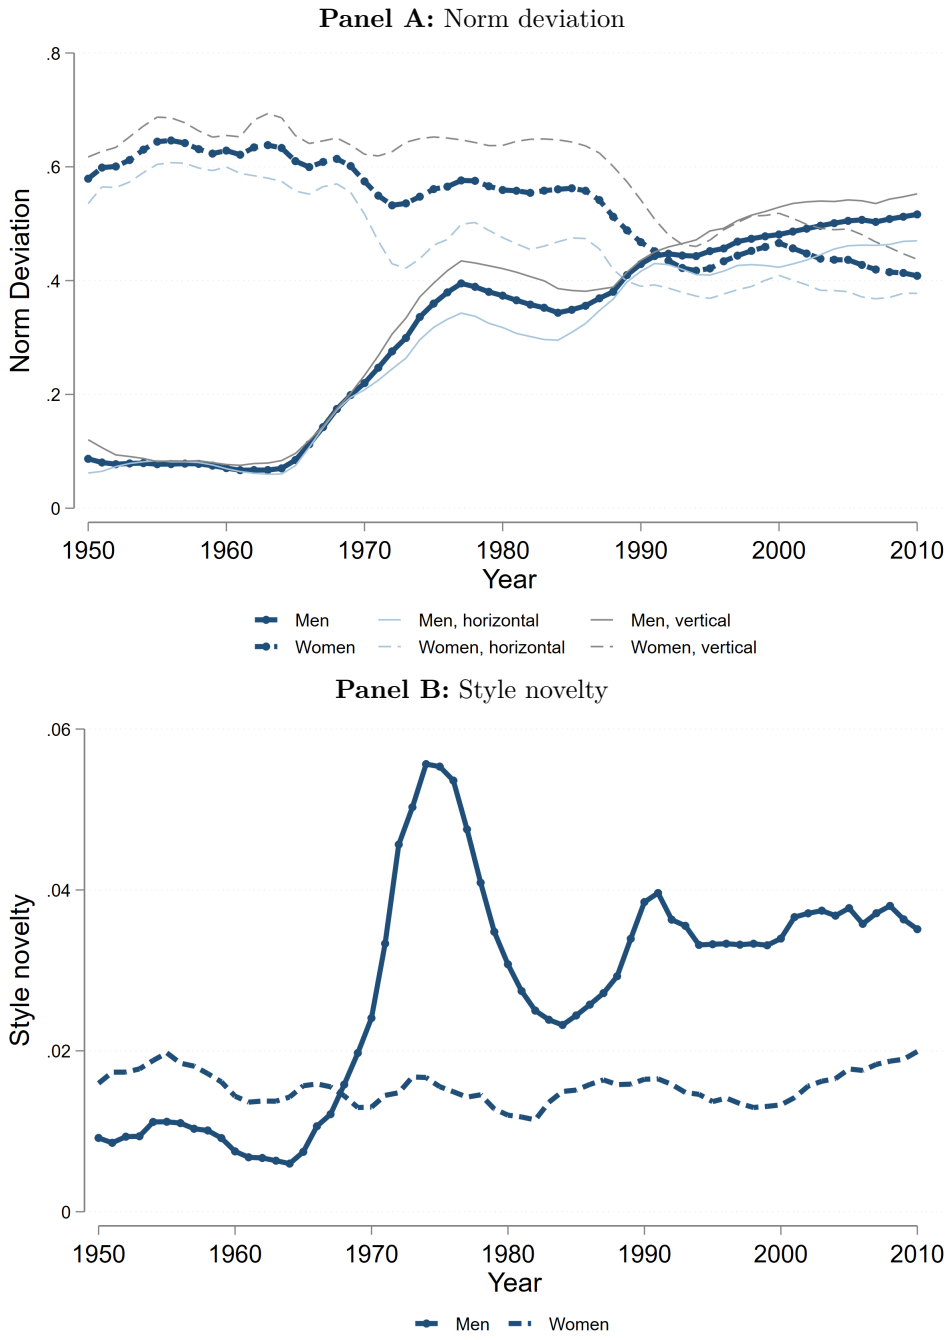
Notes: This figure compares male graduates from Preston High School, Idaho, in two cohorts twenty years apart: 1964 and 1984. The labels on the right report horizontal norm deviation scores for each displayed cohort and the vertical norm deviation score for the cross-cohort comparison: 0.03 for horizontal norm deviation in 1964, 0.61 for horizontal norm deviation in 1984, and 0.54 for vertical norm deviation when comparing the 1984 style vectors with those from the 1964 cohort. In the main analysis, vertical norm deviation is measured relative to the same high school twenty years earlier, to capture generational shifts. Appendix Figure A.7 reports analogous comparisons over alternative time horizons.

Figure 4: Style Novelty, Example



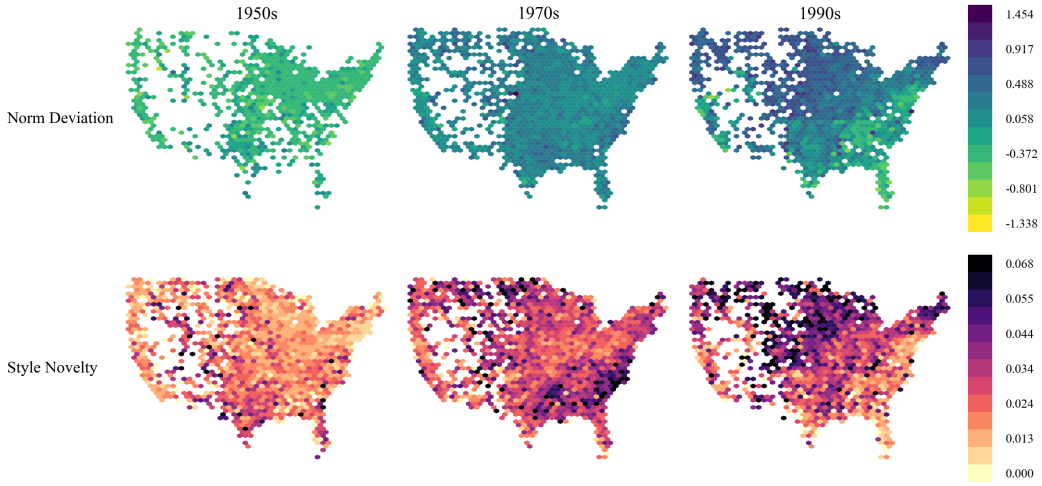
Notes: The figure illustrates our style novelty measure for the displayed female style combination. The left-most image shows the earliest example of this style in our data in 1956. The same style appears again in subsequent years and remains in the same-gender rare-style tail through the middle-image year, reflecting limited diffusion. By 1975, after broader diffusion, the image on the right no longer qualifies as novel.

Figure 5: Norm Deviation and Style Novelty over Time, by Gender



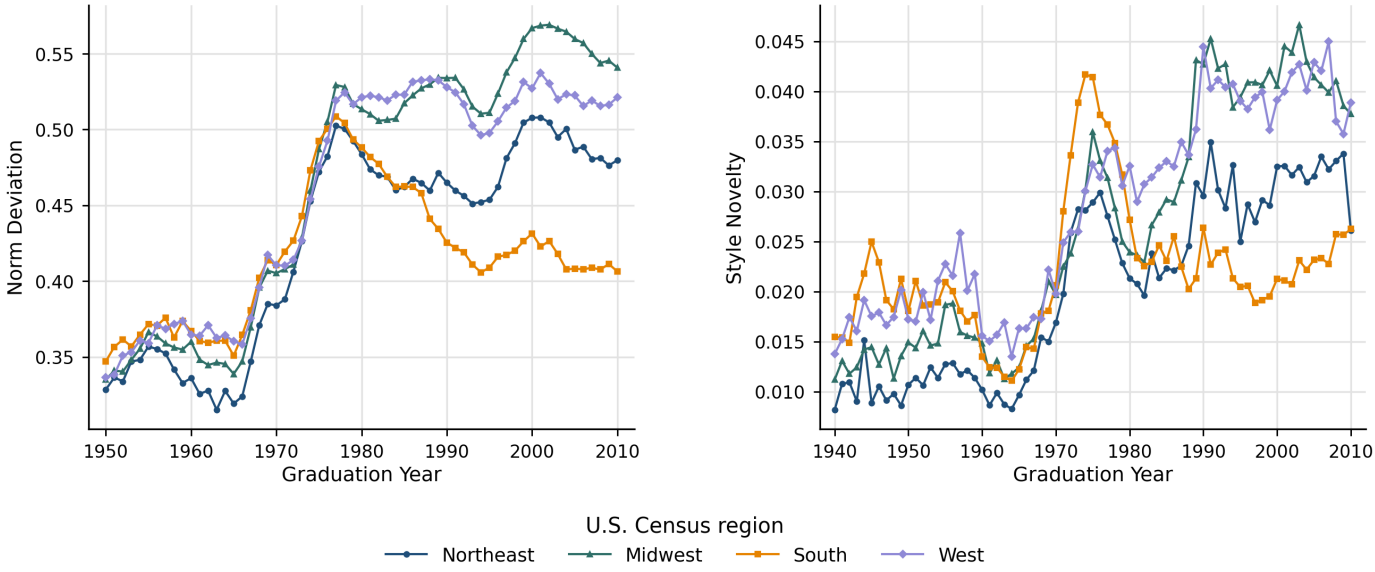
Notes: This figure plots the evolution of yearbook-based style measures by gender. Panel A reports horizontal norm deviation, vertical norm deviation, and the norm-deviation index, defined as the average of the two components. Panel B reports style novelty, measured as the share of students classified as style innovators. Each series is computed from individual-level senior portraits and averaged by graduation year separately for women and men. The sample covers more than 54 million high school senior portraits from 1930 to 2010.

Figure 6: Spatial Distribution



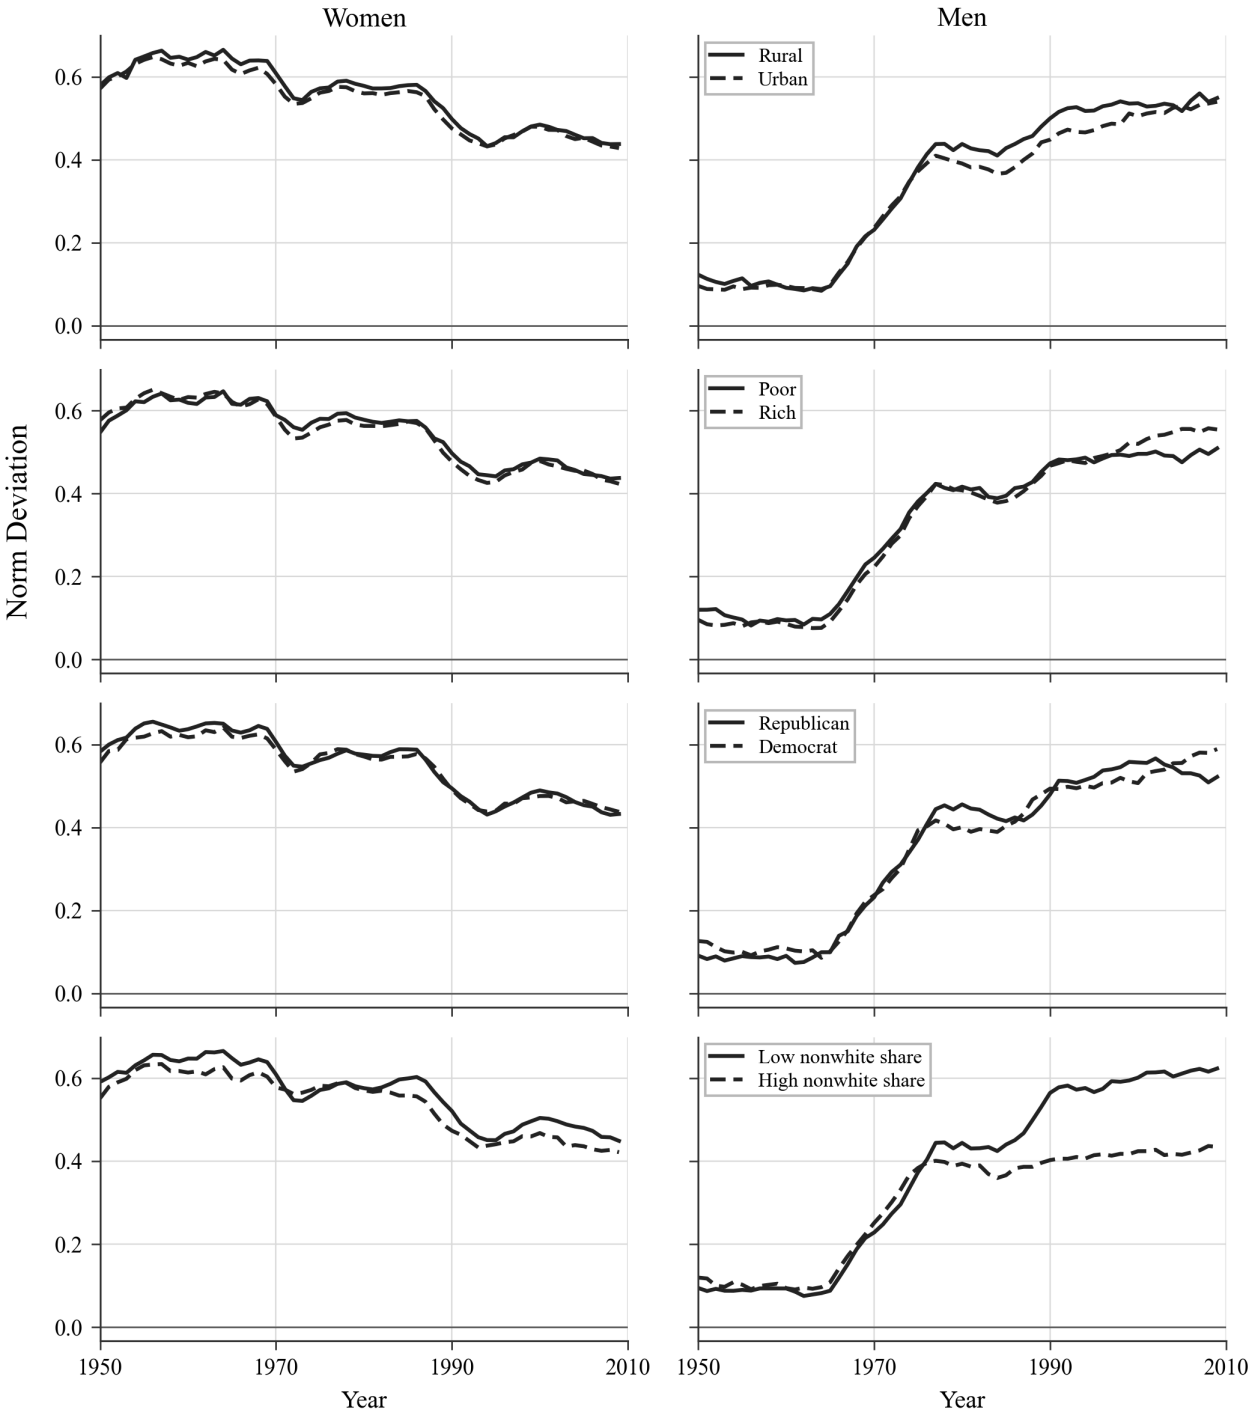
Notes: This figure shows the geographical evolution of yearbook-based style measures across decades. The top row reports norm deviation; the bottom row reports style novelty. Each point corresponds to a decade-by-high-school geolocation average, pooling male and female seniors. Norm deviation is standardized as a z-score; style novelty is plotted as a share.

Figure 7: Norm Deviation and Style Novelty across U.S. Census Regions



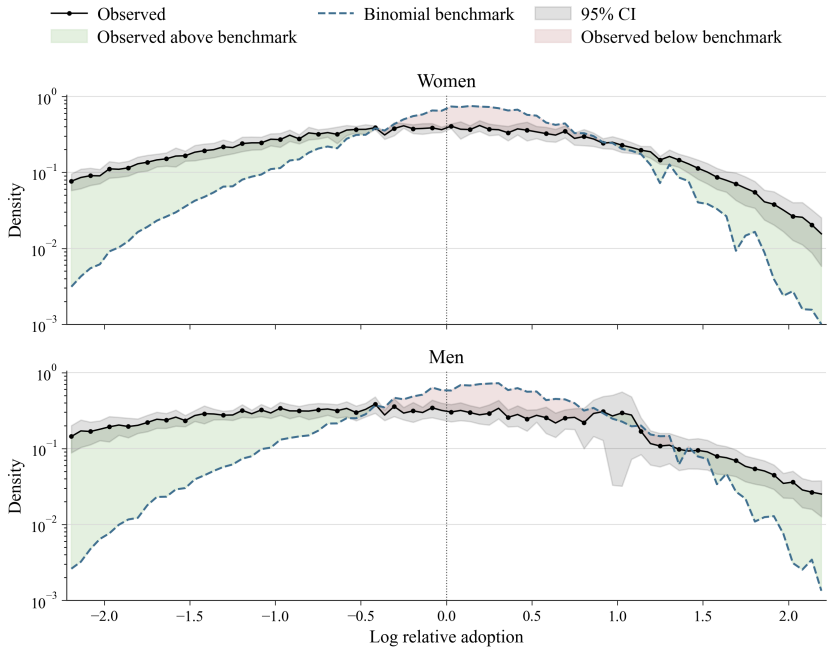
Notes: This figure plots the evolution of norm deviation and style novelty across U.S. Census regions. The left panel reports average norm deviation; the right panel reports the average share of style innovators. Observations are commuting-zone-by-graduation-year cells. Each commuting zone is assigned to a single U.S. state via a fixed commuting-zone-to-state crosswalk and then to that state's Census region; commuting zones spanning multiple states are assigned as a whole to one region rather than split across regions. Series report unweighted regional averages by graduation year.

Figure 8: Norm Deviation across Local Socioeconomic and Political Environments



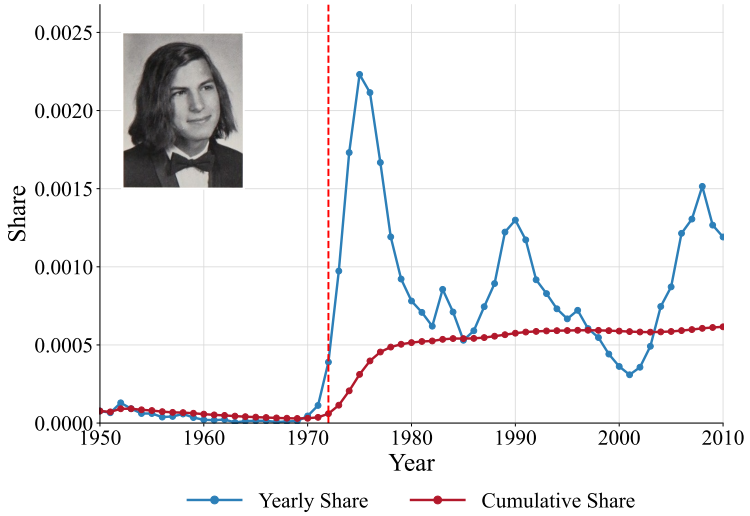
Notes: This figure plots norm deviation across commuting zones from 1950 to 2010, separately for women and men. Each row compares commuting zones in the bottom and top quartiles of a local characteristic, recomputed within year: census-derived urban population share, census-derived median family income, the county-level Democratic two-party vote margin (Algara and Amlani 2021), and census-derived nonwhite population share. The middle quartiles are omitted. Cell averages are unweighted.

Figure 9: Dispersion in Style Adoption across High School Cohorts



Notes: This figure compares observed dispersion in style adoption across high school cohorts with a binomial benchmark based on independent draws from the gender-specific national take-up rate of each style. Cohorts are school-year-gender cells. For each gender, the figure uses the ten most common full-style combinations. Adoption is measured relative to the overall gender-specific take-up rate of each style and plotted in logs with a continuity correction. The dashed line reports simulated binomial adoption; points report the observed distribution with 95% confidence intervals.

Figure 10: The Diffusion of Steve Jobs' Senior-Portrait Style



Notes: This figure traces the diffusion of the style combination observed in Steve Jobs' 1972 senior portrait from Homestead High School in Cupertino, California. The blue series reports the yearly share of senior portraits with the same style combination; the red series reports the cumulative share up to each year. The dashed vertical line marks Jobs' graduation year.

Tables

Table 1: Summary Statistics, Image Data

	Image Level Panel			CZ Level Panel		
	N	Mean	St. Dev.	N	Mean	St. Dev.
Style Attributes						
Curtained hair	54,952,980	0.255	0.436	33,407	0.222	0.119
Curly hair	54,952,980	0.284	0.451	33,407	0.269	0.128
Spiky hair	54,952,980	0.025	0.157	33,407	0.023	0.041
Bowlcut hair	54,952,980	0.103	0.304	33,407	0.104	0.088
Crew cut hair	54,952,980	0.238	0.426	33,407	0.276	0.221
Bowtie	54,952,980	0.069	0.254	33,407	0.070	0.123
Beard	54,952,980	0.004	0.060	33,407	0.004	0.007
Afro hair	54,952,980	0.006	0.078	33,407	0.005	0.018
Low neckline	54,952,980	0.139	0.346	33,407	0.121	0.129
Earrings	54,952,980	0.066	0.249	33,407	0.065	0.059
Tie	54,952,980	0.343	0.475	33,407	0.327	0.165
Glasses	54,952,980	0.093	0.291	33,407	0.105	0.070
Shirt without collar	54,952,980	0.225	0.418	33,407	0.207	0.087
Shirt with collar	54,952,980	0.656	0.475	33,407	0.678	0.113
Moustache	54,952,980	0.028	0.165	33,407	0.025	0.031
Suit	54,952,980	0.406	0.491	33,407	0.395	0.134
Necklace	54,952,980	0.254	0.436	33,407	0.233	0.095
Sweater	54,952,980	0.062	0.241	33,407	0.069	0.048
Long hair	54,952,980	0.425	0.494	33,407	0.412	0.163
Style Choices						
Horizontal norm deviation - Pooled	54,952,980	0.368	0.218	33,407	0.375	0.089
Horizontal norm deviation - Males	26,099,457	0.243	0.204	33,407	0.254	0.184
Horizontal norm deviation - Females	28,853,523	0.481	0.160	33,407	0.494	0.094
Vertical norm deviation - Pooled	32,546,590	0.478	0.235	27,431	0.481	0.092
Vertical norm deviation - Males	15,331,151	0.342	0.241	27,431	0.339	0.187
Vertical norm deviation - Females	17,215,439	0.598	0.147	27,431	0.619	0.080
Norm deviation - Pooled	32,546,590	0.430	0.208	27,431	0.433	0.085
Norm deviation - Males	15,331,151	0.314	0.216	27,431	0.308	0.183
Norm deviation - Females	17,215,439	0.533	0.134	27,431	0.554	0.077
Style Novelty - Pooled	54,901,497	0.020	0.140	33,407	0.023	0.020
Style Novelty - Males	26,071,628	0.026	0.158	33,407	0.030	0.035
Style Novelty - Females	28,829,869	0.015	0.121	33,407	0.017	0.016

Notes: This table reports summary statistics for style features and norm deviation measures at the individual level and at the commuting-zone level over the sampled period from 1930 to 2010. The N column reports the effective denominator for each row after applying the relevant measure-specific requirement to the post-filter senior portrait sample. Style features are measured from senior portraits showing heads and upper torsos and discretized as 0–1 dummy variables. Horizontal norm deviation is calculated as one minus the average cosine similarity between an individual and same-gender classmates in the same school-year cell, using cells with at least two portraits. Vertical norm deviation is calculated analogously relative to the same-gender cohort in the same high school 20 years earlier, $t - 20$. The norm deviation index is the average of horizontal and vertical norm deviation. Style novelty is a dummy for whether a student's full-style combination falls in the rare-style tail for the student's gender and year, based on cumulative occurrence counts. The cutoff is tied to 1% of eligible cumulative portraits, includes tied style combinations at the boundary, and can therefore yield a realized image-level share different from exactly one percent.

Table 2: Attitudes, General Social Survey

	Countercultural Values			Women's Autonomy			Anti-Authoritarianism		
	All	F	M	All	F	M	All	F	M
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Panel A: Norm Deviation									
Norm Deviation (std.)	0.059***	0.036	0.070**	-0.077***	-0.104***	-0.064	-0.032	-0.000	-0.061
	(0.022)	(0.031)	(0.031)	(0.028)	(0.037)	(0.042)	(0.027)	(0.038)	(0.037)
Observations	6,715	3,642	3,073	5,584	3,053	2,531	6,164	3,317	2,846
Mean Dep. Var.	-0.006	-0.041	0.027	-0.020	0.058	-0.094	0.003	0.025	-0.017
R^2	0.096	0.126	0.134	0.073	0.107	0.099	0.057	0.075	0.081
Panel B: Style Novelty									
Style Novelty (std.)	0.044***	0.035	0.048**	-0.027	-0.052**	-0.005	-0.012	0.004	-0.021
	(0.015)	(0.022)	(0.021)	(0.019)	(0.026)	(0.027)	(0.018)	(0.024)	(0.026)
Observations	6,720	3,646	3,074	5,588	3,056	2,532	6,168	3,320	2,848
Mean Dep. Var.	-0.005	-0.040	0.028	-0.021	0.057	-0.095	0.003	0.025	-0.018
R^2	0.097	0.127	0.135	0.072	0.104	0.098	0.057	0.075	0.081

Notes: The table reports weighted OLS estimates for GSS composite attitude indices among respondents aged 18–25. Panel A uses the standardized norm deviation; Panel B uses standardized style novelty. Coefficients correspond to a one-standard-deviation increase in the relevant style measure. Composite outcomes follow Kling–Liebman–Katz (2007): each component is standardized within the early-adult analysis sample, oriented so that larger values indicate more of the named construct—countercultural values, women's autonomy, or anti-authoritarianism—and then averaged. All specifications include state, graduation-cohort-year, and survey-year fixed effects, as well as demographic controls measured at the respondent's state and graduation cohort: log population, urban population share, white population share, and population density. Survey weights are applied. Standard errors are clustered at the state-by-cohort level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 3: Lifestyles, Current Population Survey

	Married		Any Child		Full-time Work		Labor Force		Some College	
	F	M	F	M	F	M	F	M	F	M
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Panel A: Norm Deviation										
Norm Deviation (std.)	-0.007***	-0.008***	-0.004*	-0.002	0.011***	-0.004**	0.008***	-0.000	0.006**	0.010***
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.001)	(0.002)	(0.002)
Observations	847,574	780,280	790,991	728,625	847,574	780,280	846,904	763,685	839,806	773,143
Mean Dep. Var.	0.676	0.638	0.703	0.522	0.546	0.870	0.681	0.935	0.476	0.486
R^2	0.051	0.090	0.064	0.093	0.043	0.027	0.066	0.010	0.084	0.036
Panel B: Style Novelty										
Style Novelty (std.)	-0.002*	-0.004***	-0.003**	-0.002	0.005***	-0.002*	0.004***	-0.000	0.006***	0.006***
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Observations	849,324	781,849	791,812	729,410	849,324	781,849	848,652	765,142	841,330	774,513
Mean Dep. Var.	0.676	0.638	0.703	0.522	0.546	0.870	0.680	0.935	0.476	0.485
R^2	0.051	0.091	0.064	0.093	0.043	0.028	0.066	0.010	0.085	0.036

Notes: The table reports weighted OLS estimates linking standardized yearbook-based style measures to adult lifecycle outcomes. Panel A uses the standardized norm deviation; Panel B uses standardized style novelty. Coefficients correspond to a one-standard-deviation increase in the relevant style measure. The dependent variables are binary indicators, estimated using linear probability models and reported separately by gender for marriage, parenthood, full-time work as a share of the population sample, labor-force participation, and college attendance. Observation counts vary across columns because the CPS sample is restricted to respondents with the relevant outcome, gender, survey-year, state-by-graduation-cohort style match, and controls. All specifications include state, graduation-cohort-year, and survey-year fixed effects, as well as state-by-cohort demographic controls. Survey weights are applied. Standard errors are clustered at the state-by-cohort level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 4: Norm Deviation, Style Novelty, and Traditionality

	Traditional Lifestyle Index		
	All	F	M
	(1)	(2)	(3)
Panel A: Norm Deviation			
Norm Deviation (std.)	-0.014***	-0.015***	-0.013***
	(0.002)	(0.003)	(0.003)
Observations	1,627,854	847,574	780,280
Mean Dep. Var.	-0.025	-0.104	0.055
R^2	0.113	0.112	0.126
Panel B: Style Novelty			
Style Novelty (std.)	-0.008***	-0.009***	-0.007***
	(0.001)	(0.002)	(0.002)
Observations	1,631,173	849,324	781,849
Mean Dep. Var.	-0.025	-0.103	0.056
R^2	0.114	0.112	0.127

Notes: The table reports weighted OLS estimates linking yearbook-based style measures to adult traditionality outcomes. The dependent variable is the Kling–Liebman–Katz (2007) traditional lifestyle index, constructed by averaging standardized measures of marriage, parenthood, gender-specific full-time-work patterns, and lower college attendance, oriented so that larger values indicate more traditional lifecycle choices. Following the Kling–Liebman–Katz convention, the index averages each respondent's standardized components over those that are non-missing, so a respondent enters the sample if at least one component is observed; the index sample can therefore exceed that of any single component. Panel A uses the standardized norm deviation; Panel B uses standardized style novelty. Coefficients correspond to a one-standard-deviation increase in the relevant style measure. All specifications include state, graduation-cohort-year, and survey-year fixed effects, as well as state-by-cohort demographic controls. Survey weights are applied. Standard errors are clustered at the state-by-cohort level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table 5: Patenting

	Applicant Share	Grantee Share	Number of Grants
	(1)	(2)	(3)
Panel A: Norm Deviation			
Norm Deviation (std.)	0.144***	0.085***	0.090***
	(0.016)	(0.015)	(0.015)
Observations	34,211	34,211	34,211
Mean Dep. Var.	0.535	0.440	0.671
R^2	0.546	0.545	0.483
Panel B: Style Novelty			
Style Novelty (std.)	0.042***	0.030***	0.025***
	(0.009)	(0.008)	(0.008)
Observations	34,211	34,211	34,211
Mean Dep. Var.	0.535	0.440	0.671
R^2	0.542	0.544	0.482

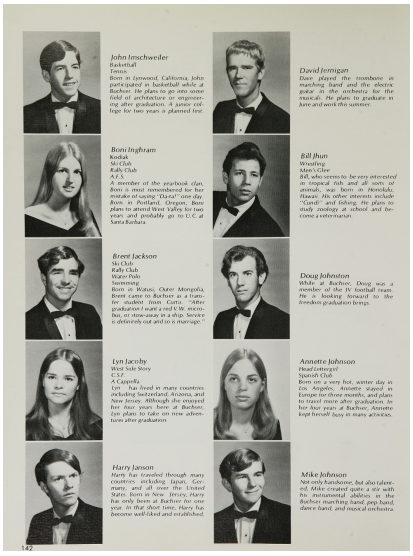
Notes: This table reports full-specification regressions of patenting outcomes on yearbook-based style measures among high school senior cohorts. The unit of observation is a commuting-zone-by-senior-year-by-gender cohort; the table stacks male and female cohorts. The analysis sample is limited to cohorts with image-based style measures, patenting outcomes from Bell et al. (2019), and the socioeconomic controls used in the specification, and therefore does not cover the full 1930–2010 image panel. Columns report standardized outcomes for *Applicant Share*, *Grantee Share*, and *Number of Grants*. Panel A uses the standardized norm deviation index; Panel B uses standardized style novelty. Coefficients therefore correspond to a one-standard-deviation increase in the relevant style measure. The Mean Dep. Var. reports the unstandardized dependent-variable mean in the estimation sample, multiplied by 1,000; regression coefficients are based on standardized dependent variables. All specifications include commuting-zone and senior-year fixed effects, a male dummy, and socio-demographic controls. Standard errors are clustered at the commuting-zone level. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

A Appendix Figures and Tables

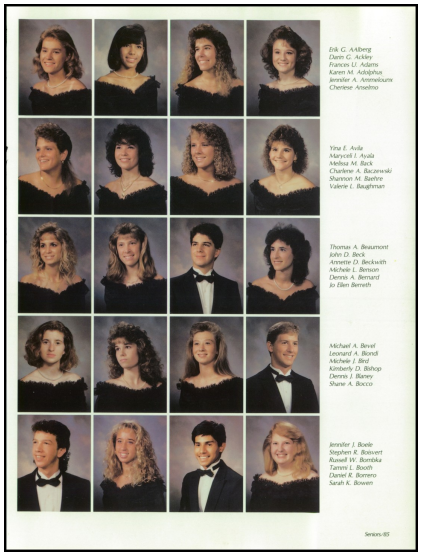
Figure A.1: Examples of Yearbook Images



Broken Arrow High School, Oklahoma, 1956



Buchser High School, California, 1970



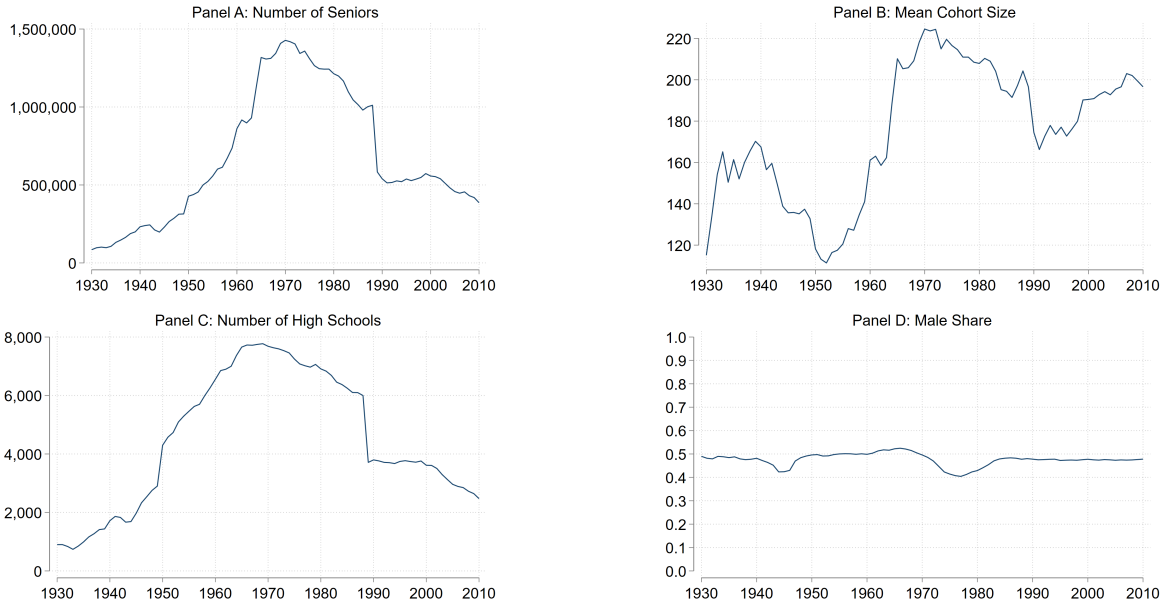
Deltona High School, Florida, 1990



Bradford High School, Arkansas, 2004

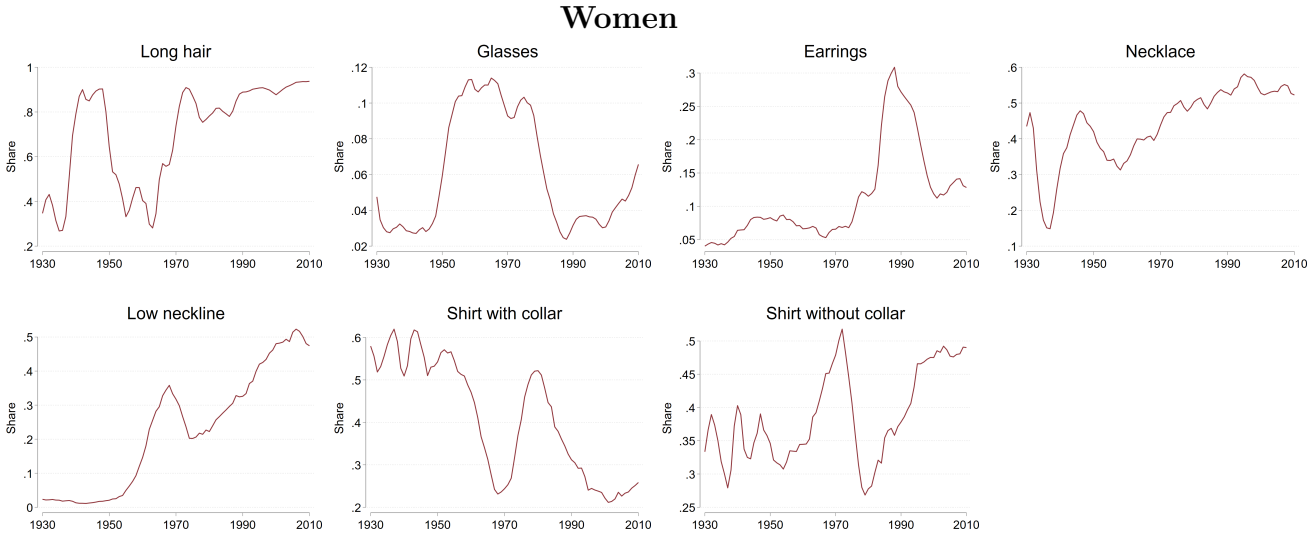
Notes: This figure shows sample pages from high school yearbooks in chronological order. The examples are taken from different years and states. The pages illustrate the format, image quality, and visual content of the yearbook materials used to construct the style measures.

Figure A.2: Sample Descriptives over Time



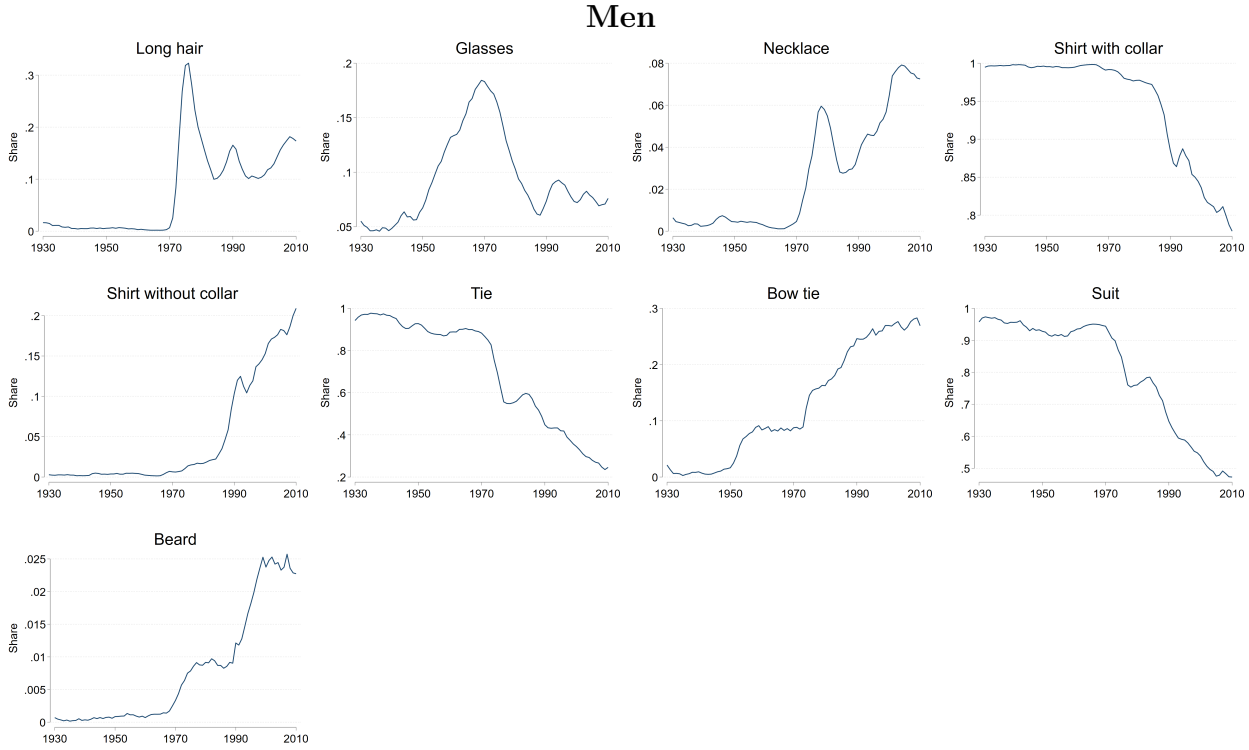
Notes: This figure reports descriptive statistics for the yearbook-image sample over time. The four panels show the total number of senior portraits, average cohort size within high-school-year cells, number of high schools observed, and share of portraits classified as male.

Figure A.3: Evolution of Style Attributes in High School Senior Yearbook Pictures



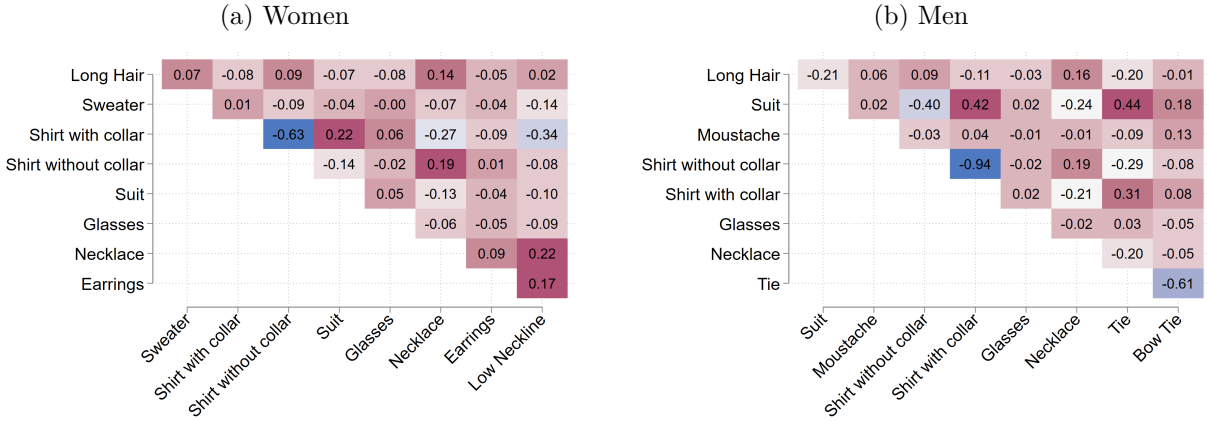
Notes: This figure plots the evolution of selected yearbook-based style attributes by gender. Each series reports the annual share of senior portraits classified as displaying the corresponding attribute. Attribute-level predictions are obtained from the Google Vertex AI classifiers and converted into binary indicators using a threshold of 0.30. The first page reports female attributes; the following page reports male attributes.

Figure A.3: Evolution of Style Attributes in High School Senior Yearbook Pictures, continued



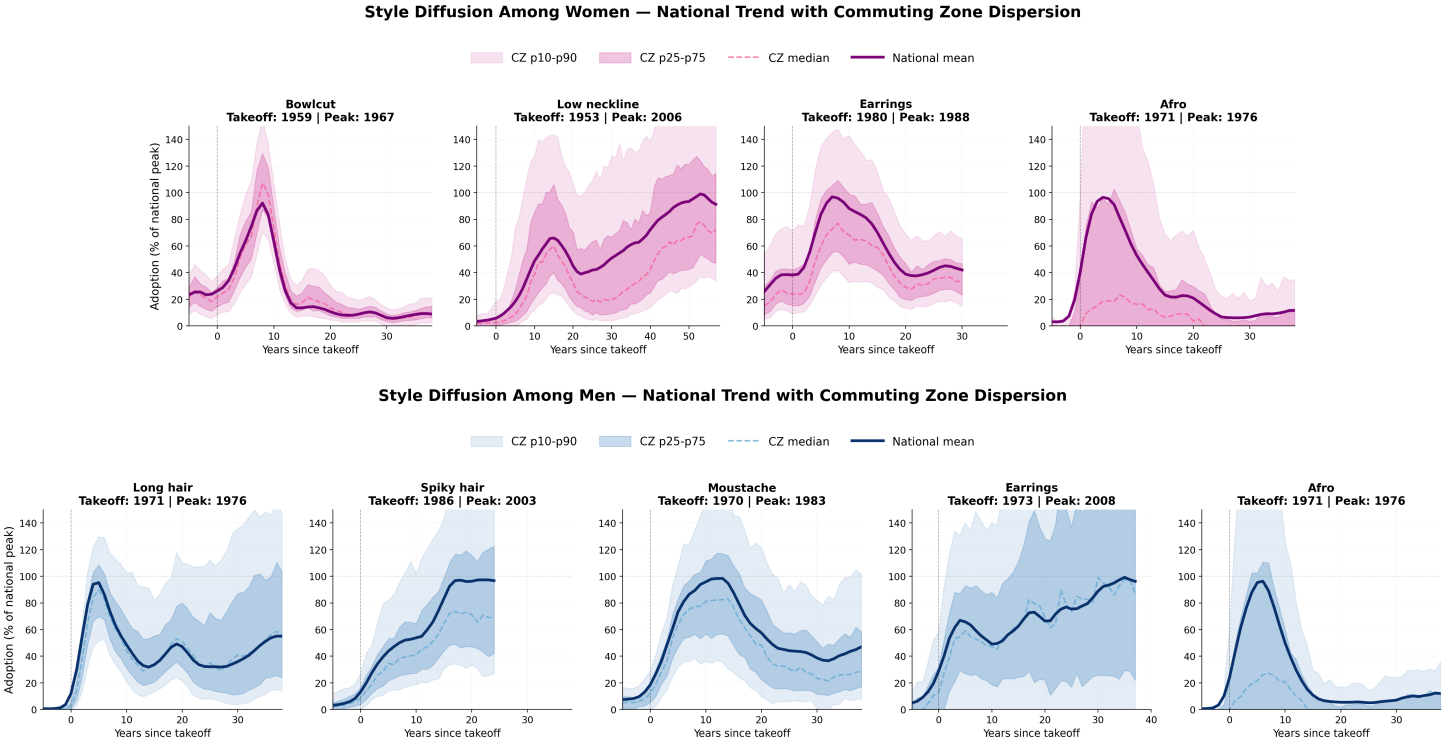
Notes: This page continues Figure A.3 and reports the corresponding male style attributes.

Figure A.4: Correlations of Style Attributes in High School Senior Yearbook Pictures



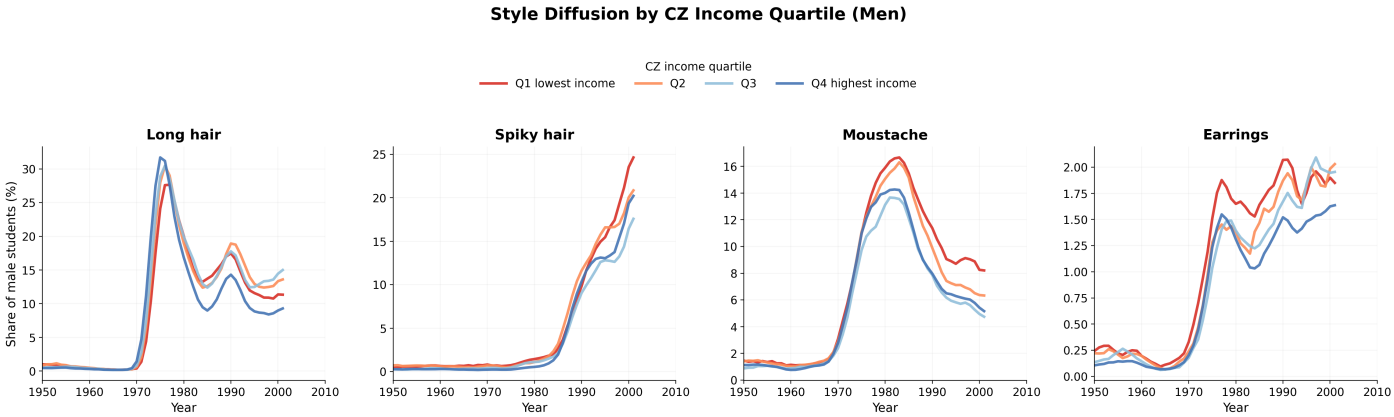
Notes: This figure reports pairwise correlations among selected yearbook-based style attributes, separately for female and male seniors. Attributes are constructed from Google Vertex AI multi-label predictions and the long-hair classifier, after converting prediction scores into binary indicators. The sample includes attributes with an average prevalence of at least 2% within gender over the period 1930–2010. Each observation is an individual senior portrait. Panel (a) reports correlations for female portraits; Panel (b) reports correlations for male portraits.

Figure A.5: Style Diffusion and Commuting-Zone Dispersion



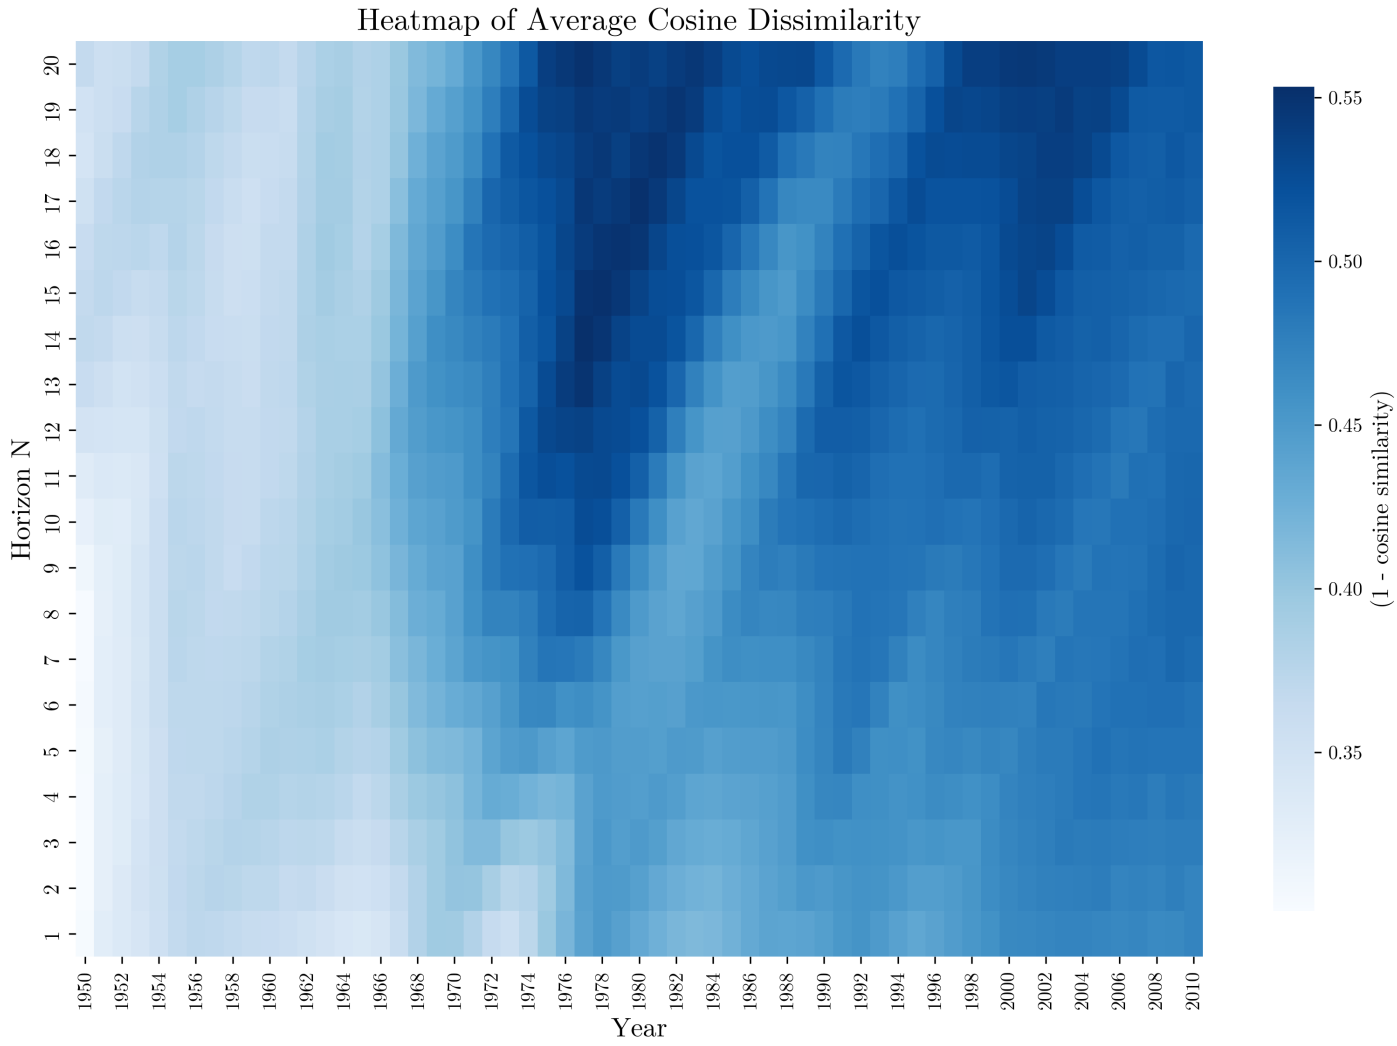
Notes: This figure shows diffusion patterns for selected style attributes around their national takeoff year, separately by gender. Adoption is normalized by the national peak adoption rate of each attribute. Solid lines report national means; dashed lines report commuting-zone medians; shaded regions show the interquartile and 10th–90th percentile ranges across commuting zones. Takeoff years are computed from smoothed national adoption series. Commuting-zone-year cells with fewer than 20 portraits are excluded.

Figure A.6: Style Diffusion by Commuting-Zone Income Quartile



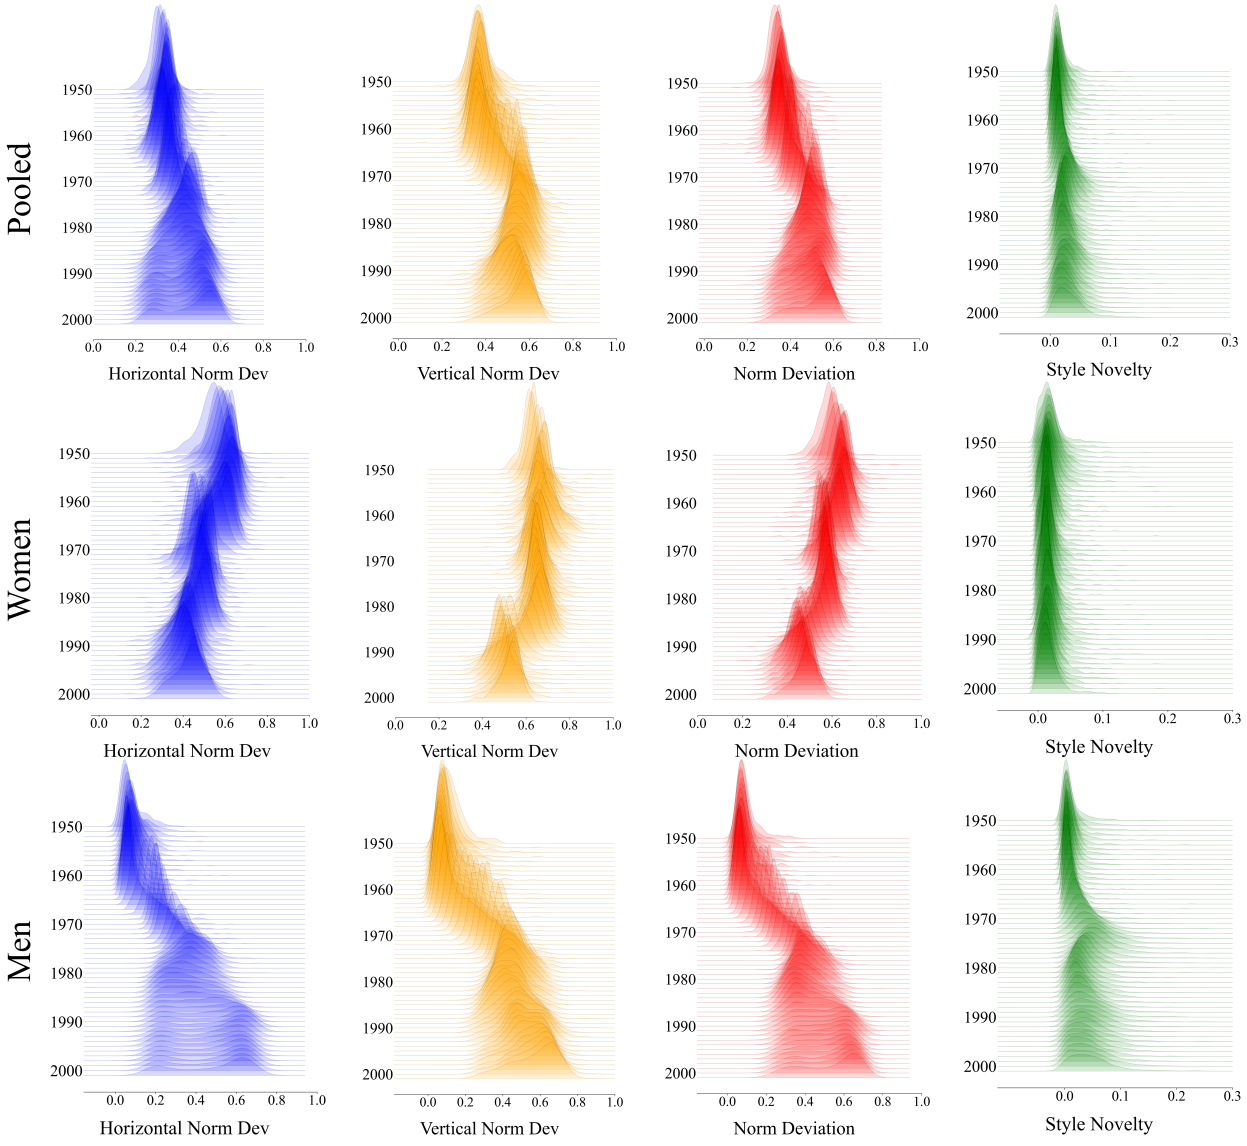
Notes: This figure reports male adoption rates of selected style attributes by commuting-zone income quartile. Quartiles are defined within year using median family income, with Q1 denoting the lowest-income commuting zones and Q4 the highest-income commuting zones. Series are smoothed annual shares of male senior portraits classified as displaying each attribute. Style attributes are constructed from yearbook-image classifier outputs and converted into binary indicators. The series run through approximately 2000; income quartiles cannot be formed for later cohorts because the underlying commuting-zone median family income data end with the 2000 Census.

Figure A.7: Vertical Norm Deviation over Time



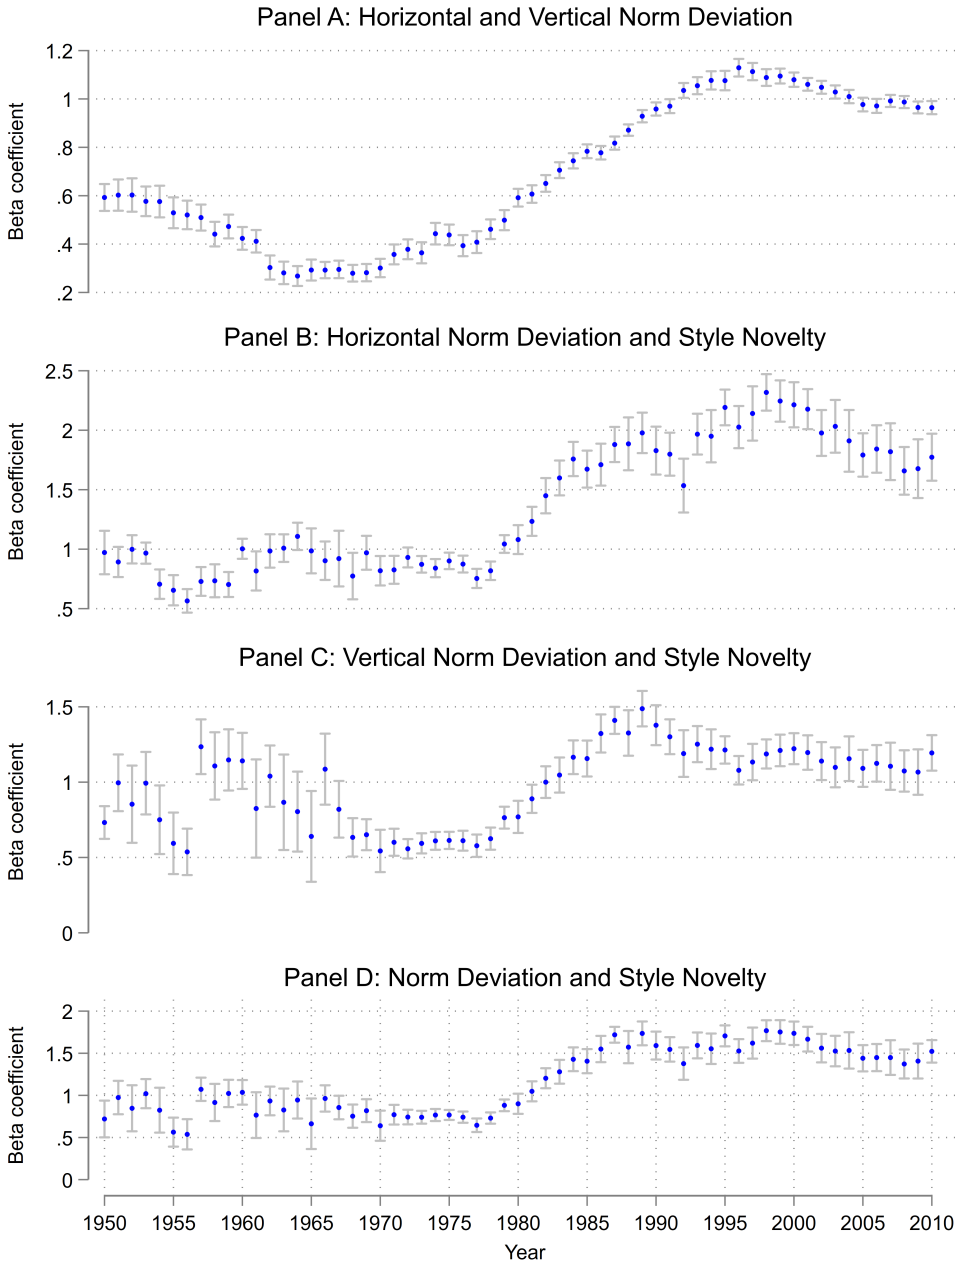
Notes: The figure shows average vertical norm deviation, defined as one minus the average cosine similarity, between high school students in each year on the x-axis and cohorts in the same high school N years before on the y-axis. Darker colors indicate higher vertical norm deviation. Vertical norm deviation is calculated separately for men and women, and then averaged. Dramatic style change, and markedly higher vertical norm deviation across horizons, is visible from 1970 onwards.

Figure A.8: Variation across Commuting Zones



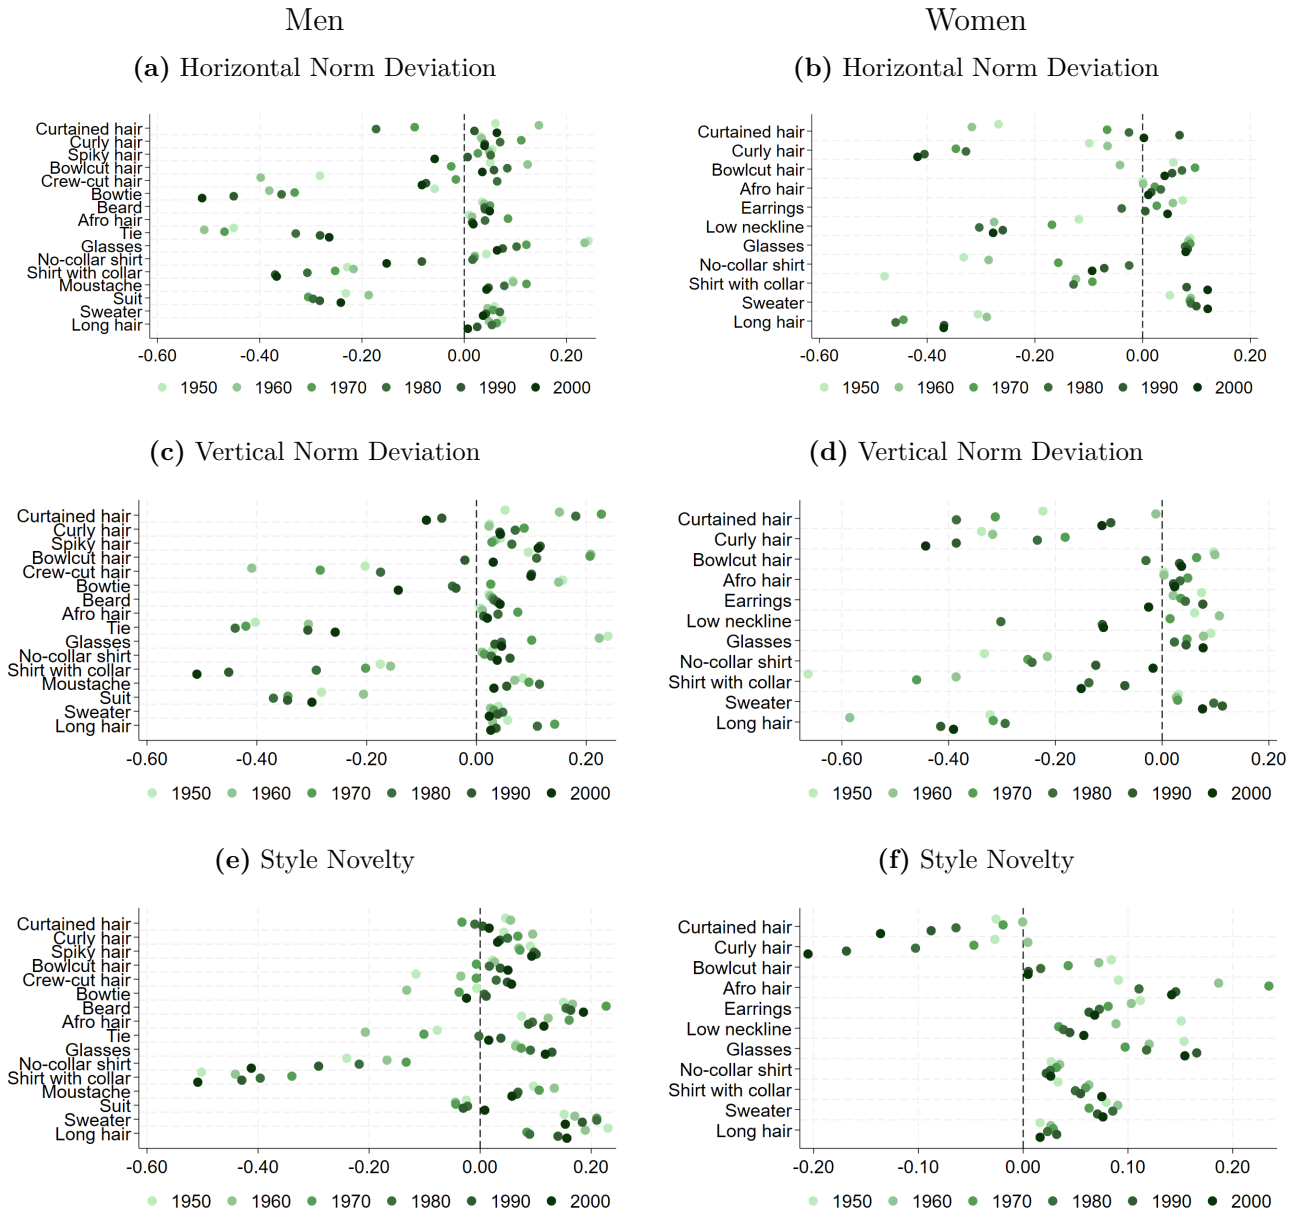
Notes: This figure plots the distribution of yearbook-based style measures across commuting zones and senior-year cohorts from 1950 to 2010. Each ridge corresponds to the cross-sectional distribution of commuting-zone-year averages in a given year. Columns report horizontal norm deviation, vertical norm deviation, the norm deviation, and style novelty. The first row pools female and male seniors; the second and third rows report gender-specific distributions for females and males, respectively. Horizontal norm deviation, vertical norm deviation, and the norm deviation are commuting-zone-year averages of the corresponding one-minus-cosine-similarity measures. Style novelty is measured as the share of seniors classified as style innovators within a commuting-zone-year cell.

Figure A.9: Correlations between Style Measures



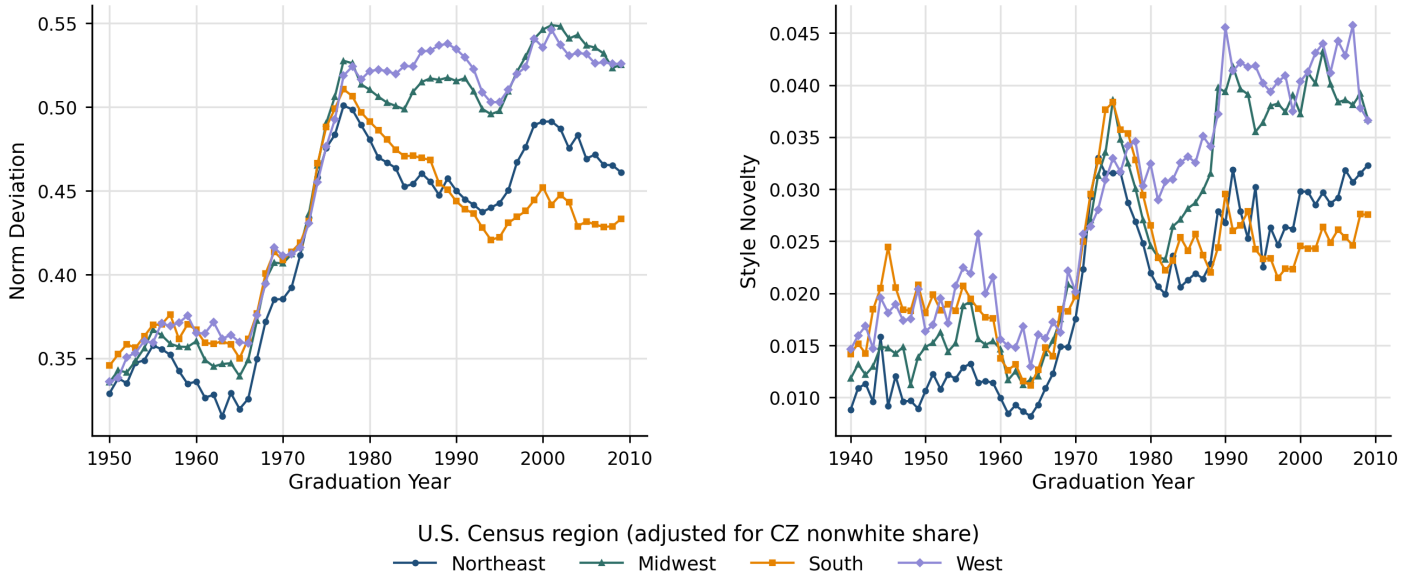
Notes: This figure reports year-by-year associations between the main yearbook-based style measures. Each point is the coefficient from a cross-sectional regression across high-school cohorts in a given graduation year. Panel A relates horizontal norm deviation to vertical norm deviation. Panel B relates horizontal norm deviation to style novelty. Panel C relates vertical norm deviation to style novelty. Panel D relates the norm deviation to style novelty. Coefficients are estimated separately by year using high-school-cohort-level observations. Vertical bars report 95% confidence intervals based on robust standard errors.

Figure A.10: Drivers of Norm Deviation and Style Novelty



Notes: This figure reports decade-specific OLS coefficients from regressions of yearbook-based style measures on style attributes, estimated separately by gender. Rows report results for horizontal norm deviation, vertical norm deviation, and style novelty. Each dot corresponds to the coefficient on an individual style attribute in a given decade. Positive values indicate that the attribute is positively associated with the corresponding style measure in that decade, while negative values indicate a negative association. For horizontal and vertical norm deviation, higher values mean larger deviations, defined as 1 minus cosine similarity; for vertical norm deviation, positive coefficients therefore correspond to lower persistence relative to the cohort twenty years earlier. The vertical dashed line marks zero.

Figure A.11: Norm Deviation and Style Novelty across U.S. Census Regions, Adjusted for CZ Nonwhite Share



Notes: This figure plots the evolution of yearbook-based norm deviation and style novelty across U.S. Census regions, adjusted for commuting-zone nonwhite population share. The left panel reports the norm deviation index; the right panel reports style novelty. Observations are commuting-zone-by-year cells, assigned to Census regions using state identifiers. Each series reports the annual regional average after residualizing the corresponding style measure with respect to the commuting-zone nonwhite population share within graduation year.[30]

[30]For each graduation year, we adjust the relevant style measure by removing the component linearly associated with commuting-zone nonwhite population share. Specifically, for each outcome, we estimate the within-year relationship between the outcome and the commuting-zone nonwhite share, and then subtract the fitted contribution of deviations from the annual mean nonwhite share. The adjusted values are then averaged by Census region and graduation year.

Table A.1: Summary Statistics for the Auxiliary Datasets

	N	Mean	St. Dev.	Min	Med	Max
Panel A: Patenting (Bell et al. 2019)						
Applicant share (per 1,000)	34,211	0.53	1.08	0.00	0.13	20.77
Grantee share (per 1,000)	34,211	0.44	0.87	0.00	0.11	17.37
Number of grants (per 1,000)	34,211	0.67	1.61	0.00	0.13	35.11
Log Population	34,211	10.27	0.88	7.82	10.15	14.74
Share Urban	34,211	0.53	0.20	0.00	0.52	1.00
Share White	34,211	0.90	0.13	0.20	0.96	1.00
Population Density (per sq. mi.)	34,211	554.29	912.47	0.50	338.93	19,984.45
Panel B: Current Population Survey						
Traditional Lifestyle Index (KLK)	1,633,366	0.01	0.61	-1.44	0.24	0.78
Log Population	1,631,226	10.56	0.95	8.31	10.48	12.67
Share Urban	1,631,226	0.69	0.15	0.20	0.70	0.97
Share White	1,631,226	0.89	0.12	0.21	0.94	1.00
Population Density (per sq. mi.)	1,631,226	682.32	691.69	2.09	487.41	3,025.92
Panel C: General Social Survey						
Countercultural Values	6,722	0.00	0.70	-1.95	0.02	1.71
Women's Autonomy	5,590	-0.01	0.82	-2.35	-0.03	1.48
Anti-Authoritarianism	6,170	0.01	0.77	-1.62	0.08	1.46
Log Population	6,724	10.70	0.86	8.66	10.56	12.67
Share Urban	6,724	0.67	0.14	0.28	0.68	0.91
Share White	6,724	0.87	0.10	0.28	0.90	1.00
Population Density (per sq. mi.)	6,724	692.16	652.06	5.61	564.08	2,940.40

Notes: Summary statistics for the three auxiliary datasets used in the paper's analyses, each on that dataset's estimation sample. Panel A (patenting; Bell et al. 2019) is the commuting-zone-by-graduation-year panel, reshaped to the person-gender level (1950–1989 cohorts); patent outcomes are reported per 1,000. Panel B is the person-level CPS ASEC; the Traditional Lifestyle Index is the KLK composite of marriage, presence of children, full-time work, and college used in the main text. Panel C is the person-level GSS (ages 18–25); Countercultural Values, Women's Autonomy, and Anti-Authoritarianism are the three KLK attitude indices. Controls are log population, share urban, share white, and population density.

Figure A.12: Patenting over the Life Cycle, by Style Novelty and Gender

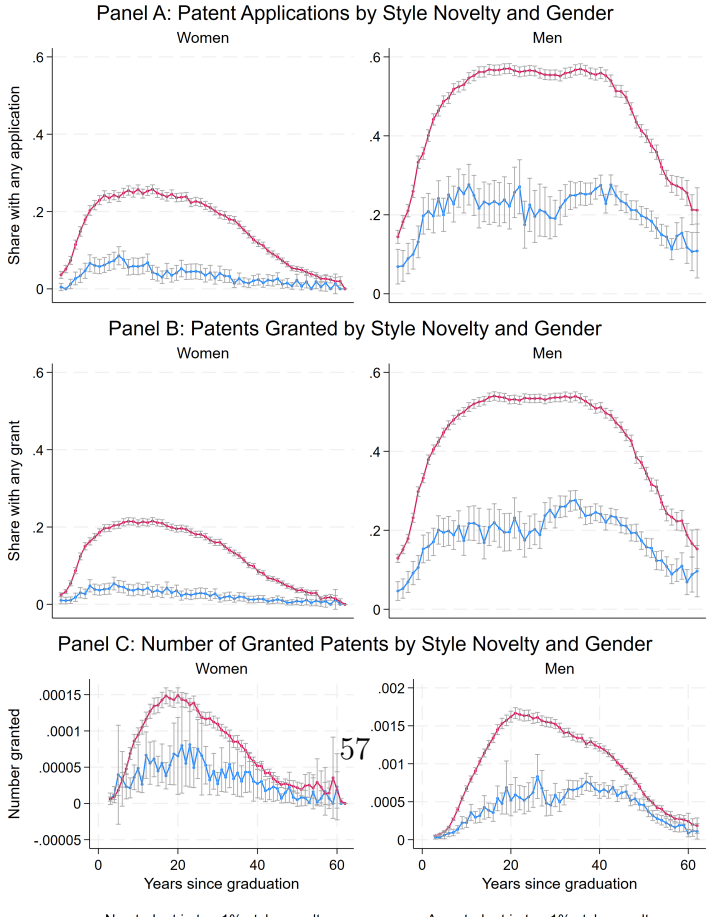
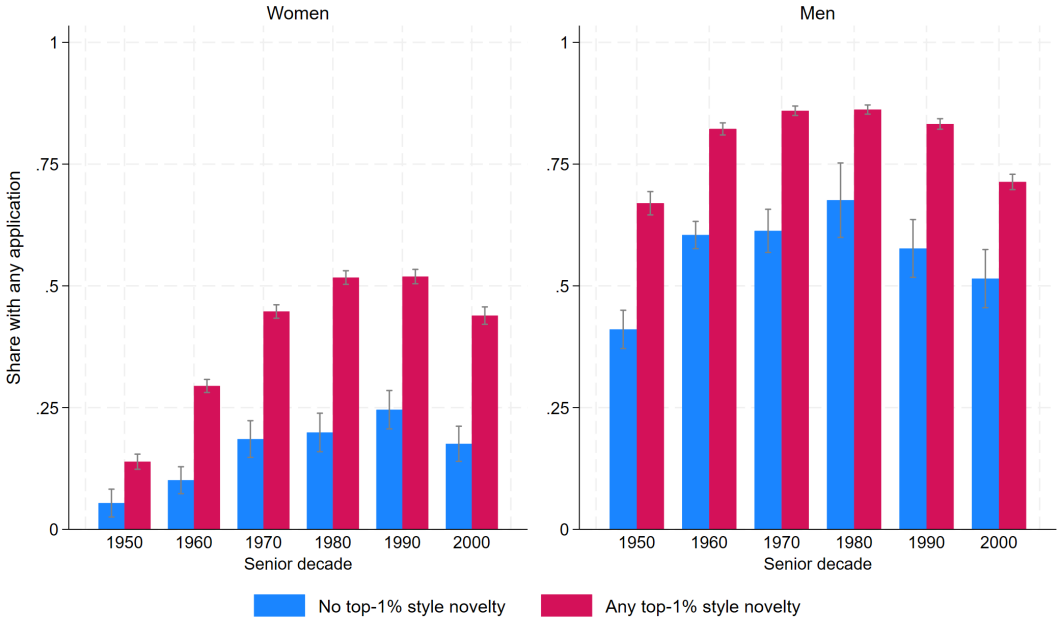


Figure A.13: Patent Applications by Decade, Style Novelty, and Gender

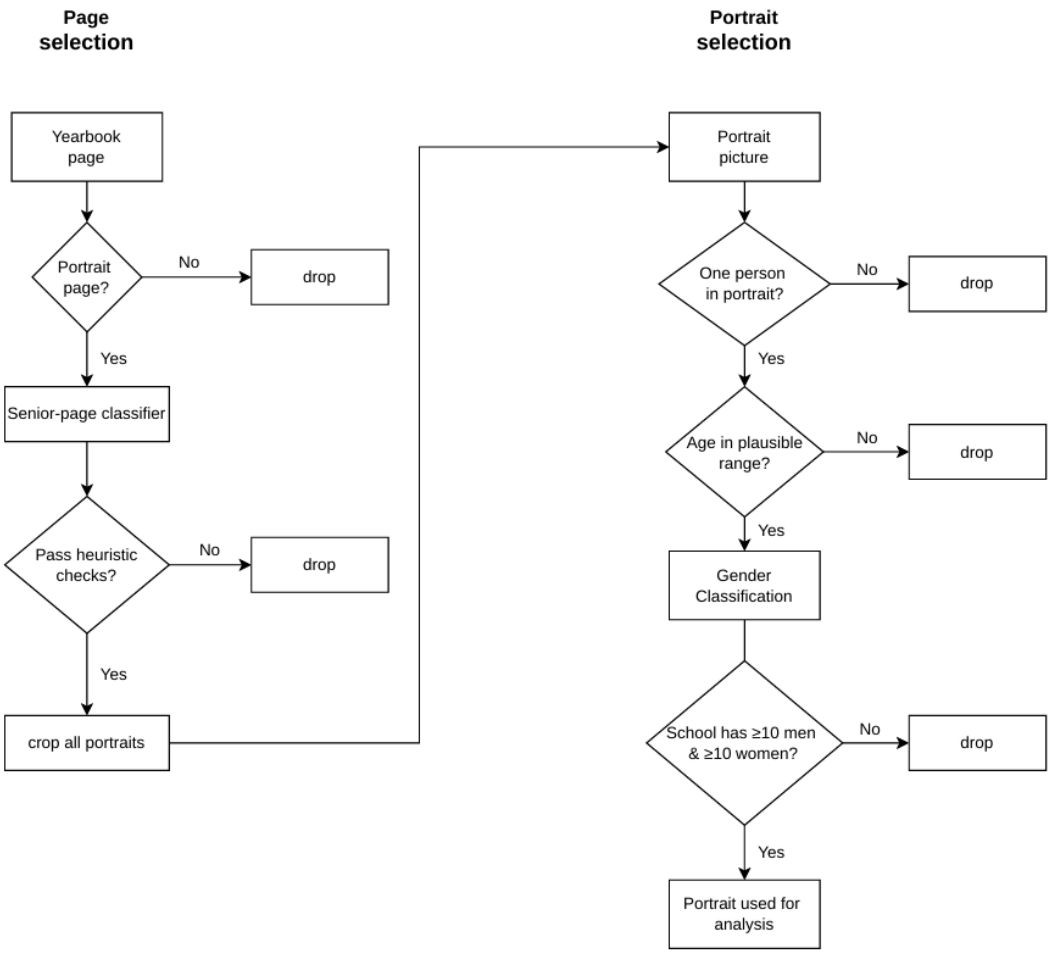


Notes: This figure reports patent application rates by senior decade, separately for male and female cohorts. The unit of observation is a commuting-zone-by-cohort cell. Bars report the share of cells with at least one patent applicant. Blue bars correspond to cells with zero gender-specific style-novelty share in the cohort year; red bars correspond to cells with positive gender-specific style-novelty share, meaning at least one style innovator. Patent outcomes are linked to high school cohorts by commuting zone and imputed year of birth, defined as graduation year minus seventeen. Vertical bars report 95% confidence intervals.

B Appendix: Database Construction

In this section, we describe the construction of our database. We first download images, identify the section containing senior pages, pinpoint where portrait pictures are on these pages, and filter to the appropriate age range. With this in hand, we classify the gender of the seniors in our portraits (Figure B.1), before proceeding to the style item classification (Appendix C).

Figure B.1: Flowcharts: Portrait Identification and Gender Classification



Notes: The flowcharts depict the steps from raw yearbook page to portrait picture used for further analysis.

B.1 Download yearbook pages from *classmates.com*

The data from *classmates.com* was acquired in two steps. First, we access classmates' "find a yearbook" query. A branching structure creates three lists of the states, cities, and high schools for which we have yearbooks. Second, based on this list of yearbooks, we iterate state by state, downloading each yearbook page image directly from the classmates.com images repository. These images are publicly available. In 2025, we expanded the dataset by incorporating yearbooks that had become publicly available on *classmates.com* but were missing from earlier download batches. The additional images originate from three sources: newly covered states, newly available high schools within already covered states, and additional yearbooks or editions for high schools already included in the sample. This update increases both the geographic coverage and the number of yearbooks observed over time. We store each yearbook with four identifying pieces of information: The state, city as listed, high school, and year. Table B.3 shows the state by state count for yearbooks downloaded, from 1930 to 2010.

B.2 Identifying portraits

The yearbook pages are processed using a portrait recognition algorithm that identifies bounding boxes for potential portraits. We then crop the detected portraits using these bounding boxes.

B.3 Selecting senior portrait pages

We identify senior portrait pages using a Google Vertex AI image classification model trained on 63,548 labeled page images drawn from more than 470 high school yearbooks, spanning the full time period of our data and covering all 50 U.S. states. To further reduce false positives, we then apply a set of conservative heuristics:

- within each school-year, we retain only a single contiguous block of pages classified as senior pages;
- among the detected portrait boxes, we retain those with dimensions consistent with the modal portrait format on the page, which helps exclude spurious or non-standard images;

- we drop pages with an unusually large number of detected portrait boxes, as these are less likely to correspond to standard senior portrait pages;
- we remove isolated positive pages and very short positive runs, which helps ensure that retained pages belong to a coherent senior section.

Table B.1 reports the resulting confusion matrix. The classifier correctly identified 5,938 out of 6,222 true senior pages (95.4%) and correctly rejected 61,292 out of 61,662 non-senior pages (99.4%). The total number of correct classifications is 67,230 out of 67,884 overall, for an aggregate accuracy of 99.04% (Table B.1). We deliberately use conservative post-classification heuristics to minimize false positives, even at the cost of accepting some false negatives. This avoids contamination of our database with images from other cohorts, even at the risk of losing a non-trivial number of observations.

Table B.1: Confusion Matrix - Senior Section Identification

		Vertex categorization		
		0	1	total
ground truth	0	61,292	370	61,662
- senior page	1	284	5,938	6,222
	total	61,576	6,308	67,884

B.4 Age classification

We used the vit-age-classifier, built on PyTorch (Raw 2023). The vast majority of our images are classified as having ages 10-29. Age classifiers trained on modern-day data tend to be upward-biased in their age estimation, especially for historical portraits (taking into account style features like suits and ties). Because of this, we impose only a mild age restriction, using images in the age range 10-49. Images below this cut-off tend to be young children; those above are typically of teachers. Table B.2 gives an overview of the age distribution in the *unfiltered* portrait data, by decade.

B.5 Gender classification

For gender classification, we adopted a two-stage process. Assigning the correct gender is crucial: a man wearing a suit and tie would not be considered innovative in 1950s America,

whereas a woman dressed in the same way would have been highly unusual, notwithstanding style icons such as Marlene Dietrich. To minimize false positives, we first classify all portraits using the Hugging Face model (prithivMLmods 2199). For observations with prediction scores below 0.9, we then use Meta Llama (2199) to resolve gender in the more ambiguous cases. To ensure that small numbers of observations do not create misleading scores, we require at least 10 men and 10 women in each graduating class.

B.6 Geocoding high school locations

To geocode the high schools in the sample, we use OpenStreetMaps API, and fill in missing information with manual searches. We find that of the public schools in our sample successfully matched to the database, 85% are still active in 2022, 9% are closed, and 6% are merged. Those not matched to the database are geo-coded by hand.

Table B.2: Predicted Age Distribution by Graduation Decade

	Predicted age group									
Graduation decade	<3	3–9	10–19	20–29	30–39	40–49	50–59	60–69	>69	Total
1930s	260 (0.0%)	32,418 (2.3%)	273,986 (19.4%)	897,328 (63.5%)	125,399 (8.9%)	56,665 (4.0%)	23,636 (1.7%)	3,347 (0.2%)	309 (0.0%)	1,413,348 (100.0%)
1940s	411 (0.0%)	61,240 (2.2%)	447,443 (16.3%)	1,823,557 (66.5%)	232,147 (8.5%)	106,582 (3.9%)	61,483 (2.2%)	7,835 (0.3%)	905 (0.0%)	2,741,603 (100.0%)
1950s	3,092 (0.1%)	173,532 (2.8%)	1,372,925 (22.5%)	3,679,599 (60.4%)	500,629 (8.2%)	209,650 (3.4%)	132,475 (2.2%)	20,490 (0.3%)	3,253 (0.1%)	6,095,645 (100.0%)
1960s	4,280 (0.0%)	560,352 (4.6%)	3,250,230 (26.7%)	6,910,361 (56.7%)	975,871 (8.0%)	337,636 (2.8%)	131,538 (1.1%)	15,704 (0.1%)	4,024 (0.0%)	12,189,996 (100.0%)
1970s	8,954 (0.1%)	344,986 (2.5%)	2,780,184 (20.3%)	9,456,396 (69.0%)	644,014 (4.7%)	269,515 (2.0%)	172,465 (1.3%)	28,010 (0.2%)	6,052 (0.0%)	13,710,576 (100.0%)
1980s	32,378 (0.3%)	475,917 (4.3%)	2,769,410 (25.0%)	6,622,737 (59.7%)	725,949 (6.5%)	238,673 (2.2%)	181,479 (1.6%)	32,273 (0.3%)	9,566 (0.1%)	11,088,382 (100.0%)
1990s	44,455 (0.8%)	199,270 (3.5%)	1,324,352 (23.4%)	3,482,113 (61.4%)	405,855 (7.2%)	135,822 (2.4%)	67,186 (1.2%)	8,886 (0.2%)	2,612 (0.0%)	5,670,551 (100.0%)
2000s	61,607 (1.2%)	152,678 (3.0%)	1,323,332 (25.7%)	3,185,916 (61.8%)	296,828 (5.8%)	88,040 (1.7%)	37,441 (0.7%)	5,262 (0.1%)	1,041 (0.0%)	5,152,145 (100.0%)
2010s	19,868 (1.1%)	52,126 (2.8%)	531,567 (28.2%)	1,153,521 (61.2%)	87,904 (4.7%)	24,110 (1.3%)	11,903 (0.6%)	2,137 (0.1%)	175 (0.0%)	1,883,311 (100.0%)
Total	175,305 (0.3%)	2,052,519 (3.4%)	14,073,429 (23.5%)	37,211,528 (62.1%)	3,994,596 (6.7%)	1,466,693 (2.4%)	819,606 (1.4%)	123,944 (0.2%)	27,937 (0.0%)	59,945,552 (100.0%)

Notes: This table reports the distribution of predicted ages for individual portraits, separately by graduation decade. Age groups are obtained from the age-classification model described in the text. Each cell reports the number of portraits assigned to the corresponding age group, with the row percentage in parentheses. The sample includes all detected portrait images before applying the final senior-age restriction used in the main analysis.

Table B.3: Yearbooks by State

State	Yearbook Count	State	Yearbook Count
Alabama	6,265	Montana	2,738
Alaska	293	Nebraska	5,561
Arizona	2,842	Nevada	1,045
Arkansas	7,555	New Hampshire	1,212
California	20,469	New Jersey	8,223
Colorado	4,692	New Mexico	1,873
Connecticut	4,682	New York	20,900
Delaware	626	North Carolina	5,567
District of Columbia	219	North Dakota	2,034
Florida	6,819	Ohio	23,484
Georgia	6,327	Oklahoma	13,285
Hawaii	393	Oregon	5,572
Idaho	1,969	Pennsylvania	24,032
Illinois	20,746	Rhode Island	1,410
Indiana	16,327	South Carolina	2,843
Iowa	9,917	South Dakota	4,569
Kansas	8,485	Tennessee	5,977
Kentucky	5,284	Texas	23,500
Louisiana	3,366	Utah	1,474
Maine	3,944	Vermont	1,500
Maryland	3,970	Virginia	5,496
Massachusetts	4,506	Washington	6,173
Michigan	16,540	West Virginia	3,008
Minnesota	9,094	Wisconsin	12,377
Mississippi	2,788	Wyoming	1,455
Missouri	10,485		
Total			363,911

Notes: This table reports the distribution of high school yearbooks in our sample across U.S. states. The sample covers yearbooks from 1930 to 2010 and is constructed from digitized yearbooks available on *classmates.com*. State-level coverage reflects both the availability of digitized yearbooks and constraints encountered during the data collection process.

C Appendix: Style Feature Classification

Images were manually labeled to create a training set. We attached labels to each image as detailed in Table C.1:

Table C.1: Style Attributes and Options

Attribute	**Options**
Gender	Female / Male
Hair	Long / Short
Tie	No Tie / Tie / Bow Tie
Clothing	Dress / Suit / Shirt without Collar / Shirt with Collar / Sweater
Mustache	Yes / No
Beard	Yes / No
Glasses	Yes / No
Earrings	Yes / No
Necklace	Yes / No
Mortar Board	Yes / No
Low Neckline	Yes / No
Hair Style	Afro / Spiky / Curtained / Bowl Cut / Crew Cut / Curly

In addition to the clothes and fashion style labels, we also want to classify hairstyle. To do so, we use three separate classifiers: a) hair length (long/short), b) multi-label hairstyle classifier (6 styles),[31] and c) a multi-label classifier for other style items. We train these classifiers on hand-labelled data. Gender is classified separately and is not one of these Vertex classifiers; that procedure is described in subsection B.5. For the hair classifier, we use 2,127 images—1,702 for training, 213 for validation, and 212 for testing. For the multi-label style classifier, we use 6,540 images—5,228 for training, 656 for validation, and 656 for testing.

Experimentation showed that a single multi-label classifier for all features had poor performance; training separate classifiers for each style item would have been prohibitive and would not result in markedly greater performance. The best combination of performance and cost involves using these three classifiers.

For each image, the classifiers return the probability of each tag occurring for every option above—for example, each image has a probability p of "glasses" and a probability $1-p$ "no

[31] Note that we combined the categories wavy and curly after experimentation.

glasses". For the hairstyle and style multi-label classifiers, we use a threshold of 0.3. The gender label, which does not use a Vertex classifier, uses a cutoff of 0.5; the decision process for gender classification is described in subsection B.5. Predictive performance and trade-off between precision and recall is given below. We do not show performance for the long-short hair classifier since it was not trained on Vertex. Its performance in the human audit sample is discussed below.

Table C.2: Validation of Style-Classification Models

Classifier	Attributes	N	Prevalence	ROC-AUC	PR-AUC	Precision@0.3	Recall@0.3
Multi-label classifier	10	13,210	0.175	0.948	0.871	0.783	0.808
Hairstyle classifier	7	9,247	0.178	0.873	0.649	0.556	0.691
Gender classifier	1	1,321	0.470	0.968	0.937	0.960	0.958

Notes: This table reports validation statistics for the three image-classification models used to construct the style measures. The multi-label classifier predicts clothing and accessory attributes; the hairstyle classifier predicts hair-related attributes; and the gender classifier predicts gender. Validation is based on manually coded images matched to model predictions. Prevalence is the share of positive labels. ROC-AUC and PR-AUC are the areas under the receiver-operating-characteristic and precision–recall curves, respectively. Precision and recall are evaluated at a 0.3 classification threshold for the style attributes and at the 0.5 operating threshold for the gender classifier, at which the gender precision and recall are essentially unchanged (0.961 and 0.955). The multi-label row covers the clothing and accessory attributes with matched hand-coded ground truth in this validation sample; a few additional attributes that enter the style vector are not separately reported in this table.

D Appendix: Human Audit

D.1 Classifier Accuracy

Instead of relying on validation performance of the classifiers using part of the training dataset, we conduct a true out-of-sample analysis by recruiting evaluators on Prolific. These evaluators are then asked to assign labels to images. We determine the modal choice of the humans first, and then compare the human labels with our classifier output. In total, we asked human evaluators to classify up to 2,500 images per attribute.[32] We find a high level of accuracy overall, ranging from 79-98% by category. Agreement among the human evaluators is higher than algorithm-human agreement, but the differences range from less than 20% (necklace) to less than 1% (moustache).

Table D.1: Out of Sample Accuracy of Feature Detection

Attribute	Sub-categories	Algo Accuracy (Specific)	% of Human Agreement	Algo Accuracy (Aggregated)	Training
Glasses	Glasses	95.98%	98.77%	96%	2500 images per category (i.e. glasses yes/no)
Ties	Tie	85.20%	98.31%	90%	1338 images per category (i.e. tie/bow tie/no tie)
	Bowtie	93.87%	98.99%		
Jewelry	Necklace	76.96%	96.53%	84%	2500 images per category (i.e. necklace yes/no)
	Earrings	90.27%	97.76%		
Hair	Curtained	76.74%	81.44%	83%	2500 images per category (i.e. short/long hair)
	Curly	67.44%	86.55%		
	Spiky	95.98%	95.98%		
	Bowl cut	94.29%	93.95%		
	Crew cut	71.46%	90.70%		
	Long hair	76.96%	87.10%		
	Afro	97.89%	97.76%		
Clothes	Shirt without collar	70.61%	87.06%	79%	1250 images per category (i.e. suit yes/no)
	Shirt with collar	72.52%	90.11%		
	Suit	79.70%	97.59%		
	Sweater	86.89%	93.53%		
	Low neckline	85.41%	90.06%		
Facial Hair	Beard	97.89%	99.58%	98%	525 images per category (i.e. beard yes/no)
	Moustache only	98.52%	98.94%		

We asked Prolific respondents to hand-label images, and compare their modal choice with the output of our Vertex classification models. The third column gives the extent of agreement, with precise percentages in the Accuracy (Specific) column. Agreement by category is in the Accuracy (Aggregated) column; agreement between humans is reported in the fourth column.

[32] We excluded mortar board from the audit as the classifier achieved 100% accuracy in our testing.

D.2 Image Similarity Perceptions

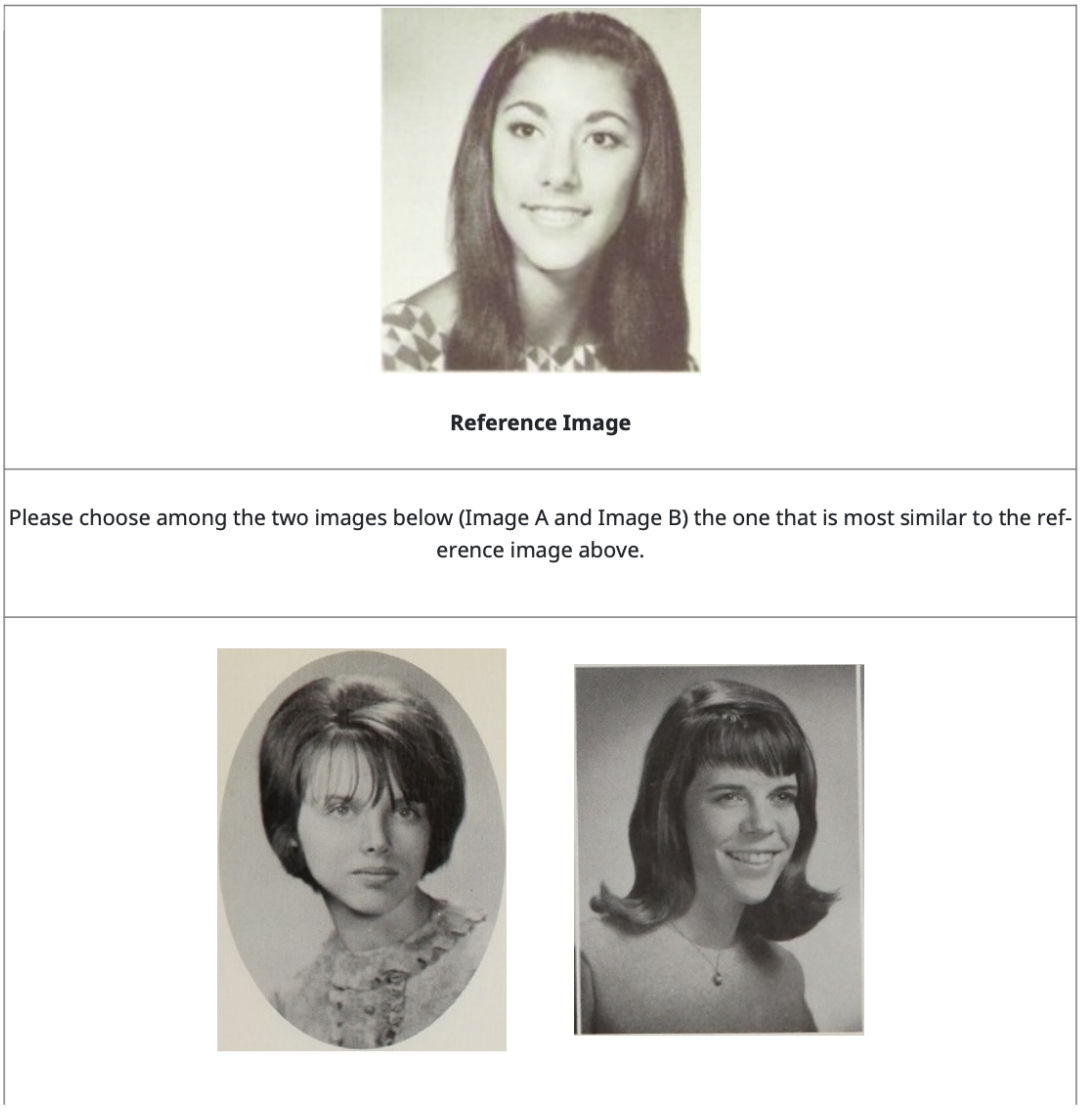
Does the style similarity as captured by our vector scoring correlate with similarity as perceived by humans? In this section, we compare the algorithm's similarity predictions with choices of human auditors. To this end, we conduct a survey using Prolific. We used 83 randomly selected images of men and women from our sample. Each participant was asked to choose which of two images, labeled Image A and Image B, most closely resembled a given reference image (Figure D.1). We then calculated the cosine similarity between the style attributes of the reference image and each of the two selected images.

How much do humans agree with each other? And how often does our procedure based on cosine similarity agree with the human choice? Figure D.2 shows the degree of agreement when we chose random pairs of images A and B, relative to a fixed reference image. Both humans and machine-human record agreement that is approximately twice as high as disagreement, with variation depending on the case.

To assess the comparative performance of the cosine similarity calculation against human judgments more rigorously, we examine how often humans/the algorithm agree with the modal choice. The procedure is detailed in Table D.2. Each row in Table D.2 corresponds to a single set of comparisons, as in Figure D.1. For each comparison set, we collected responses from 60 distinct human evaluators. Initially, we determined the mode of the human responses for each question. The mode represents the most frequently chosen image—either Image A or Image B—highlighting the collective preference of the human evaluators. We then quantified the degree of consensus among human responses by calculating the percentage of agreement with the modal choice across all evaluators for each question. This metric is documented in column *% Human Agreement*. For instance, if the modal choice is selected by 45 out of 60 human evaluators for a particular set of two-way comparisons, the *% Human Agreement* for that question would be 75%. This indicates that three-quarters of the evaluators concurred with the majority choice, suggesting a high level of agreement among participants. The last column *Algorithm Choice* shows the choice implied using the cosine similarity calculation. Here, we calculate the cosine similarity of images A and B with the reference image. The image with the higher cosine similarity score is considered the algorithm's preferred choice.

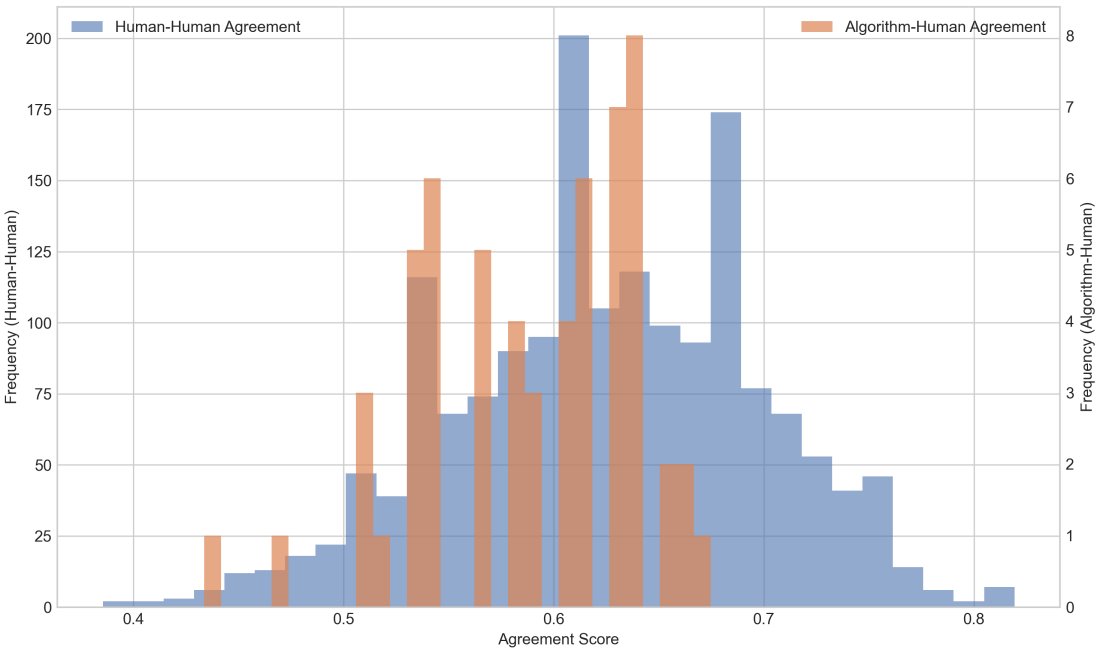
Who agrees more with human judgments? Other humans, or the algorithm? Table D.3 gives the range of agreement between humans, and the extent to which the algorithm agrees with the modal human choice. Humans do not agree all of the time, with agreement ranging from 50 to 96.6% of comparisons. The average of 73% for humans is only slightly above the value for the algorithm's agreement (66%), which falls firmly in the range of agreement

Figure D.1: Sample Screen - Prolific



Notes: The figure shows examples of reference and comparison pictures. Through the online survey provider Prolific, we recruited 60 evaluators, who were asked to judge similarity of style choice. In particular, they were asked to decide whether picture A or B is more similar to the reference image.

Figure D.2: Agreement between Humans and Humans/Algorithm



Notes: Each value on the x-axis indicates the % of answers where human-human or human-algorithm agree with each other. The height of the bars represents the frequency of agreement. The data are comprised of 1700 human-human and 1700 human-algorithm comparisons.

Table D.2: Agreement between Human Coders and Algorithm

Question code	**Human Coder**				**Modal Human**	**% Human**	**Algorithm**
	(1)	(2)	...	(60)	**Choice**	**Agreement**	**Choice**
Q02	Image B	Image B		Image B	Image B	86.67	Image A
Q03	Image B	Image B		Image B	Image B	71.67	Image B
Q04	Image B	Image B		Image B	Image B	91.67	Image A
Q06	Image A	Image A		Image A	Image B	71.67	Image B
Q07	Image B	Image B		Image B	Image B	78.33	Image A
Q08	Image A	Image A		Image A	Image A	88.33	Image A
Q09	Image A	Image B		Image A	Image A	65.00	Image A
Q10	Image B	Image A		Image B	Image B	66.67	Image B
Q11	Image B	Image A		Image A	Image A	50.00	Image A

probabilities of humans.

Table D.3: Summary: Human vs Machine Evaluations

	Accuracy Relative to Modal Human Score in %
Average Human Agreement	72.97
[min, max]	*[50.00, 96.61]*
Algorithm	66.27

E Appendix: Additional Results

E.1 Additional Results on Shocks

We also examine whether major economic and political shocks are reflected in our measures of style-based norm deviation. Specifically, we consider two episodes that plausibly affected local social and economic conditions: exposure to the China shock in the 1990s and 2000s, and Vietnam War casualties during the late 1960s and early 1970s.

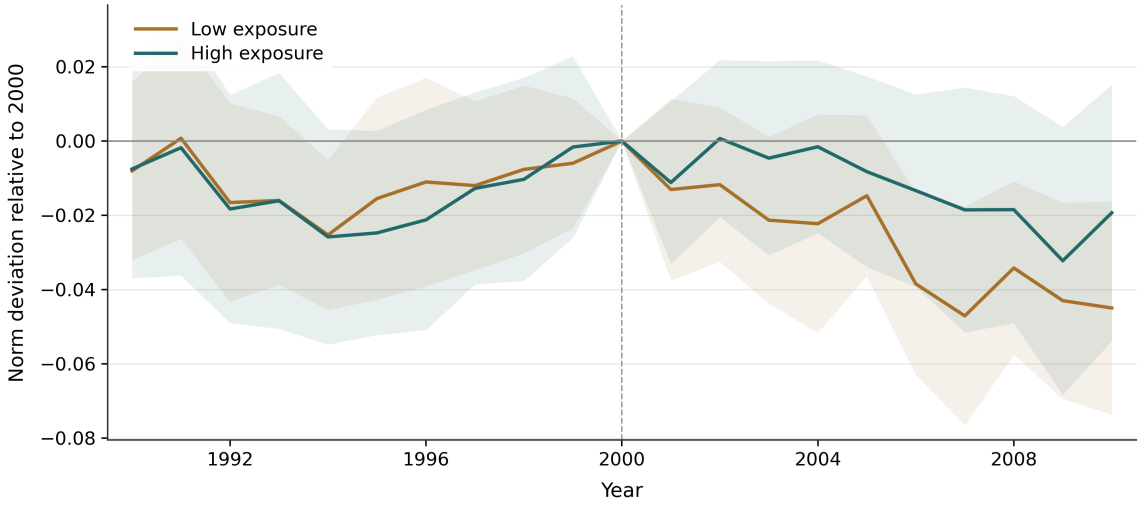
Figure E.1 compares areas with low and high exposure to these shocks. Panel (a) compares commuting zones in the bottom and top quartiles of local exposure to import competition from China (the commuting-zone import-penetration measure of Autor et al. 2013), plotting the evolution of the norm deviation index relative to its value in 2000. The two series do not display a clear break or systematic divergence around 2000. If anything, the post-2000 evolution is gradual and noisy, with confidence intervals that overlap substantially across high- and low-exposure commuting zones.

Panel (b) performs an analogous difference-in-differences-style exercise for Vietnam War casualties. We classify states by casualty intensity—using state-level death counts from the National Archives Combat Area Casualties Current File (National Archives and Records Administration 2199)—and compare low- and high-casualty states around 1964, plotting norm deviation relative to its 1964 level. Norm deviation rises sharply after the mid-1960s in both groups, consistent with the broader aggregate increase documented in the main text. However, the paths for high- and low-casualty states remain very similar, and we do not observe a sharp or persistent differential shift in high-casualty states.

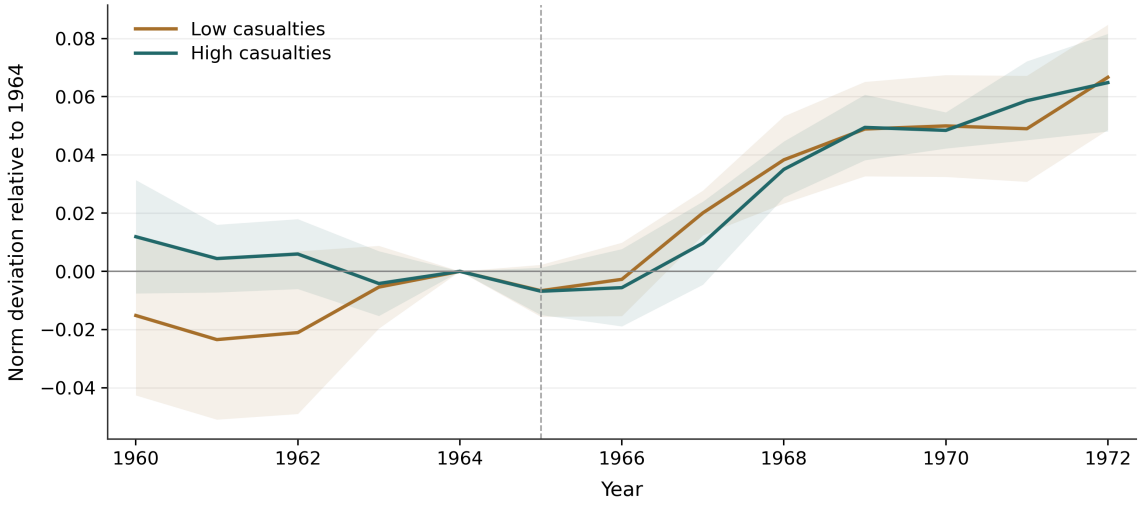
Overall, these patterns suggest that the large changes in style-based norm deviation documented in the paper are not driven mechanically by these two shocks. The evidence is descriptive and should not be interpreted as a causal test. Rather, it indicates that neither local exposure to the China shock nor Vietnam War casualty intensity is associated with large, systematic differential movements in our norm deviation measures.

Figure E.1: Norm Deviation around Economic and Political Shocks

(a) China shock exposure



(b) Vietnam War casualties



Notes: This figure examines whether major economic and political shocks are reflected in yearbook-based norm deviation. Panel (a) compares commuting zones in the bottom and top quartiles of local exposure to import competition from China, from Autor et al. (2013). The outcome is the norm deviation index, plotted relative to its value in 2000. Panel (b) compares states in the bottom and top quartiles of Vietnam War casualty intensity. The outcome is the norm deviation index, plotted relative to its value in 1964. The vertical dashed line marks 1965. Shaded areas report 95 percent confidence intervals. The evidence is descriptive and does not indicate large or systematic differential movements in norm deviation around either shock.

E.2 Trend Transitions: Within-School Homogeneity and Cross-Trend Persistence

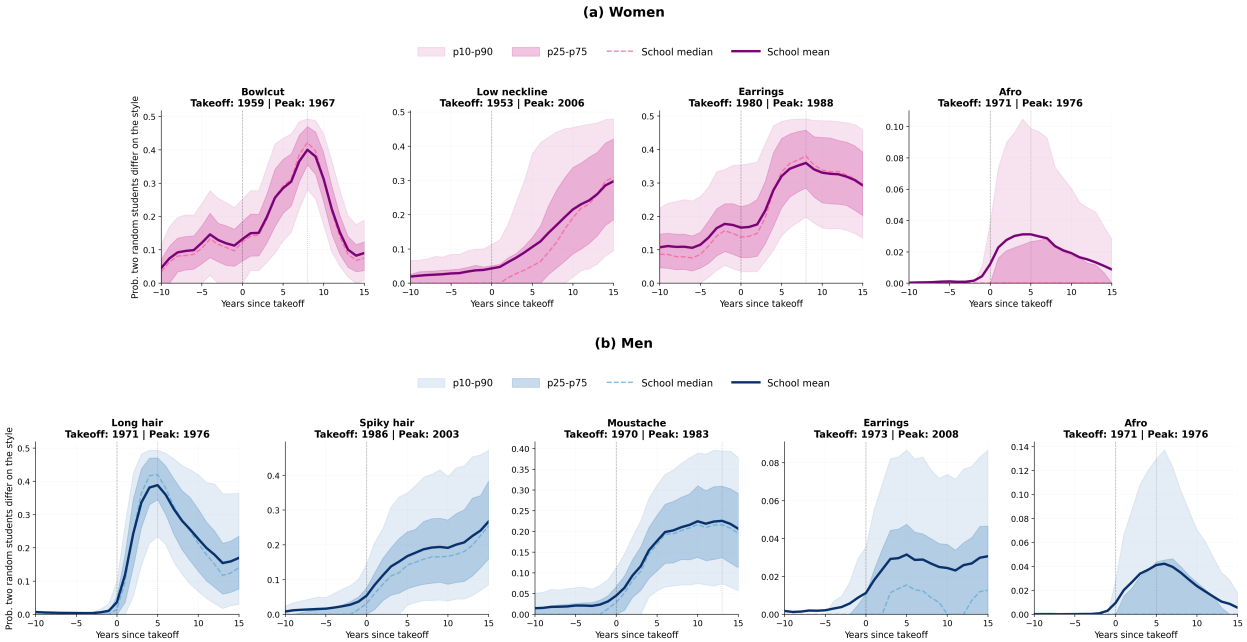
When a new style arrives and a school switches over, do seniors stay uniform, or do they diverge while the change is happening? For each style we detect the national take-off year (as in Figure A.5) and measure, school by school, the probability that two randomly chosen same-gender seniors differ on that style, $2p(1-p)$, where p is the school-year adoption share (0 = everyone agrees, 0.5 = the class is split half and half).

Figure E.2 shows the result in event time. Schools divide temporarily and then re-converge. These were minority waves, never the majority: even at its peak, men's long-hair adoption stays well below half of the class. Within-school disagreement $2p(1-p)$ rises from near-uniform (≈ 0.005) to a peak of ≈ 0.39—close to its maximum of 0.5—and falls back to ≈ 0.18 as the wave recedes; women's bowlcuts trace the same arc (0.10 to 0.40 and back to 0.16). Schools thus diverge substantially as a style takes off, then re-converge as it settles. Styles that settle at partial popularity (low neckline, mustaches) retain a moderate split; styles that never take off (afros, earrings for men) barely move.

A related question is whether adoption propensity is a persistent local trait: are some commuting zones consistently slow to adopt national trends, no matter what they are? Treating each style by gender as a trend (18 in total: the nine men's and nine women's attributes of Figure A.4, 1930–2010), we measure for every commuting zone and trend the correlation between the commuting zone's two-year change in adoption and the national two-year change. We then correlate these persistence measures across trends. Correlations are positive throughout, averaging 0.41 for women and 0.38 for men, and a single common component accounts for close to half of the variation (48% and 47%, respectively). An alternative definition based on whether a commuting zone is above or below its own mean when the nation is above or below its own yields the same conclusion.[33]

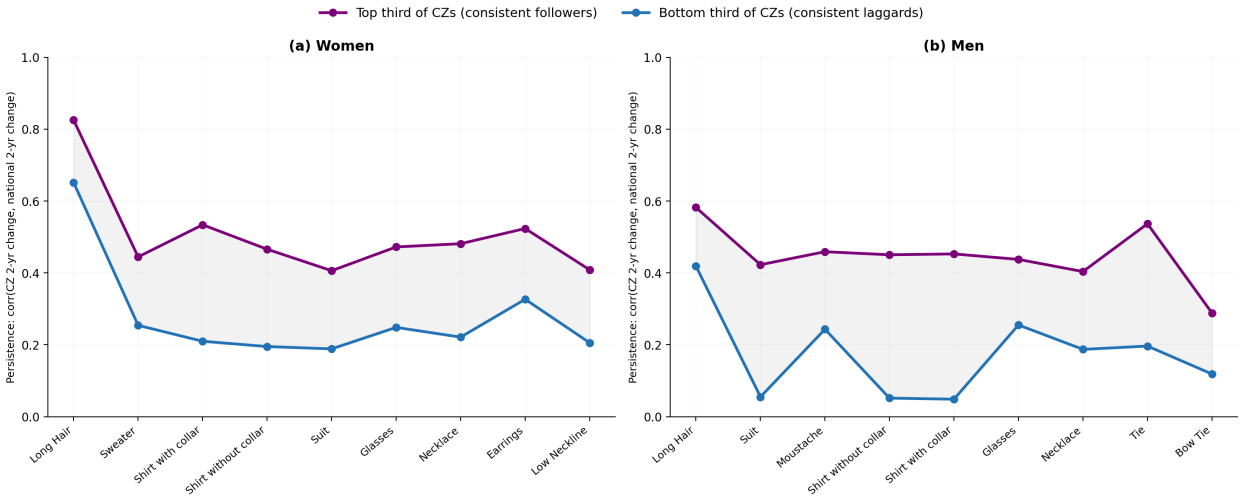
[33] Figure E.3 shows this directly: ranking commuting zones by their general co-movement with national trends (leave-one-out, so it never uses the style in question) and comparing the top and bottom third, the "follower" group tracks the nation more closely than the "laggard" group for every one of the nine styles, in both genders.

Figure E.2: Within-School Division on a Style around Its National Take-off



Notes: Event study around each style's national take-off year (take-off detection as in Figure A.5; year 0 = take-off). For each school and year, we plot the probability that two randomly chosen same-gender seniors differ on the style, $2p(1-p)$, where p is the school's adoption share (0 = all agree, 0.5 = evenly split). Lines show the mean and median across schools; bands show the p10–p90 and p25–p75 ranges. School-year cells with at least 20 individuals of the relevant gender; event-time points with at least 15 schools.

Figure E.3: Consistent Followers vs. Laggards: Some Commuting Zones Track National Trends More Closely, Whatever the Trend



Notes: Built from the per-commuting-zone persistence measures (the correlation of a commuting zone's two-year change in adoption with the national two-year change, by style attribute). We score each commuting zone's *general* tendency to follow national trends by leave-one-out: for a given style, the score is the commuting zone's average persistence over the *other* eight styles only, which removes mechanical circularity. We then split commuting zones into the top third ("consistent followers") and bottom third ("consistent laggards") on that score, and plot, for each style, the mean persistence within each group. The followers line lies above the laggards line for every style in both panels (mean gap 0.23 for women, 0.27 for men), so fast or slow adoption is a property of the place, not of the particular trend. Women (left); men (right).

E.3 "From Jerry to Terry"

E.3.1 Motivation

This section presents a simple validation exercise that examines whether individuals whose portraits diverge from local style norms also tend to have less common first names. The purpose is to assess the face validity of the style-based norm deviation measure by comparing it to a seemingly independent cultural signal: first-name distinctiveness. The results show that although name and style norm deviations co-move, the relationship is far from one-to-one.

First names are most often chosen by parents at birth and reflect cultural preferences, tastes, or identity at the time of naming. In contrast, senior-portrait styles are more likely to reflect personal choices made many years later by the individuals themselves. If our measures of style norm deviation capture meaningful departures from local stylistic norms, we would expect to observe some modest correlation with independently constructed measures of name distinctiveness, as a potential signal of individuality. At the same time, because these two domains capture different moments and different forms of identity expression, the correlation should not be large. For example, if a student's name is unusual, that may reflect a family or cultural preference for individuality. And if that same student also expresses individuality years later in how they dress, we should observe some relationship.

We evaluate the relationship for three norm-deviation measures, each constructed analogously on the name and style sides: a similarity-based measure relative to the contemporaneous state-gender cohort norm for the relevant senior year or matched SSA cohort (horizontal norm deviation), a lagged-cohort measure benchmarked against the comparable state-gender cohort 20 years earlier (vertical norm deviation), and an average of the two (norm deviation index). We note explicitly that these relationships are strictly descriptive: the goal is interpretability, not causal inference or prediction, and we make no claims about how naming practices and style choices might influence each other.

E.3.2 Data Construction

For computational tractability, we draw senior portraits from a random sample of up to 250 high-school yearbooks in each sampled year, every two years from 1950 to 2010 across all states. Full names are extracted from page images using a vision-language model to parse page layouts and identify portrait–name pairs, then cleaned to isolate first names. After portrait–name matching, similarity attachment, and filtering, this yields a matched input

sample of 752,444 senior portrait–name records with style norm-deviation scores across 7,484 yearbooks. The figures and regressions below use the complete-case validation sample of 746,744 observations with the required name-side and style-side norm-deviation measures.

Both name-side and style-side norm deviation are benchmarked within state-gender cohorts for the relevant senior year or matched SSA cohort, a choice we justify below, and which is also constrained by SSA data being released only at the state level. Style measures use the senior year, while name measures use the matched SSA cohort, proxied by senior year minus 18. We construct three measures for both style and name. On the style side, style horizontal norm deviation is defined as 1 - cosine similarity between an individual's portrait embedding and the average embedding of their state–year–gender cohort. Style vertical norm deviation is defined analogously, but the cohort reference is the comparable state-gender cohort 20 years earlier, so style vertical norm deviation captures how unusual a portrait would have looked relative to peers from a different historical moment. Style norm deviation index is the average of style horizontal norm deviation and style vertical norm deviation when both are observed.

The distinctiveness of a first name is measured using birth-count data from the U.S. Social Security Administration (SSA), which reports counts of every name given at least five times by state of birth, sex, and year of birth. Ideally, one would observe the actual distribution of names within each graduating cohort at the time the yearbooks were produced. As these data are unavailable, we rely on birth records from 18 years before graduation as a proxy. This approach cannot capture subsequent migration or individuals legally or informally changing their first names. However, cohort sizes are sufficiently large that the aggregate distribution of names is unlikely to shift meaningfully between birth and high-school graduation. Rates of in- and out-migration, as well as the low prevalence of formal name changes among minors, imply that such movements should introduce only minor perturbations to the overall name-frequency distribution. We therefore view birth-based name counts as a reasonable and stable approximation for measuring name typicality within cohorts. Because the SSA does not report names below this five-birth threshold, we assign names present in the yearbooks but absent from SSA data a count of four, the largest frequency consistent with non-reporting.

A natural question is why we benchmark name distinctiveness at the state-gender level rather than at the school-cohort level, which is the unit used to define style horizontal norm deviation elsewhere in the paper. Two reasons motivate this choice. First, the measurement objects differ fundamentally between names and styles. Style is multidimensional and decomposable: a portrait can be characterized along hairstyle, clothing, accessories, and other

visual attributes, so two students with otherwise distinct ensembles can still share components, and cosine distance in embedding space captures these partial overlaps. Names have no analogous decomposition. Two students either share a first name or they do not, and there is no meaningful intermediate notion of "similarity" between, for example, "Michael" and "Matthew." Constructing a school-level name distinctiveness measure would therefore amount to asking only how rare each student's exact name is within their immediate class, which is a much coarser object than the embedding-based style measure and discards the cultural-distribution information that makes name distinctiveness interesting in the first place. As a robustness check, however, we do construct a school-level analogue based on *phonetic* (Soundex) name similarity—which captures the intermediate resemblance that exact-name matching misses—in Appendix E.3.4.

Second, the within-school name sample is too thin to identify population-level distinctiveness even if a sensible measure could be constructed. With on average roughly 100 portraits per yearbook (about 50 per gender), a typical school cohort contains on the order of 50 distinct names per gender; the great majority appear only once, and even the most common names within a single class are observed only a handful of times. By contrast, the underlying state-gender cohort distribution often contains many hundreds of distinct names: for example, we observe 815 distinct names in the 1990 Indiana female cohort, from which only a small subset reaches any given school. A school-level frequency count therefore conflates two very different sources of rarity: names that are genuinely rare in the population, and names that are common in the population but happen to appear only once in a given yearbook. Consider the 1990 female class at Richmond High School in Richmond, Indiana, which contains 71 distinct names among 80 female seniors; Jennifer, the most common name in this state-cohort, appears only once in the class and would be coded as very distinctive under a within-school measure, alongside genuinely rare names like Sherene. Both would receive identical school-level scores, even though the former is among the most common names in the underlying population and the latter is genuinely uncommon. The SSA-cohort benchmark resolves this by drawing on the full population distribution from which each school's names were sampled.

Let y denote senior year and $b = y - 18$ the matched SSA cohort. For each state-gender SSA cohort (s, g, b), we define name horizontal norm deviation as:

$$\text{Name HND}^i_{s,g,b} = 1 - \frac{\log(c^i_{s,g,b})}{\log\Big(\max_{j \in G_{s,g,b}} c^j_{s,g,b}\Big)},$$

where $c^i_{s,g,b}$ denotes the cohort-specific count of name i in the matched SSA cohort. The measure indexes how far the log-count of a name falls below the log-count of the cohort's most frequent name, so that higher values correspond to more distinctive names. The log transformation mitigates the extreme right-skew in name counts and prevents the most common names from dominating the normalized scale.

Name vertical norm deviation is constructed analogously, but uses the SSA count of the same name in the comparable state-gender cohort 20 years earlier:

$$\text{Name VND}^i_{s,g,b} = 1 - \frac{\log(c^i_{s,g,b-20})}{\log\Big(\max_{j \in G_{s,g,b-20}} c^j_{s,g,b-20}\Big)}.$$

That is, we compare each name to the matched SSA cohort 20 years earlier, normalizing by the log-count of that earlier cohort's most frequent name. Vertical norm deviation captures how unusual a senior's name would have appeared in a previous generation rather than relative to their own peers, with higher values indicating greater cross-generational drift. Missing or zero lagged counts are replaced with four before scaling. Name norm deviation index is the average of name horizontal norm deviation and name vertical norm deviation when both components are observed.

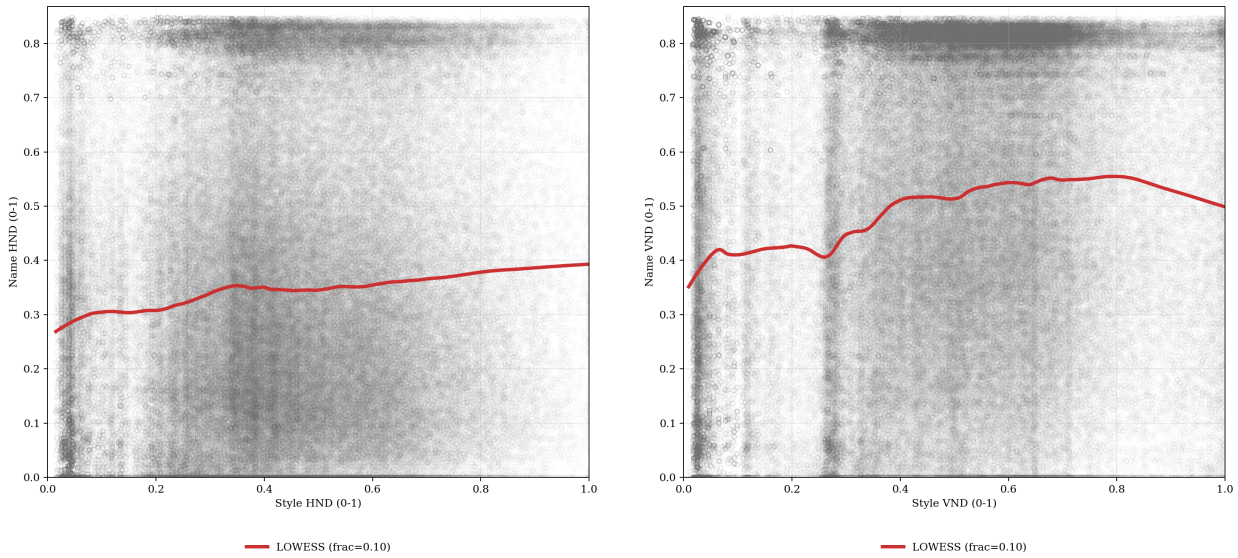
E.3.3 Results

Because the scatterplot of name horizontal norm deviation against style horizontal norm deviation contains a large number of observations and substantial overplotting, we summarize the relationship using a LOWESS smoother with a span of 0.10. Figure E.4 shows the LOWESS curves for both horizontal norm deviation and vertical norm deviation. Both smoothers indicate a clearly positive but shallow association between the corresponding name-side and style-side measures, with the vertical norm deviation curve visibly steeper than the horizontal norm deviation curve.

To quantify the magnitude of the association, we regress each name-side measure on its style-side counterpart, including additive state, senior-year, and gender fixed effects to absorb geographic, time-period, and gender differences in naming and styling practices. Because raw horizontal norm deviation, vertical norm deviation, and norm deviation index units lack a natural interpretation, we standardize both dependent and regressor variables on the complete-case row-level validation sample. Each coefficient then gives the standard-deviation movement in the name-side measure associated with a one-standard-deviation increase in

the style-side measure. We first establish the result for horizontal norm deviation across alternative fixed-effect specifications, then extend the analysis to vertical norm deviation and norm deviation index in the preferred specification.

Figure E.4: Name vs Style Norm Deviations



(a) Horizontal norm deviations

(b) Vertical norm deviations

Notes: The figure plots the relationship between name-side and style-side norm deviation measures. Panel A reports horizontal norm deviations (HND); Panel B reports vertical norm deviations (VND). Gray points show a random 200,000-observation plotting sample from the valid row-level data. Red curves are LOWESS smoothers estimated on all 746,744 valid observations with span 0.10. The sample covers 1950–2010 and all states in the final row-level dataset.

Table E.1 reports standardized estimates for horizontal norm deviation across four fixed-effect specifications. In the preferred specification (column 4), a one-standard-deviation increase in style horizontal norm deviation is associated with a 0.023 standard-deviation increase in name horizontal norm deviation. We read this as descriptive co-movement rather than evidence of a directional effect. Across columns, coefficients attenuate monotonically as fixed effects are progressively added, but they do not change sign, indicating that the relationship is not driven by compositional differences across states, time periods, or genders.

Table E.1: Name Horizontal Norm Deviation Regressed on Style Horizontal Norm Deviation, by Fixed-Effect Specification

	(1)	(2)	(3)	(4)
Dependent variable: Name HND (z-score)				
Style HND (z-score)	0.076***	0.069***	0.063***	0.023***
	(0.007)	(0.005)	(0.005)	(0.004)
	[11.395]	[13.363]	[11.815]	[5.719]
R^2	0.0058	0.0293	0.0354	0.0433
Within R^2	—	0.0048	0.0038	0.0004
Observations	746,744	746,744	746,744	746,744
State FE	No	Yes	Yes	Yes
Senior Year FE	No	No	Yes	Yes
Gender FE	No	No	No	Yes

Notes: Each column reports an OLS regression of standardized name horizontal norm deviation (HND) on standardized style HND. Columns differ in the set of additive fixed effects included; the fully controlled specification in column (4) absorbs state, senior-year, and gender means. Variables are standardized using the complete-case row-level validation sample of 746,744 observations before estimation. Standard errors clustered by state are reported in parentheses; t-statistics are reported in brackets. *Significance:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Table E.2 extends the analysis to vertical norm deviation and norm deviation index in the preferred specification. The qualitative pattern from Table E.1 holds for the other two measures: estimates are positive and statistically significant. More substantively, the table reveals a cross-measure ranking. Vertical norm deviation is the strongest descriptive correlate of name distinctiveness, with a standardized coefficient of 0.038, roughly 1.7 times the corresponding horizontal norm deviation coefficient. The same ranking appears in R^2 and within-R^2, though the absolute magnitudes remain small across all measures.

Table E.2: Name Norm Deviation Regressed on Style Norm Deviation, Across Measures

	HND	VND	NDI
Dependent variable: Name measure (z-score)			
Style measure (z-score)	0.023***	0.038***	0.032***
	(0.004)	(0.003)	(0.004)
	[5.719]	[11.926]	[8.375]
R^2	0.0433	0.0852	0.0707
Within R^2	0.0004	0.0010	0.0008
Observations	746,744	746,744	746,744
State + Senior Year + Gender FE	Yes	Yes	Yes

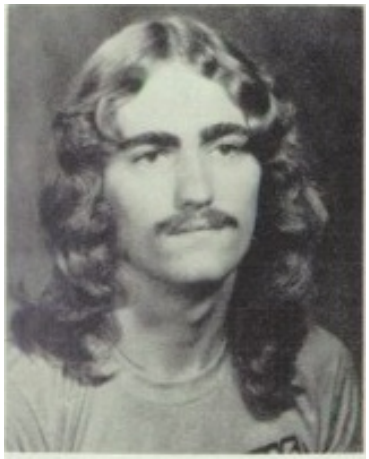
Notes: Each column reports a separate OLS regression of the indicated standardized name-side norm-deviation measure on its standardized style-side counterpart: horizontal norm deviation (HND), vertical norm deviation (VND), or norm deviation index (NDI). All regressions include additive state, senior-year, and gender fixed effects. Variables are standardized using the complete-case row-level validation sample of 746,744 observations before estimation. Standard errors clustered by state are reported in parentheses; t-statistics are reported in brackets. *Significance:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Although highly significant, the magnitude of the relationships remains small in substantive terms. Using the preferred specification, a one-standard-deviation increase in style horizontal norm deviation (0.204 raw units) is associated with only about a 0.006 increase in raw name horizontal norm deviation, or roughly 0.6 percentage points. Consider the 1986 male senior cohort from Kentucky: this difference corresponds to moving from "Jerry" to "Terry," respectively, the 28th and 29th most frequent names in that cohort out of 419 names observed. While useful as an illustrative benchmark, this shift represents only a modest increase in name distinctiveness, underscoring that stylistic and naming distinctiveness capture related but non-overlapping dimensions of cultural expression.

Figure E.5: Visual Illustration of a One-Standard-Deviation Increase in Style Horizontal Norm Deviation



(a) Center Grove High School senior
IN, 1976; male
Style horizontal norm deviation = 0.494



(b) Elwood Community Willkie High School senior
IN, 1976; male
Style horizontal norm deviation = 0.697

Notes: The figure shows two male senior portraits from Indiana in 1976. The difference in Style horizontal norm deviation between the two students is 0.203, which is close to one standard deviation of the Style horizontal norm deviation measure in the sample (0.204). The comparison illustrates the magnitude of stylistic change associated with an approximately one-standard-deviation increase in Style horizontal norm deviation.

However, the pooled relationship masks substantial decade-level heterogeneity. In State + Year + Gender FE regressions estimated separately by decade, the horizontal norm deviation coefficient is small but positive in the 1950s and 1960s, rises sharply to 0.159 in the 1970s and 0.144 in the 1980s, then attenuates to around 0.085–0.090 in the 1990s and 2000–2010 (Table E.3). The vertical norm deviation pattern is flatter, ranging from 0.014 to 0.070 across decades without a comparable mid-sample peak. All coefficients remain positive and statistically significant throughout. This pattern indicates that the link between unusual names and unusual styles is historically contingent rather than a mechanical regularity, and disciplines the pooled estimates; the 1970s and 1980s peak in HND in particular shows that the pooled value averages over distinct regimes rather than reflecting a stable underlying relationship.

Table E.3: Standardized estimates of name on style across decades

Dependent variable: Name measure (z-score)						
Regressor	1950s	1960s	1970s	1980s	1990s	2000–2010
Style HND (z-score)	0.056***	0.044***	0.159***	0.144***	0.090***	0.085***
	(0.011)	(0.008)	(0.013)	(0.008)	(0.010)	(0.011)
Style VND (z-score)	0.070***	0.014***	0.030***	0.052***	0.068***	0.055***
	(0.011)	(0.005)	(0.008)	(0.007)	(0.009)	(0.010)
Style NDI (z-score)	0.071***	0.028***	0.102***	0.113***	0.089***	0.079***
	(0.011)	(0.006)	(0.011)	(0.008)	(0.010)	(0.012)
Observations	83,889	112,504	149,234	119,722	111,400	169,995
State+Senior Year+Gender FE	Yes	Yes	Yes	Yes	Yes	Yes

Notes: Each cell reports a standardized OLS coefficient from a separate decade-specific regression. The dependent variable is the indicated standardized name-side norm-deviation measure; row labels identify the standardized style-side regressor. The three measures are: horizontal norm deviation (HND), vertical norm deviation (VND), and norm deviation index (NDI). All regressions include additive state, senior-year, and gender fixed effects. Variables are standardized using the complete-case row-level validation sample before splitting by decade. Standard errors clustered by state are reported in parentheses. *Significance:* *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

Taken together, these results imply that individuals whose style departs more from the norms of their cohort tend to have less common names, but the association remains limited. Name and style norm deviations co-move, yet style-side measures contain substantial independent variation beyond naming practices and thus provide non-redundant information about individual-level cultural distinctiveness. Overall, the evidence is consistent with the view that style and naming choices reflect related but distinct aspects of cultural individuality.

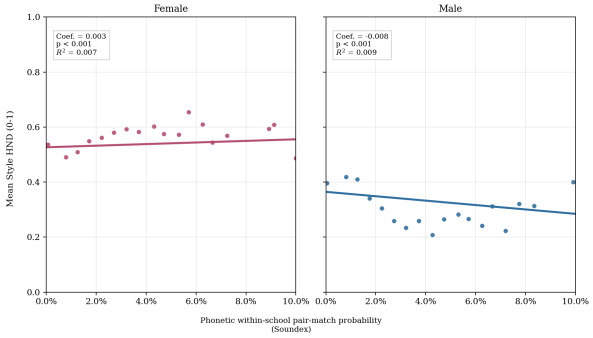
E.3.4 School-level Name Individualism

So far, we have used state-level rarity of first names because overlap in first names is rare at the high school level. Conceptually, the closest analogue to the norm deviation index for styles is to compile name similarity at the level of the graduating high school cohort, too. This we do here. Because exact matches of first names are quite rare, we use similarity based on how they sound, using the Soundex algorithm (meaning that the US version considers

Robert and Rupert identical).

Figure E.6 gives the results. We binscatter soundex match rates and horizontal norm deviations in style choices separately for men and women. While the correlation is not zero, it is very low.

Figure E.6: Name vs Style Norm Deviations



Notes: Each point is a bin mean over state-school-gender-year cells with at least two nonblank first names. The sample includes 14,815 cells: 7,415 female cells and 7,400 male cells. The x-axis is the probability that two distinct students randomly drawn from the same cell have first names with the same Soundex code. The y-axis is mean Style HND in the same cell, reported on its original 0-1 scale. Lines report separate OLS fits by gender using all cells; observations above 10 percent on the x-axis are included in the OLS estimates but omitted from the plotted range. Sample range: 1950-2010, 250 yearbooks every other year.

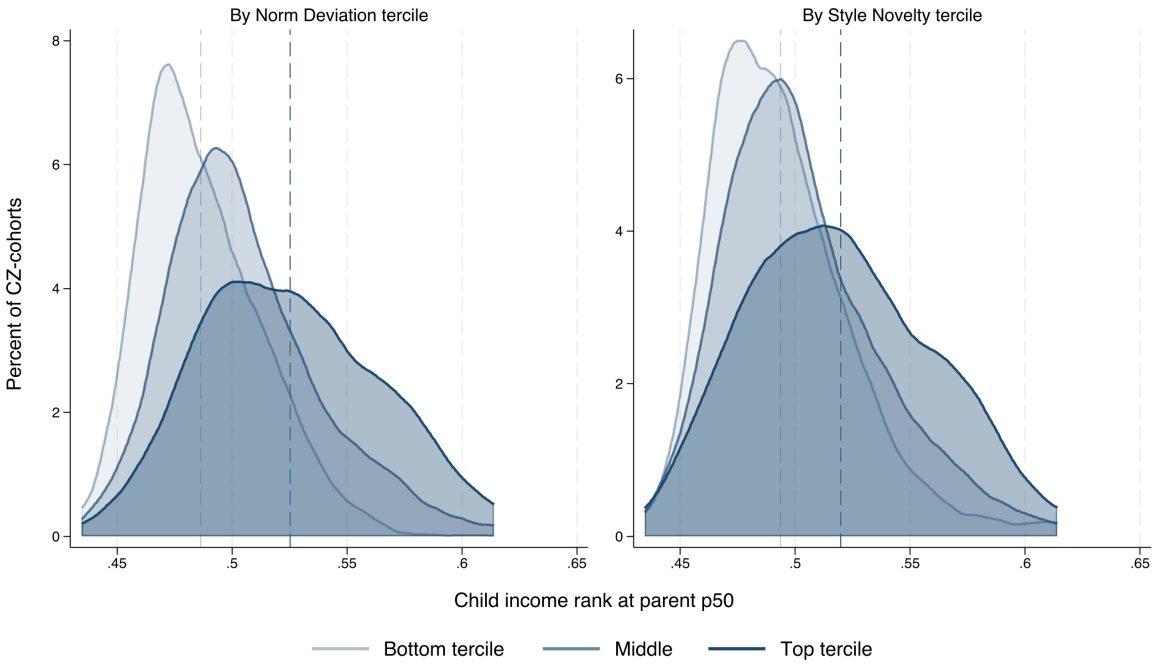
E.4 Style Measures, Mobility, Migration, and Voting

This section relates our high-school style measures to three further outcomes that lie outside the family- and labor-market margins studied in the main text: intergenerational income mobility, interstate migration, and voting. The mobility evidence is a descriptive commuting-zone comparison, whereas the migration and voting results come from individual-level regressions that follow the within-state specification used for the life-cycle outcomes, absorbing graduation-cohort, survey-year, and state fixed effects, standardizing the norm deviation index and style novelty, controlling for log population, urban share, white share, and population density, and clustering standard errors at the state $\times$ cohort level.

E.4.1 Intergenerational Mobility

We first ask whether places whose seniors departed more from local style norms also display greater upward income mobility. We merge our commuting-zone style measures to the Opportunity Insights mobility statistics (Chetty et al. 2026), taking as the outcome the mean adult income rank of children raised at the 50th percentile of the national parental income distribution, These estimates vary by commuting zone and birth cohort; we match each commuting-zone–cohort cell to our style measures for the corresponding senior class (birth cohort = graduation year $-$ 18), across the 1978–1992 birth cohorts for which the Opportunity Insights estimates are available. plots the distribution of this measure across terciles of norm deviation (left) and style novelty (right). Commuting zones in the top tercile of either measure lie visibly to the right of those in the bottom tercile, with a one-standard-deviation increase in norm deviation corresponding to roughly one additional percentile of child income rank. The pattern is descriptive and cross-sectional: it reflects differences across commuting zones rather than within-state variation, and it attenuates toward zero once state fixed effects absorb broad regional differences, so we read it as a suggestive association rather than a causal channel.

Figure E.7: Intergenerational Mobility across Terciles of Norm Deviation and Style Novelty



Notes: The outcome is the commuting-zone mean adult income rank of children raised at the 50th percentile of the national parental income distribution (pooled across child gender), taken from the Opportunity Insights commuting-zone-by-birth-cohort mobility statistics (Chetty et al. 2026). The unit of observation is the commuting zone × birth cohort (birth cohort = graduation year − 18), for the 1978–1992 cohorts covered by these estimates. Each panel shows the distribution of this measure, expressed as a percentage of commuting-zone-cohort cells, across terciles of the relevant style measure: norm deviation (left) and style novelty (right).

E.4.2 Migration

Table E.4 relates the style measures to interstate migration in the Current Population Survey, where the outcome is an indicator (scaled by 100) for living in a different state than one year earlier. Both norm deviation and style novelty predict a higher probability of an interstate move, and the association is concentrated among men: a one-standard-deviation increase in norm deviation raises the move probability by about 0.14 percentage points overall and 0.21 percentage points for men, with somewhat smaller estimates for style novelty. Because the style measure is assigned by current state of residence, movers are characterized by the norms of their destination rather than their origin; the estimates should therefore be read as an association between local norm deviation and residential mobility, not as the effect of an individual's own high-school environment.

Table E.4: Style Measures and Interstate Migration, Current Population Survey

	Interstate Move (past yr, %)		
	All	Female	Male
	(1)	(2)	(3)
Panel A: Norm Deviation			
Norm Deviation (std.)	0.142**	0.076	0.209**
	(0.064)	(0.079)	(0.090)
Observations	1,345,994	702,718	643,276
Mean Dep. Var.	3.905	3.699	4.115
R^2	0.016	0.015	0.017
Panel B: Style Novelty			
Style Novelty (std.)	0.081**	0.049	0.115**
	(0.037)	(0.044)	(0.053)
Observations	1,348,593	704,106	644,487
Mean Dep. Var.	3.904	3.699	4.114
R^2	0.016	0.015	0.017

Notes: Person-level CPS ASEC, ages 25–40, merged to yearbook style measures by high-school state × graduation cohort. The outcome is an indicator (×100) for an interstate move in the year before the survey (IPUMS `MIGRATE1`); the sample is restricted to valid movers/stayers. Independent variables are standardized. All models include graduation-cohort, survey-year, and state fixed effects and control for log population, share urban, share white, and population density; weighted by `asecwt`; standard errors clustered at state × cohort. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.4.3 Voting

Table E.5 turns to political behavior, using General Social Survey respondents aged 25–64. We consider two outcomes, each coalesced across the 1968–2012 presidential items a respondent was asked: whether the respondent voted in the most recent presidential election, and, among two-party voters, whether they voted Democratic. Neither norm deviation nor style novelty is significantly related to turnout or partisan choice, for the pooled sample or separately by gender. High-school style deviation thus tracks a range of family and labor-market outcomes but bears no detectable relationship to later electoral participation or partisanship.

Table E.5: Style Measures and Voting, General Social Survey

	Voted in Last Pres. Election (%)			Voted Democratic, two-party (%)		
	All	Female	Male	All	Female	Male
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A: Norm Deviation						
Norm Deviation (std.)	-0.906	-0.733	-1.116	-1.008	-1.600	0.089
	(0.826)	(1.101)	(1.259)	(1.053)	(1.400)	(1.626)
Observations	34,691	19,194	15,497	23,022	12,935	10,087
Mean Dep. Var.	67.932	68.901	66.875	47.792	50.745	44.440
R^2	0.052	0.052	0.062	0.042	0.056	0.046
Panel B: Style Novelty						
Style Novelty (std.)	-0.544	-0.997	-0.101	0.125	-0.476	1.007
	(0.469)	(0.621)	(0.703)	(0.585)	(0.752)	(0.907)
Observations	34,821	19,263	15,558	23,114	12,984	10,130
Mean Dep. Var.	67.967	68.942	66.905	47.776	50.718	44.440
R^2	0.052	0.053	0.062	0.041	0.057	0.045

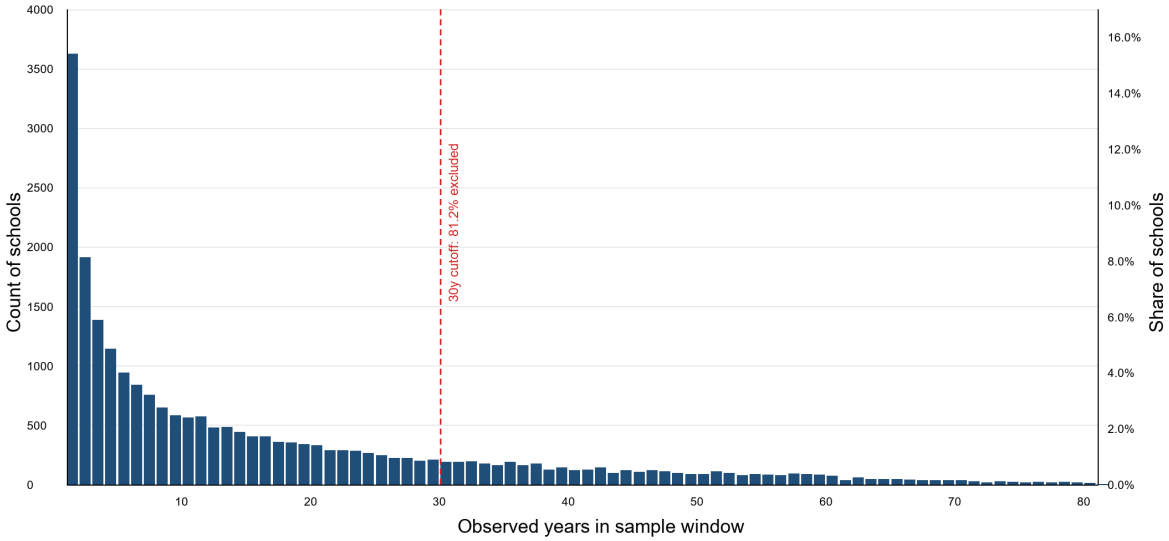
Notes: GSS respondents ages 25–64, merged to yearbook style measures by state × graduation cohort (cohort = survey year − age +17). Outcomes are coalesced across the 1968–2012 presidential items each respondent was asked: voted in the most recent presidential election (`vote##`); and, among two-party presidential voters, voting Democratic (`pres##`). Independent variables are standardized. All models include graduation-cohort, survey-year, and state fixed effects and control for log population, share urban, share white, and population density; weighted by `gss_weight`; standard errors clustered at state × cohort. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

E.5 Robustness: Long-run Schools, Uniforms, and Gender-Composition Restrictions

E.5.1 Long-run Schools

The composition of schools observed in the sample changes over time. Figure E.8 shows the distribution of the number of years over which we observe a given high school.

Figure E.8: School-year coverage in main sample.

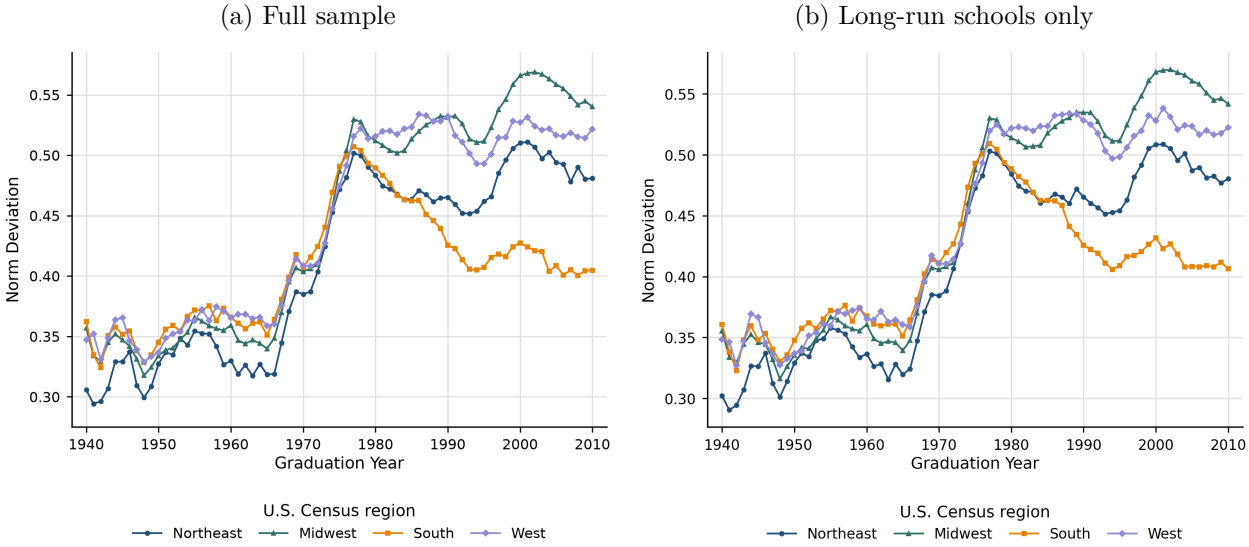


Notes: This figure presents the distribution of the number of years we observe a high school. Only 18.8% of High Schools are observed for at least 30 years.

To check the robustness of our results to changing sample composition, we perform a robustness check in which we exclude schools which we observe for less than 30 years. This ensures that the results shown in the main analysis are not a relic of additional schools entering the sample.

Out of 23,219 schools, we observe 18.8% for 30 years or longer. These 4,343 schools account for 31,136,127 (56.7%) individual yearbook portraits of our overall sample. Figure E.9 presents the results for the robustness check for norm deviation, in comparison to the main results from Figure 7. Panel (b) shows that the results do not change in their ordering of the census regions, nor in the overall or census region-specific trends over time.

Figure E.9: Norm Deviation per census region.



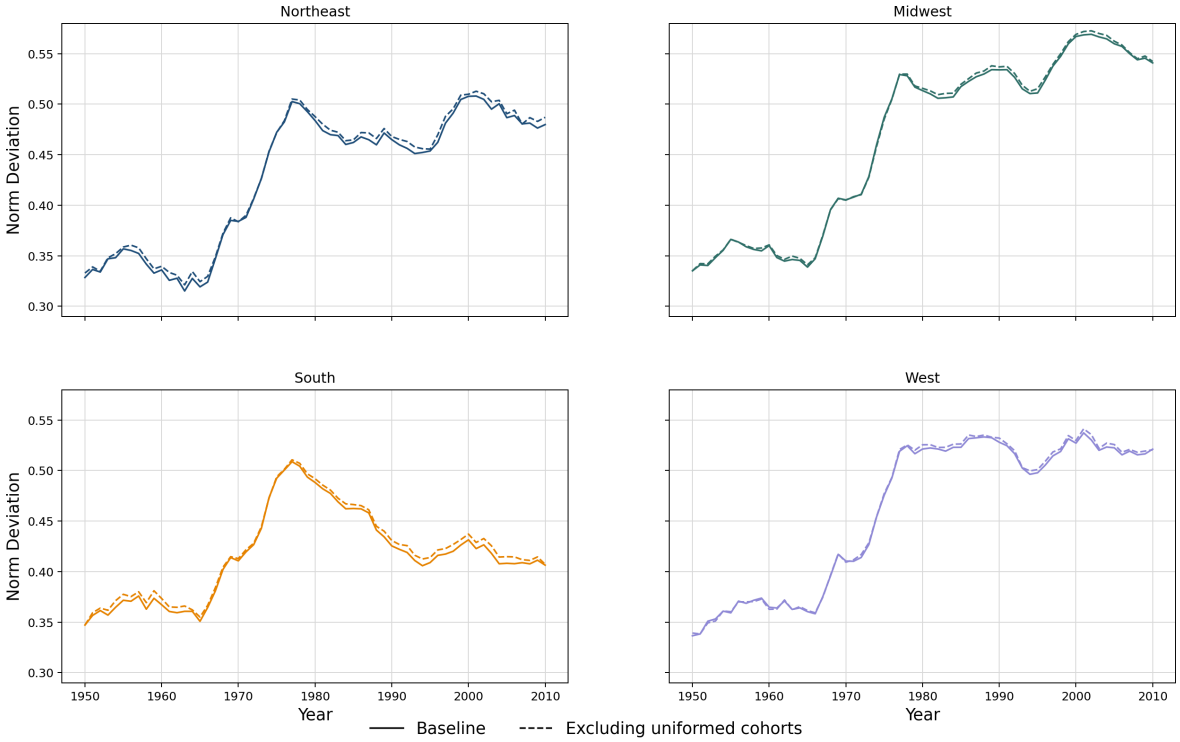
Notes: This figure shows the results for the Norm Deviation per U.S. Census Region. Panel (a) shows the main results on the full sample, as shown in Figure 7. Panel (b) shows the results excluding schools observed for less than 30 years.

E.5.2 School Uniforms

Next, we examine the effect school uniforms have on our results. To do so, we exclude those cohorts wearing school uniforms from the sample.

For each school-year-gender cell we examine the individual level apparel identifiers: bow tie, low neckline, tie, shirt without collar, mortarboard, shirt with collar, suit, and sweater. This gender cell is flagged if at least 80% of students share the same clothing. A cohort is identified to wear a uniform if that condition is true for both genders. In doing so, we exclude 6,445 (1.8%) cohorts and 932,477 (1.7%) yearbook pictures from the sample. Figure E.10 presents the norm deviation per U.S. census region for the main result and for the sample excluding cohorts wearing uniforms. It is apparent that the results do not change even if excluding observations wearing school uniforms.

Figure E.10: Norm Deviation across U.S. Census Regions: Schools without Uniforms



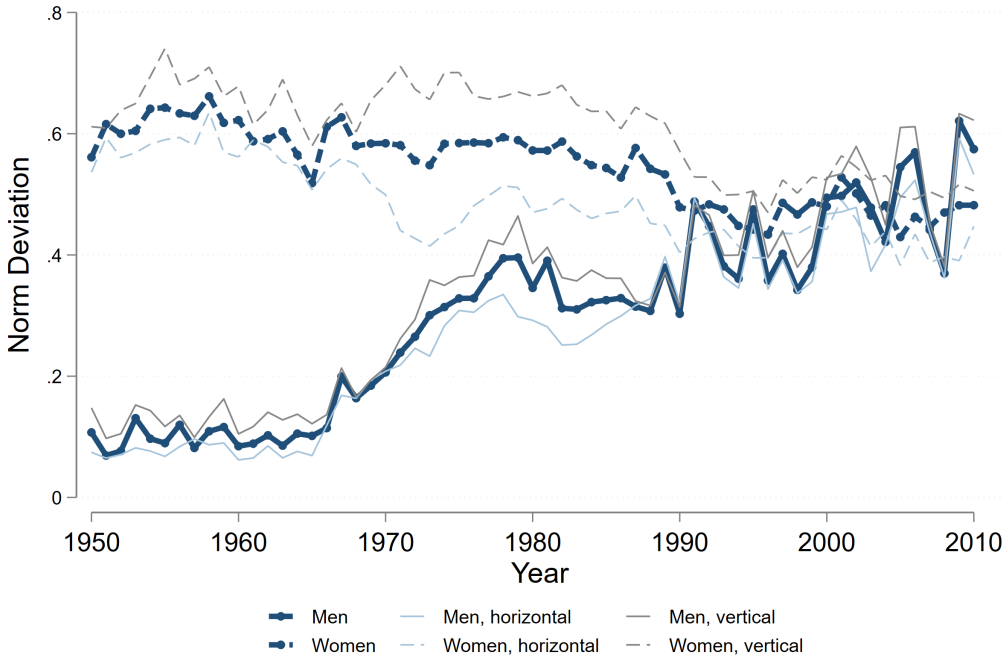
Notes: This figure presents the norm deviation trend per U.S. Census region for the full sample and when excluding cohorts with school uniforms.

E.5.3 Gender Composition

Finally, we examine a 15 percent sample of the approximately 93,000 yearbooks excluded by the gender-composition restriction. The gender-composition restriction excludes school-year cohorts that do not contain at least 10 male and 10 female seniors. The robustness check draws a 15 percent sample from the excluded yearbooks and recomputes the norm-deviation measures.[34] The resulting series, shown in Figure E.11, are essentially unchanged.

[34]We do not report style novelty for this exercise because it is defined from cumulative style frequencies. We would need to score the entire set of excluded schools to apply our method to style innovation.

Figure E.11: Norm Deviation among Yearbooks Excluded by the Gender-Composition Restriction



Notes: This figure plots norm deviation by gender in a 15 percent sample of yearbooks excluded by the gender-composition restriction. The excluded yearbooks come from school-year cohorts that do not contain at least 10 male and 10 female seniors. The figure reports horizontal norm deviation, vertical norm deviation, and their average. Style novelty is not reported because it is defined from cumulative style frequencies and is not directly comparable in this non-random subsample of excluded yearbooks.

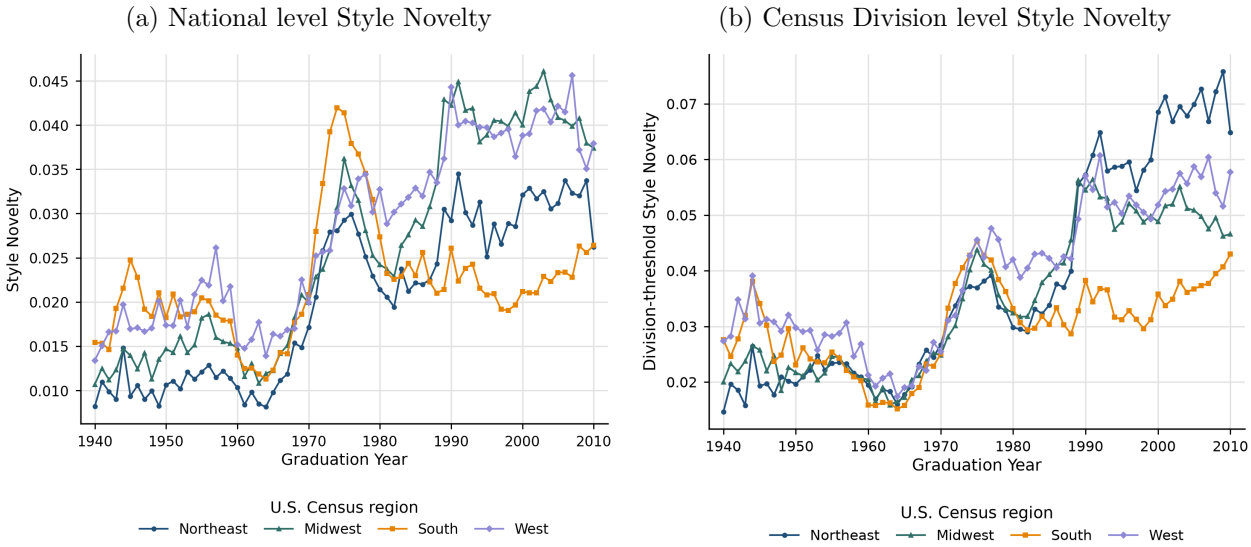
E.6 Local Style Novelty Measurement

For the main results we calculated Style Novelty by comparing styles across the nation. A concern may be that style has different dynamics more locally. To accommodate this concern, we calculated Style Novelty on the level of the nine U.S. Census divisions.[35] Specifically, we recomputed the rare-style threshold separately within each Census division, year, and gender cell, so that a style is classified as novel only if it is rare relative to other same-gender students in the same census division up to that year. We then aggregate these division-based novelty indicators to the same Census-region time series used in Figure 7. The resulting patterns in Figure E.12 are, however, different in ordering. Especially the Northeast has higher shares of novel styles when estimating novelty locally compared to nationally. Comparing the two

[35]New England, Middle Atlantic, South Atlantic, East North Central, East South Central, West North Central, West South Central, Mountain, and Pacific.

panels, the Northeast moves from being relatively non-novel under the national measure (left panel) to the region with the most novel styles when novelty is estimated locally (right panel). This development takes place especially in the late 1980s and early 1990s.

Figure E.12: Style Novelty calculated on Census Division level per Census Region.

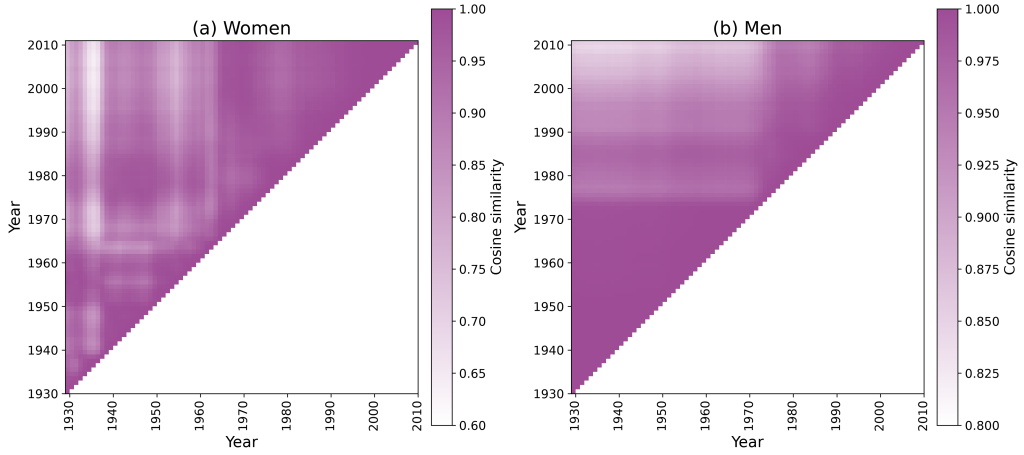


Notes: The left panel reports the baseline Style Novelty measure, where rare styles are defined relative to the national cumulative style distribution within gender and year. The right panel reports the robustness measure, where the rare-style threshold is recomputed within each Census division, gender, and year, and then aggregated to Census regions for comparability with the baseline figure. Both panels plot commuting-zone-by-year averages by U.S. Census region.

E.7 Re-Appearing Trends

Our sample also enables us to analyze the recurrence of the overall style composition over time. Figure E.13 compares, for every pair of years and separately by gender, the national mix of the nine style attributes, using cosine similarity between their adoption rates in each year. Darker blocks indicate periods in which the overall mix of style attributes was more similar across years.

Figure E.13: Year-by-year similarity of national style composition.



Notes: Style similarity for every pair of years, separately for women and men. For each year and gender, we compute the national adoption rate of the nine gender-specific style attributes and compare year pairs using cosine similarity. Darker cells indicate more similar national style composition across the two years. The matrix is symmetric, so only one triangle is shown.

The figure shows that women's styles exhibit greater variation over time than men's, with the 1930s and 1940s less similar to later decades. Men's styles, by contrast, show higher persistence and lower variation overall.