Image(s) *

Hans-Joachim Voth [†] David Yanagizawa-Drott[‡]

July 2, 2024

Abstract

From clothes and hairstyles to fashion accessories, humans use a great range of stylistic elements to express themselves, impress others, demonstrate their individualism, or show that they belong to a group. We present new methods to use images as a high-frequency, granular source for the analysis of cultural change. Despite its central importance as a form of social interaction and self-expression, and a rich body of theoretical work, empirical work on style choices is rare. We measure similarity over time and space, tracking the timing and location of influential style innovations. To illustrate our methods, we systematically exploit data from more than 14 million high school yearbook pictures of graduating US seniors to analyze persistence and change in style. We detect the collapse of high conformity as well as persistence across generations in the late 1960s. Style polarization increases sharply across commuting zones from the 1970s onwards. We also develop a novel measure of style innovation and show that it predicts patenting by cohorts later in life, suggestive of broader societal trends facilitating innovation across a range of domains. Overall, our results highlight the usefulness of images as a source for cultural economics.

Keywords: Culture, images, persistence, cultural change, similarity, counter-culture, hippies.

JEL Classification: A12, Z13, N34

*We gratefully acknowledge financial support via SNF 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, the ETH-Zurich AI workshop, Mannheim, Stanford, the 2024 UBS Advisory Board workshop in Zurich, and the NBER cultural economics meetings for feedback, and Igli Bajo, Emilia Bullano, Josep Cunquero Orts, James Godlonton, Ewan Rawcliffe, Max Read, Simone Ricci, Rossella de Sabbata, 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

Cultural economics has made great strides in recent decades, shedding new light on a range of behaviors, attitudes, and outcomes (Akerlof and Kranton 2000; Fernández and Fogli 2009; Giuliano and Nunn 2018; Guiso et al. 2016). The Cambridge Dictionary defines 'culture' as 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". Empirical studies of culture have largely focused on the second aspect. However, culture in the broader sense transcends attitudes and beliefs, and encompasses 'a way of life'. Recent work in cultural 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 as a data source (Adukia et al. 2023; Ash et al. 2022; Caprini 2023; Ludwig and Mullainathan 2024). The idea of using 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. Recent empirical evidence shows that humans react strongly, rapidly, and instinctively to images of other people (Todorov 2017).[1] Instead of analyzing faces, we focus on 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. However, they are not choices made in a vacuum, but in a social context. In fact, 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). Style choices can matter to others, have externalities, and reflect society-wide trends (Hancock et al. 2013). For example, the length of hair for young men especially can create primal reactions and has often been fought over in the courts (Graham 2004).[2] Social norms often restrict stylistic expression: Is it permissible to deviate from what

[1] A dedicated neural system in the brain allows for the processing of visual information about faces (Kanwisher et al. 2002; Haxby et al. 2000).

[2] As recently as September 2023, a black Texas high school student was suspended because he wore his hair in locks, strands of hair rolled in a circular fashion, and pinned to the head. NYT, "Black High School Student Suspended Over His Hair Length Sues Texas Leaders", September 23, 2023. Twenty-four US states have statutes banning the discrimination against students based on hair style. A classical example from popular culture is the 1960s cult movie "Easy Rider", which depicts the murder of two freedom-seeking

others wear? Does a woman need to cover her hair? Can men wear wigs or tights? 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 use recent advances in machine learning to study cultural change through the lens of stylistic choices, as reflected in portrait pictures. Our data are more than 14 million senior portrait pictures from approximately 112,000 yearbooks, taken of graduating students in thousands of US high schools during the period 1930 to 2010. Images represent high-dimensional objects; for interpretation and use in analytical applications, we need to obtain low-dimensional representations of the information they contain (Ludwig and Mullainathan 2024). We extract a vector of style attributes from each image, capturing observable choices by individuals. To do so, we train a deep neural network to identify key style characteristics, from hair length and hair style to the use of necklaces, the wearing of ties, or the depth of necklines. In this way, we obtain a sparse vector reflecting individual style choices. We then analyze similarity between style vectors to speak to issues of cultural differences, innovation, and change.

To measure how similar the style choices of two individuals are, we use the cosine similarity of the corresponding vectors. This allows us to calculate three key metrics related to culture: individualism, persistence, and style novelty. First, *individualism* captures whether local culture permits people to make choices that *differ from those of their peers*. Such an absence of conformity could reflect similarity of underlying preferences, or a high level of "norm tightness", where people fear ostracism for daring to be different (Gelfand et al. 2011). Individualism is known to vary markedly in the cross-section of individuals, across cultures, and over time (Bazzi et al. 2020; Hofstede 2011; MacFarlane 1978). We measure individualism using within-school cosine similarity for men and women, tracking its evolution at a granular level annually. This allows us to observe rapid local changes and at the national level, as well as to document how they correlate with socioeconomic factors.

Second, *persistence* captures how similar style choices in a given cohort are when compared to seniors at the same high school *a generation earlier*. We use cohorts twenty years before as our baseline comparison group. A highly persistent local culture implies high similarity over time – senior year students from the year in question will, on average, look very similar to those graduating twenty years earlier. In contrast, local cultural change implies that persistence is low. We provide evidence on how strong persistence has been across cohorts, how it differs across the country, identifying areas and time periods of rapid cultural

young men at the hands of rural rednecks for whom the heroes' hair is the main driver for visceral hate.

change. We also investigate how persistence is correlated with local socioeconomic factors such as income, race and education. Since our image data start in 1930, we begin measuring persistence in 1950, for every year until 2010, and for consistency reasons we present results on all measures for that time span.

Third, we measure innovation in the form of *style novelty*. This captures style combinations that essentially *did not exist before*. Individuals can deviate not only from previous generations in the same locality (persistence) or from their peers (individualism). They may engage in new behavior, previously not observed in the same high school, state, or country. Innovation is arguably a strong form of individualism and a signal that someone dares to be different, especially in the presence of social image concerns and if the risk of getting ostracized by peers is high. For each student's style, we ask how rare it has been and is (up to and including the present year). If fewer than one percent of students ever chose this style before (or now), we call the style combination novel. While clearly related, this metric is distinct from individualism and persistence. Someone may show high individualism (wearing a set of clothes, accessories, and a hairstyle different from their classmates) and low persistence (different from their parent's generation in the same high school), but not wear a novel style because the same style was frequently adopted in other high schools in earlier years across the country, for example.

With these three new measures in hand, we document several new facts. First, until the middle of the 1960s, culture as reflected in high school portraits was remarkably stable across time: individualism for men was low, persistence was high. For women, the opposite pattern pertained – individualism was high, and persistence, relatively low. Innovation was rare. Style choices were also 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. Intergenerational transmission was very high. Among females, we document greater variation, both over time and across space. Style novelty was more common than among males, and stable.

Starting around 1965, we begin to observe rapid cultural change. Consistent with the notion of *counterculture*, men increasingly deviated from the looks of the parental generation, as well as recent cohorts before them. Importantly, there was not a uniform shift in styles locally, which would just signal a new manifestation of conformist behavior. Instead, similarity among peers declines precipitously. At the same time, levels of individualism and persistence between men and women begin to *converge* – by the 1990s, men and women show near-identical levels. At the same time, the pace of innovation accelerates. We detect a range of novel styles never seen before in the visual record of high school images in our data. Cultural change accelerated following the Woodstock festival and the "Summer of Love" in 1968. We show that counterculture – in the form of a high level of individualism, low persistence of styles vis-à-vis the parent generation and rapid style innovation – peaked in the mid-seventies. The increase in style novelty is even more dramatic: the likelihood that someone chooses a novel style is more than three times greater in the second half of the 1970s, when innovation peaks, than before 1965.

Since the peak of counterculture, we can discern several patterns. First, after the local trough of the 1970s, persistence slightly recovered for men and individualism for men dipped dipped during the late 1970s and early 1980s, before rising once more from the mid-1980s onwards. Since the 1990s, male individualism has stabilized at a relatively high level. For women, a brief recovery of individualism in the late 1970s and 1980s gave way to another decline; at the same time, persistence among women surged. Overall, male and female style choice shows remarkable convergence over the course of the second half of the 20th century – as equality between the sexes became the norm before the law and in popular culture, style choices in two important dimensions became near-identical.

At the same time, style novelty overall increased. In fact, our most recent data from 2010 shows the highest average level of style innovation in our sample.[3]

These average trends mask a dramatic shift in spatial patterns over time. We document this shift by aggregating data to the commuting zone by cohort level. Before the mid-1960s, the country was largely homogeneous. Style expressions of high schoolers were very similar to each other, regardless of whether they lived in Alabama, New York, or Texas. In contrast, since the peak of counterculture in the mid-1970s, there is much greater spatial segregation across the country – national averages have lost predictive power for what happens in any one high school or commuting zone. In recent decades, the distribution of both individualism and persistence across commuting zones is characterized by bi-modality: some areas show low persistence and high individualism, while others combine high persistence and low individualism. This dramatic bifurcation is not easily explained by socioeconomic factors or political preferences, but forms part of a broader pattern of polarization of US society since the 1970s (Gentzkow et al. 2019).

[3] 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, orders of magnitude more than styles we actually have observed, which also helps explain why the last year of data can be the most innovative.

Next, we examine the socio-economic covariates of stylistic choice over time and in the cross-section. Overall, raw correlations show higher median household incomes correlate with greater individualism and lower persistence. A higher share of non-white individuals predicts less individualistic style choices, greater persistence, and more style novelty. This pattern, however, does not adjust for many potential unobservable factors. When we examine patterns correcting for potential unobservable factors through the use of year and commuting-zone fixed effects, we document a striking pattern – *lower* median income (or socio-economic status as measured through an index) is associated with *greater* individualism and style innovation. Persistence is also reduced in areas with lower incomes. Correlations with the share non-white mostly become insignificant.

Finally, we show that style innovation and technological innovation are closely linked. Cohorts in areas creating many new styles also file more patent applications and are granted more patents. While individualism and persistence do not strongly predict technological innovation, style innovation does. This does not mean that wearing unique clothing causes more patents. Instead, it suggests that high school environments with a loose culture that permits (or perhaps even encourages) standing out with new stylistic expressions are also home to students that invent novel solutions and products later in life.

We make several contributions to the existing literature. First and foremost, we build on the rich literature on cultural persistence and change. When fundamental economic and environmental factors change, norms of behavior can shift and reduce persistence (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). We contribute to this literature by studying style choices as a cultural phenomenon, by making use of recent advancements in machine learning. 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 education predict cultural change. Moreover, our new tools allow us to predict downstream technological innovation, a key driver of economic growth. Therefore our paper points to the possibilities of using new tools from machine learning to speak to old 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 there is no shortage of theoretical papers in economics analyzing 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 – arguably because of a dearth of high-frequency, granular data.[4]

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 concludes.

[4] 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.

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 of themselves, showing themselves 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, presenting 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 individualism, 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 would also often describe sporting events and teams, depict high school clubs and associations, portray teachers, and show ads from local firms.

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 1960s, which extended into the 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, like 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. 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, including subcultures and stylistic expressions that were on display on the popular cable television channel Music Television (MTV) - from street style to ghetto chic, and from goths to punks and 1980s power dressers. Our data, covering yearbooks until 2010, allows us to empirically establish evidence across many decades along the dimensions of individualism, persistence and style innovation.

3 Data

American high school yearbooks have a long lineage. From the early 20th century onwards, 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, and a high share of them published yearbooks containing portraits. Figure A.1 gives an example of such a publication from 1959, for Tift High School in Tift, Georgia. Most images are relatively small, and only portray 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. For the period 1930 to 2010, it contains a total of over 350,000 yearbooks. It covers thousands of high schools in 44 states from 1930 onwards. This number increases further into the 1980s, where it peaks, and is then reduced in the most recent decades. The company acquires physical yearbooks and digitize them, but its selection process is not disclosed. The data is hence unlikely to be fully representative of the population at large, whether at the national, state, or commuting zone level.

Initially, we selected yearbooks based on a simple sampling rule: for each state, we took the top 25 cities by maximum coverage (number of yearbooks available), and for these, we selected the top 3 high schools by yearbooks covered. This was done through the publicly available image repository *images.classmates.com*. We then aimed to download the universe of yearbooks available for each state on the website; however, due to apparent accessibility limits that changed during our sampling process, we were only able to do this in alphabetical order up to the state of New York. Our sample will therefore be richer for states up to the letter N in the alphabet. Table B.3 breaks down yearbooks covered by state in our sample, from year 1930 to 2010. In total, the final data contains approximately 112,000 yearbooks. While the final sample is not random and cannot be said to be representative at different geographical levels, we do not sample but use the universe of images of seniors. Moreover, when we conduct regression analysis such as correlations with socioeconomic characteristics or patents at more aggregate levels, we show results with commuting-zone fixed effects, exploiting variation within locations over time.

We then run a portrait recognition algorithm to identify in which section of a yearbook pictures of students are displayed. The algorithm scans for faces and a sequence of rectangles on the page, with a darker border compared to the background. We identify sections with seniors by using a symbolic algorithm to decide where the section for seniors begins. To this end, we look for at least four pages of consecutive images of similar size and color mix. We also require the word "senior" to be at the start of the section and exclude all sections containing the words "junior", "faculty", or "teachers". Finally, we use information on color and size to identify senior images (which are more likely to be in color and often larger). Human audits confirm we correctly identify about 99.6% of all available senior portraits. Figure A.2 shows the number of high schools covered by the data as well as the total number of available images over time. In total, we have some 14.5 million images in this national dataset. These images are then used to train classifiers (Google Vertex AI) for various style attributes such as gender, hair length, clothing, and accessories, with each classifier predicting the likelihood of each attribute through a continuous score. Model predictions from Google Vertex AI are discretized to become dummy variables using a sensible 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 a more in depth explanation on data construction.

We use two main working datasets in this paper. The first principal source of data is the 14.5 million image level data, which is the starting point for our style measures calculations. The key variables are style attributes (model predictions), individualism, persistence and style novelty for each individual. We describe how these are constructed below. The second source of data is a commuting zone level dataset that contains averages for style attributes and style measures coming from image data, along with a matched dataset on patenting behavior and innovation. Table 1 shows summary statistics for these two panels, related to style attributes and style measures (individualism, persistence, style novelty); while Table 2 describes the CZ level patenting variables available. As control variables in regression and correlation analyses we make use of *Historical, Demographic, Economic, and Social Data* from ICSPR Study 2896. These data contain socio-economic variables recorded every ten years, at the county level.

4 Method

Humans can typically judge the similarity of images instinctively and quickly (Ginosar et al. 2015). Here, we present simple methods to make such comparisons tractable and demonstrate their reliability, using a human audit. We also give an overview of main patterns in our data.

4.1 Object Identification and Image Vision

While the human eye processes images rapidly, we require a systematic approach to evaluate image similarity. One potential approach is pixel-based. It uses so-called quality measures of images, which are then compared through indicators such as the structural similarity index measure, SSIM (Wang et al. 2004). While it can be implemented easily, it is strongly affected by changing technology (color film, lens characteristics, printing technology) and photographic style (depth of field, contour or frontal lighting, etc.). Results also suffer from a lack of interpretability. The second approach, which we favor, 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 individual identifiable features, reflecting *style choices*. This approach has the advantage of being easily interpretable. Even sparse vectors allow us to distinguish thousands of different "styles" (i.e. combinations of style elements). One limitation of our source is that portraits only depict seniors head and upper; fashion change below the waistline (changes from long skirts to miniskirts, say) are not captured.[5]

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 MTurk workers to

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

make 18,878 image-category decisions, using a set of 6,540 randomly chosen images from our high school corpus for style classification. For gender classification, we used 2,188 images, and for the hair style classifier, 2,127 images.

We then trained convolutional neural networks (CNNs), using Google Vertex AI, to extract style features, using two multi-label classifiers.[6] We then evaluate the predictive performance of our algorithms.

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 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 moustache identification, marking the best result in our dataset; at the opposite end of the spectrum, our algorithm agrees 70.6% of the time with the (median) human codings. Across categories, we average 79% (clothes) to 98% (facial hair) out-of-sample accuracy.

4.2 Descriptives and Trends of Style Attributes

Figure 2 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 less than 10% of the sample by 2010. The share of 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 returning to single-digit levels. Clean-shaven portraits are dominant in the 1950s and 60s, and their share falls in the 1970s, before bouncing back from the mid-eighties 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.3 shows the correlation pattern of style attributes. Suits are positively correlated with ties and shirts with a collar, as one

[6] In addition, we use two single-label algorithms. One is 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), and the other is a gender classifier. Further details of our procedure are in Appendix C.

would expect for the case of men. Long hair is remarkably uncorrelated 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 combine these style characteristics into unique combinations and examine their evolution over time. In Figure 3, we show the composition of styles in our sample for men and women. In the 1940s, around 85% of men were in the same group – suit, tie, shirt with a collar, no glasses, short hair, no jewelry, and were clean-shaven (style 2496). By the 1960s, the share of suit and tie had declined to around 75%; another 8-10% had short hair, suit, bow tie, collared shirt, and no facial hair (style 3406). Until the late 1960s, four styles alone account for nearly 90% of the sample. By the early seventies, the influence of the more traditional styles (2496 and 3406) 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 "other" (style 999999) is one of the most common types in our sample by the 1970s already, and counts as one of the three most important categories 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, fragmentation of styles increased sharply over time – from a minority of each class that fell outside the top few categories, to one of the largest single groups.

4.3 Individualism, Persistence and Style Novelty

We calculate three key metrics designed to capture visual culture. They are based on the notion of similarity relative to a particular reference group.

Individualism. We calculate *individualism* in a school s, of gender g, in year y, as:

$$S_i = \frac{\sum_i^N \left(\frac{1}{|C|} \sum_{j \epsilon C} \frac{v_i \times v_j}{||v_i||||v_j||} \right)}{N} \tag{1}$$

where v_i is the style vector for individual i; C represents the set of style vectors for all other individuals in the same school in year y; $|C|$ is the cardinality of set C, excluding individual i, and N is the total size of the cohort. This formula calculates the average cosine similarity of individual i's features against those of their same gender peers within the same high school cohort, providing a measure of intra-cohort homogeneity.

Figure 4 illustrates the approach using senior portraits from Attica High School in New York. 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 calculate the first image's style vector cosine similarity with all others in this cohort, and then iterate over all images and take the average. Cosine similarity takes value 1 if two vectors are identical, and value 0 if they are orthogonal. Here, the first image has an individualism score of 0.098, indicating that style choices (tie, suit, hair length) are mostly identical among men in the same cohort.

Persistence. We define persistence analogously, but now use the set of all style vectors of individuals from the same institution from twenty years earlier as the comparison group C. 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 5 illustrates the persistence calculation, comparing two high schools – East Providence High School in Rhode Island, and De La Salle High School in New Orleans. In East Providence, comparing the cohorts of 1966 and 1986, we see a lot of change. Suits, ties, and short hair have given way to long hair, open shirts, and no ties. The persistence score is 0.056, indicating little similarity across generations.

Now compare this result with the pattern among young men in Panel B, who attended De La Salle HS in New Orleans. All wear bow ties, both in the 1960s and the 1980s, a tuxedo with a white shirt, no glasses, and no earrings. While hair is shorter in the 1960s, and bow ties are more exuberant in the 1980s, the similarity is unmistakable. Using cosine similarity, we derive a persistence score of 0.83, capturing the markedly higher level of similarity in style choices over time than in East Providence.

Style Novelty. To calculate our third measure, style novelty, we compare the style choices of each high school senior with those of all graduates who came before them plus the current one (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 the current generations' style choices. To this end, we calculate the count of high school senior images in all preceding years that use the same style, as a total. 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 takeup frequency of style s for gender g in year y. We then determine whether a style ranks in the top 1%:

$$I_{s,g,t} = \begin{cases} 1 & \text{if } C_{s,g,y}(t_0, t) \leq Q_1\left(\{C_{s',g'}(t_0, t) \mid s' \in \mathcal{S}, g' \in \{m, f\}\}\right) \\ 0 & \text{otherwise} \end{cases} \tag{3}$$

where $\mathcal{S}$ is the set of all styles ever observed, and Q_1 indicates the cut-off value for the top 1% of the (expanding) cumulative style distribution, ranked from lowest to highest. Figure 6 illustrates our method. The style on the left is the first example of its kind in our dataset. For the men, it includes long wavy hair, glasses, bow tie, a tuxedo, and a mustache. The first example we find in California in 1970. By 1974, it has diffused more widely, but it is still rare – choosing it places an individual in the top 1% of cumulative rarity in style choice. By 1978, however, the style has spread so widely that it no longer qualifies for the top 1%.

Not all choices that affect senior portraits are necessarily made at the level of the individual. Consider Figure 5. Clearly, the men's choice of bow tie and tie may not independent of each other, 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.

4.4 Validation

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 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,[7] 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.3 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. For some images, humans agree only 40% of the time, and for others, more than 80%. Human-machine agreement is similarly variable. When we compare agreement of individual humans or of the algorithmic prediction with the median human assessments, we find an average agreement rate among human coders is 72.97%, only slightly higher than the 66% agreement rate between humans and the algorithm (Table D.3). Inter-human agreement ranges from 50-97%; human-human and human-machine assessment are not statistically distinguishable. This suggests that our coding passes a version of the "Turing Test", where AI responses are not identifiable as coming from a machine, rather than from a human. It also implies that had we used human coders, our results would likely have been substantively similar.

5 Applications

In this section, we apply our new methods 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; they correlate to an important extent with patenting activities.

5.1 Patterns of Cultural Persistence 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 similar each student is to other individuals from his or her class ("individualism"). We also compare each student in year t with

[7] *"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."*

people graduating from the same high school 20 years earlier ("persistence"). Figure 7 plots the similarity scores for both types of analysis over time (Panel A), as well as our measure of style innovation (Panel B).

Women start out with high levels of individualism in the 1950s, while men register very low values. With each high school, the style choices of men were near-identical, while those of women varied a great deal. 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, individualism scores converge. Men start to show much more individualism 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 persistence trends. Males start out at very high levels – each generation looked 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 become the comparison group for contemporary high school students that we find evidence of renewed, rising persistence in style choices. In contrast, women showed relatively low persistence for much of the post-war era. It is only in the late 1980s that these values start to rise, converging with male levels by the mid-90s. By 2010, male persistence is actually lower than female persistence, 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. From 2000, however, these rates start to rise, and both men and women end the sample period at all-time-highs in terms of style innovation.

We can aggregate these trends to national averages, standardizing the scores for each gender. Figure A.4 shows the results. Aggregate persistence falls sharply from the late 1960s onwards, driven by big changes for men. It then recovers as high school students in the 1990s imitate the post-68 style of their parent's generation. Individualism rises at the same time, and reaches an all-time peak in the late 1970s, before settling at a high average level. Style novelty at the aggregate level shows sharp swings around a rising trend from the 1970s onwards, ending at the highest-ever values by 2010.

In Figure A.6, we examine persistence over time more closely. Instead of the (arbitrary) 20-year horizon used in Figure A.4, we explore the similarity of high school portraits at different horizons. Darker colors show greater persistence; light colors indicate great differences. On average, longer horizons are associated with lower similarity. Until the counterculture era, similarity scores for all years from the one immediately preceding to as far back as 20 years earlier showed high values, 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 that the drop occurs around 1970 and then persistence rebounds around 1980: high school cohorts around 1970 have styles that substantially from those that went to high school in 1960, but those that went to high school in 1980 increasingly started to mimic those that went to the same school in 1970. Put simply, this implies that the '1970-style' became influential long-term, remaining very popular around ten years later. More broadly, there are visible patterns of a form of auto-correlation in persistence as we move 'northeast' in the figure, tracking longer time horizons and later cohorts simultaneously. A concrete example of this is, again, the '1970-style' was fashionable for multiple years in the 1980s; the 12-year horizon persistence is high around 1982 and the 15-year horizon persistence is high around 1985. This style appears to have fashionable for at least 8-10 years. Thus, our measure of persistence 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. It is beyond the scope of our paper to study the determinants of such cycles, or test specific theories from the literature, but it is clear that this type of data can open the door for such inquiry.

What about variability over space? Figure 8 shows our three indicators by decade, plotted at the grid-cell level across the US. The 1950s and 1960s are characterized by great uniformity across space in terms of individualism, persistence, and novelty. From the early 1970s, there is a growing regional divide, with Northern states increasingly individualistic relative to a more conformist and less innovative South. For individualism, it takes until the 1980s for clear geographical patterns to emerge; for persistence, it takes until the 1990s. Novelty surges and variability across the US increases from the 1970s, but there is no clear geographical polarization.

In Figure 9, we examine distributions by year and commuting zone, exploiting the granularity of our data. The low individualism and high persistence of the 1950s and early 1960s is reflected in most commuting zones having very similar levels overall – there were little differences between the similarity scores across high school students, and compared with earlier decades. By the late 1960s, the sharp ridges of high conformity and persistence 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 individualism areas, mirrored by low/high persistence. 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.

Figure A.5 splits the data by gender. Again, we see substantial heterogeneity. The bimodality in individualism and persistence 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. The gender differences we detect here are substantial. Future work could explore their economic roots.

5.1.1 Time-varying Correlates of Style Change

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 examine the contribution of individual factors we use LASSO regressions. LASSO modifies standard OLS regressions by means of a regularization that shrinks the coefficients of less important variables to zero (Mullainathan and Spiess 2017). It is useful when dealing with high-dimensional datasets with many correlated predictors. Mathematically, it chooses β so as to minimize the loss function given by:

$$R(\beta) = \sum_{i=1}^{n}(y_i - \beta_0 - \sum_{j=1}^{p} \beta_j x_{ij})^2 + \lambda \sum_{j=1}^{p} |\beta_j| \tag{4}$$

where the first term penalizes the lack of accuracy and the second term penalizes the introduction of new covariates. Lambda (λ) is a hyperparameter and it is chosen by cross-validation, using the root mean squared error (RMSE) to assess prediction performance. It controls the strength of the penalty and determines the degree of sparsity in the model. As λ increases, the penalty on the coefficients also increases, causing more of the coefficients to shrink towards zero. This results in a model with fewer predictors, which can improve the interpretability and stability of the model. If $\lambda = 0$, the LASSO estimator is equivalent to OLS.

We regress the dependent variable (individualism/persistence) on the vector of style features using the LASSO technique for each decade, at the individual level. Note that if a coefficient is not shown for a variable, it was shrunk to zero. Given the number of different styles, it could be that only a few of them are driving persistence and individualism. After penalizing for adding more variables, most style features survive (i.e. the regularization parameter λ is greater than zero). This reflects that the style characteristics we chose are "important" to some extent, i.e. have predictive power, and are not "spanned" by other features.

Some variables are clearly more important than others. Moreover, differences between the coefficient for a variable across the decades reflect that a particular feature can be more important in some decades that in others. Figure 10 plots the coefficients of LASSO regressions explaining the contribution of individual style characteristics to the overall individualism score of an image, for men and women (Figure A.7 and A.8) shows the same for persistence and style novelty). Take the case of shirts without a collar, for example. In the 1950s and 1960s, wearing a shirt without a collar increased a young man's individualism score – it made them less similar to other, randomly chosen classmates. From the 1970s onwards, as more and more graduating seniors appear in yearbook pictures like this, this effect diminishes; by the 1990s, take-up rates are so high that the same style choice is associated with greater conformity. Bow ties are the strongest predictors of low individualism – and increasingly so in recent decades. In contrast, sweaters, long hair, glasses, and facial hair are always positively associated with greater individualism. For women, a bowlcut in female high-school senior portraits predicts style novelty, in all decades except the 1960s, when the style is particularly popular, and thus is negatively associated with novelty.

Similar changes in the predictive power (and sign) of attributes occur for other measures of cultural change in senior portraits. For example, crew cuts for men used to increase persistence in the 1950-70s; by the 1990s, they are reducing it. Conversely, curtained hair has gone from reducing persistence to increasing it. For women, low necklines used to reduce persistence; by the 1980s, it is adding to the persistence score of images. The changing coefficients derived from our LASSO exercise highlight one important advantage of our analytic approach. The effect of individual style choices on our measures of persistence and individualism is conditional on other people's choices; we can use them as time-varying, context-sensitive indicator of the social signal embedded in style choices.

5.1.2 Comparison of Style Individualism with Individualism from First Names

Important recent work has used unusual first names as an indicator of individualism (Bazzi et al. 2020). Is our indicator of individualism capturing the same variation? To examine this question, we collected the first names of 5,000 graduating seniors from 220 high schools during the 1960s in our sample. We investigated the share of rare names by using a list of the most common 200 names during the 1960s, coming from the most popular names by decade ranking available on the US Social Security Administration website. We then calculate name individualism as one minus the share of common names.

Figure A.9 shows a binscatter of high-school level individualism derived from images (on the y-axis) against name individualism (x-axis). The two are unrelated – the beta coefficient is -0.04 and the $R^2 = 0.01$. This suggests that we capture an orthogonal dimension of individualism.

5.2 Inventions and Style Novelty

"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 as such? We use the data on commuting-zone level innovation from Bell et al. (2019). They document both the number of inventors per commuting zone and birth cohort in the US, as well as the importance of these patents (proxied by patent applications). Remarkably, their data contain patenting rates for each year of age. To compare patenting rates across high- vs low-style innovation areas, we first aggregate our data to the commuting zone level. Next, we create a dummy variable that takes the value of 1 if there is substantial style innovation in a commuting zone. As explained in detail above, we define top innovation as a new style that has so far been adopted by less than 1% of high school graduates in history (including the current year). If the share of such styles in a commuting zone is non-zero, we consider it innovative. We then examine whether the cohorts graduating from high schools in such a commuting zone are more likely to apply for patents, or have a patent granted to them.

To fix ideas, consider Figure 11. 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 top 1% of style novelty, with 0.075% 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 if the pattern visible in the case of Jobs can be generalized in our data. Figure 12 shows the patent application rates by age for graduates of with and without style innovation. As is readily apparent, the differences are substantial. Both men and women who graduate from areas with high levels of fashion innovation apply for more patents than those from areas without style innovators. They also receive many more of them (Figure A.11). The difference is considerable from age 20 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, by age 30, before declining gradually. Next, we abstract from the age dimension and collapse style innovation and patenting rates at the commuting zone/year level. We do this for both genders jointly, and plot the pattern by decade in Figure 13. High style innovation areas consistently show higher patenting frequencies, outperforming by 15-20 pp.

In Table 5, we examine the determinants of patenting more broadly, adding style indicators to a rich set of covariates. When we add individualism, persistence, and style novelty, we can explain 1.5-2% of the variation in patenting rates (col. 2). Style novelty is positively correlated; the partial correlations with persistence and individualism are negative (but our measures of style innovation are correlated among each other). We would expect the patenting variables to scale with population mechanically and due to agglomeration forces. As we add controls for population in a flexible manner, the coefficient on individualism becomes insignificant (col. 3+4).[8] Only style novelty is consistently positive across outcomes and specifications in panels A and B, although the magnitude and statistical significance of the coefficient fluctuate. In our most restrictive linear specification with a large set of fixed effects (col. 7), the coefficient is small and insignificant. We examine potential interactions with population size in the specification of column 8. Interestingly, style novelty predicts patenting later in life only when the population size in the commuting zone is relatively small and becomes weaker as population size increases. While speculative, this interaction effect suggests that the dimension of culture that we capture, a form of highly individualistic expression through novel style choices, may matter the most in relatively small communities.

The relationships in Table 5 are not "causal" – putting an earring on a young man in 1956 will not necessarily induce patenting in later life. It is, however, compatible with an

[8] We present simple variable definitions and regression specifications for ease of interpretation noting that the true functional form could be different.

interpretation that emphasizes the importance of broad cultural outlook, with tolerance towards unusual behaviors and choices in one domain encouraging an inclusive and open-minded approach in many activities. School policies that encourage innovation and a liberal outlook on life may well causally increase style innovation and technological innovation later in life. In this sense, the style choices documented in high school yearbook images are suggestive of the role of broader social and cultural forces in encouraging innovation.

6 Conclusion

Imagine walking through the halls of a typical American high school in the 1950s. You would be faced with a sea of uniformity – clean-shaven faces and short, neatly trimmed hair for the boys. Senior portraits would almost all show the men in suit and tie, looking earnestly at the camera. At the same time, while many girls are wearing modest and conservative clothing, others experiment with new hairstyles, jewellery, or low necklines. Fast forward to the late 1960s, and you would witness a dramatic shift. Suddenly, long hair and rebellious fashion choices for men dominate the corridors and yearbook pages. Some of the men sport long hair, sideburns, and a moustache, others dress conservatively, in the style of their parents; women on the 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.

Undoubtedly, the counterculture movement that swept the US had a profound effect on the relationship between the sexes, between citizens and the state, the notion of work, and politics in general – to name but a few dimensions. But how can we quantify such cultural change? Visual culture and self-representation through style and fashion choices are arguably important elements, but there has so far been little empirical analysis of this dimension of culture.

We analyze millions of senior portraits spanning decades, extracting key style attributes. This allows us to create a data-driven narrative of cultural evolution. 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 introduce a set of methods and tools that allow rigorous and systematic analysis of images as a source for cultural change. To do so, we use sparse feature vectors capturing key attributes of style. We train algorithms to identify style features in images using convolutional neural networks. The vector representation in turn facilitates the use of standard measures of similarity. These can be used to map three key dimensions of culture: a lack of conformity ("individualism"), persistence, and style novelty. With this, we contribute to a growing trend in economics to use image data systematically (Adukia et al. 2023; Ludwig and Mullainathan 2024).

Secondly, we apply our new methods to a large dataset of US high school senior images. We show marked convergence in style choice between men and women. Where in the 1950s, women had expressed themselves through a wide range of style choices, and men had mostly adopted a uniform look, by the 1990s both the levels of individualism and of persistence had converged completely. For men, we trace the decline and fall of conformity in senior portraits, showing that the cultural revolution of the sixties and early seventies not only led to a sharp decline of conformity within each local high school. It also destroyed the – previously high – level of persistence, when portraits of the parent's generation were broadly similar to those of their children. We find that the late sixties and early seventies emerge as a dramatic discontinuity in the way in which young Americans on the cusp of adulthood presented themselves to the world, and to posterity. As the old high-conformity, high-persistence steady state collapsed, it gave way to a unique upsurge of non-conformity within the vast majority of high schools and growing generational cleavages in style choice. The subsequent decline in innovation and individualism was temporary; individualism stabilized at high levels from the 1990s, and style novelty has surged to a new high by 2010.

For women, the trends are an almost precise mirror-image. High levels of within-school variation decline from the late 1960s. While they bounce back in the late 1970s, the trendline indicates long-term decline. At that moment in time, persistence remains the same, but rises sharply by the late 1980s. By the late 2000s, female persistence levels are even slightly higher than for men. Just as in the case of men, style innovation is rising gradually over time, ending the sample period at all-time-highs.

Importantly, we also document the emergence of sharp polarization across commuting zones. Some high schools maintained a low individualism, high persistence pattern of style choices, while others scored high on individualism and low on persistence. This divergence has grown over time, and has a strong geographical component, with low individualism prominent in the American South. In contrast, Bertrand and Kamenica (2023) used survey responses and consumption pattern to argue that polarization has not increased in recent decades in the US. Our results suggest that a new source, combined with greater geographical and temporal resolution, can complement this perspective.

We also show that style choices of high school seniors are correlated with an economically important activity – patenting. In commuting zones where more high school seniors adopted novel, unusual styles, patenting rates are higher. We show that this finding holds at the cohort-commuting zone level. While not causally identified, this points to important synergies 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 such imitation can simultaneously lead to rapid change and long periods of uniformity has so far been challenging to document empirically. Our methods allow us to trace, using uniquely granular and high-frequency data, how important style changes emerge, and the extent to which they break with pre-existing patterns. We also demonstrate that daring to “dress different” is associated – both geographically and at the cohort level – with “think different”, as it correlates with the frequency of important technological innovations.

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.

Akerlof, George A, and Rachel E Kranton. 2000. "Economics and Identity." *The Quarterly Journal of Economics* 115 (3): 715–753.

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.

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.

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.

Fernández, Raquel, and Alessandra Fogli. 2009. “Culture: An Empirical Investigation of Beliefs, Work, and Fertility.” *American Economic Journal: Macroeconomics* 1 (1): 146–177.

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

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. 2018. "Ancestral Characteristics of Modern Populations." *Economic History of Developing Regions* 33 (1): 1–17.

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.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2016. "Long-Term Persistence." *Journal of the European Economic Association* 14 (6): 1401–1436.

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

Haxby, James V, Elizabeth A Hoffman, and M Ida Gobbini. 2000. "The Distributed Human Neural System for Face Perception." *Trends in Cognitive Sciences* 4 (6): 223–233.

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.

Kanwisher, Nancy, Josh McDermott, and Marvin M Chun. 2002. "The Fusiform Face Area: A Module in Human Extrastriate Cortex Specialized for Face Perception." *The Journal of Neuroscience* 17 (11): 4302–4311.

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.

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.

Mullainathan, Sendhil, and Jann Spiess. 2017. “Machine Learning: An Applied Econometric Approach.” *Journal of Economic Perspectives* 31 (2): 87–106.

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.

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.

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

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

Todorov, Alexander. 2017. *Face Value: The Irresistible Influence of First Impressions*. Princeton University Press.

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 Yearbook

Style Feature	Model Dummies	
	Image A	Image B
long_hair	1	1
curtained_hair	1	0
curly_hair	0	0
spiky_hair	0	0
bowlcut_hair	0	0
crew-cut_hair	0	0
afro	0	0
beard	0	0
moustache	0	0
earrings	0	0
necklace	0	0
glasses	1	0
low_neckline	0	1
sweater	0	0
suit	0	0
shirt_collar	1	0
shirt_nocollar	0	0
bowtie	0	0
tie	0	0


Image A


Image B

Notes: We lsit the style features of images A and B in tabular form. The presence of a style is indicated by 1, absence by 0. Values in the table derived from a Google Vertex AI multilabel model prediction, after desaturation, with a threshold of 0.3. Cf. Appendix C for details.

Figure 2: Style Attributes in High School Senior Yearbook Pictures

(a) Males

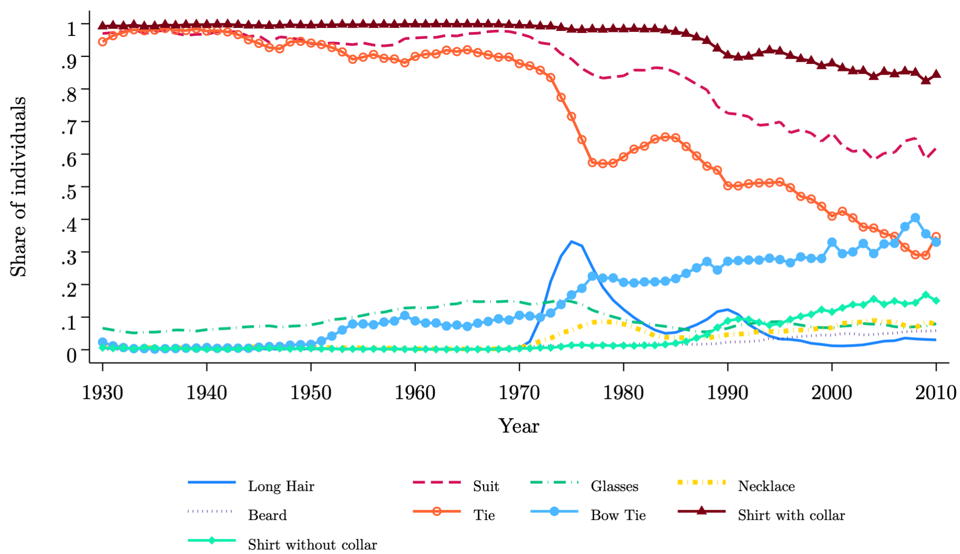

(b) Females

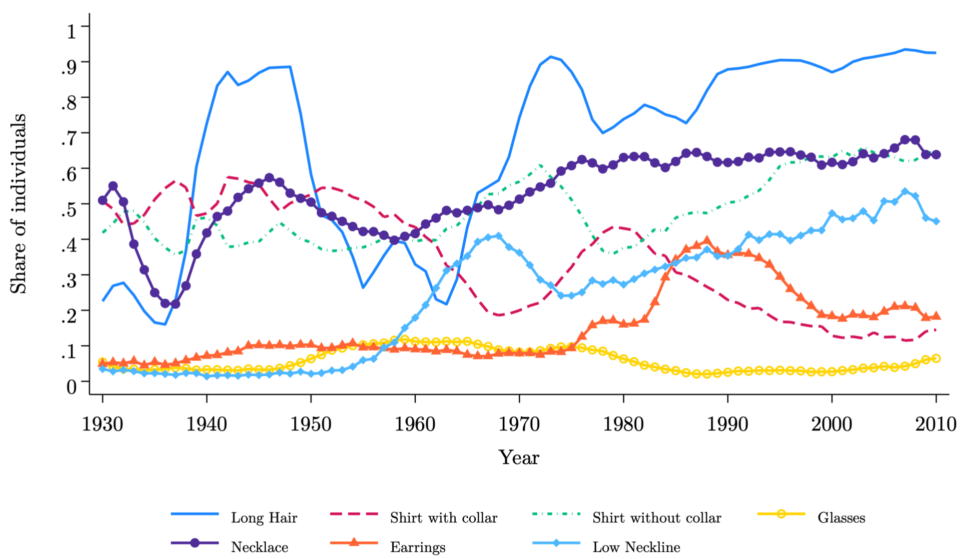

Notes: This figure breaks down the evolution of style features over time by gender. Every style feature prediction at the individual (image) level is transformed into a 0/1 dummy using a 0.3 threshold. Predictions are derived from a Google Vertex AI multi-label model and hair classifier. Values are averaged by year. Cf. Appendix C for details.

Figure 3: Styles in High School Senior Yearbook Pictures

(a) Males

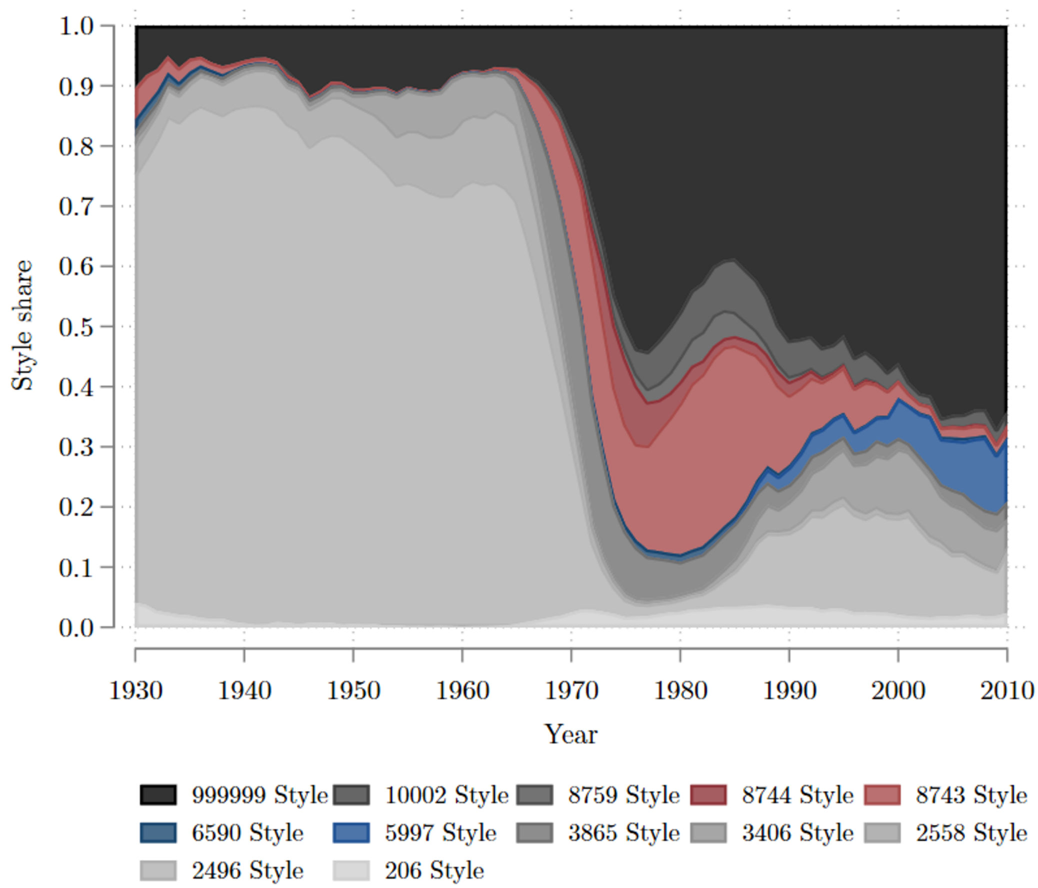

Feature	10002	8759	8744	8743	6590	5997	3865	3406	2558	2496	206
Curtained Hair	■	■	■	■	□	□	□	□	□	□	□
Curly Hair	□	□	□	□	■	□	□	□	□	□	□
Spiky Hair	□	□	□	□	□	■	□	□	□	□	□
Bowl Cut	□	□	□	□	□	□	■	□	□	□	□
Crew Cut	□	□	□	□	□	□	□	■	■	■	□
Bowtie	■	□	□	□	□	■	□	■	□	□	□
Beard	□	□	□	□	□	□	□	□	□	□	□
Afro	□	□	□	□	□	□	□	□	□	□	□
Tie	□	■	■	■	■	□	■	□	■	■	■
Glasses	□	□	□	□	□	□	□	□	■	□	□
Shirt without Collar	□	□	□	□	□	□	□	□	□	□	□
Shirt with Collar	■	■	■	■	■	■	■	■	■	■	■
Moustache Only	□	■	□	□	□	□	□	□	□	□	□
Suit	■	■	■	■	■	■	■	■	■	■	□
Sweater	□	□	□	□	□	□	□	□	□	□	□
Long Hair	□	□	■	□	□	□	□	□	□	□	□

(b) Females

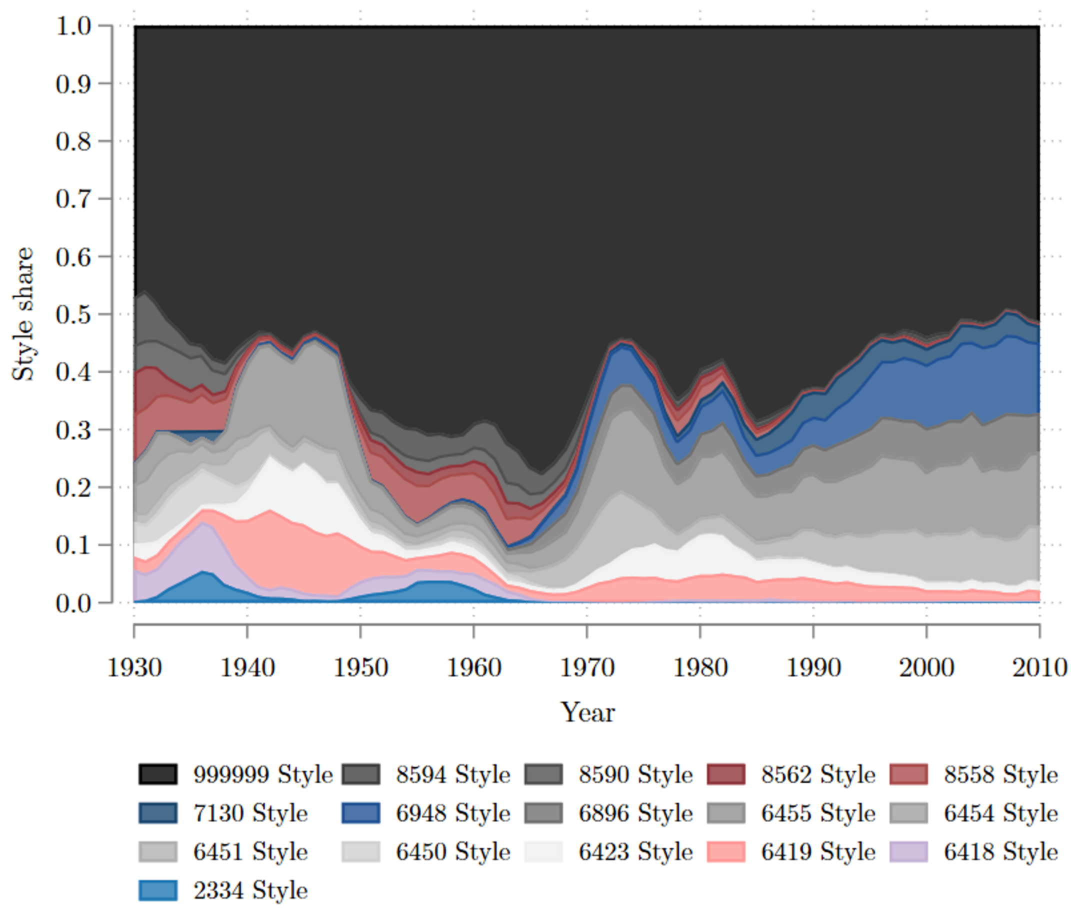

Feature	8594	8590	8562	8558	7130	6948	6896	6455	6454	6451	6450	6423	6419	6418	2334
Curtained Hair	■	■	■	■	□	□	□	□	□	□	□	□	□	□	□
Curly Hair	□	□	□	□	■	■	■	■	■	■	■	■	■	■	□
Bowl Cut	□	□	□	□	□	□	□	□	□	□	□	□	□	□	■
Bowtie	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Afro	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Low Neckline	□	□	□	□	■	■	■	□	□	□	□	□	□	□	□
Earrings	□	□	□	□	■	□	□	□	□	□	□	□	□	□	□
Tie	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Glasses	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Shirt without Collar	■	■	□	□	■	■	□	■	■	■	■	□	□	□	□
Shirt with Collar	□	□	■	■	□	□	□	□	□	□	□	■	■	■	■
Suit	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Necklace	■	□	■	□	■	■	■	■	■	□	□	■	□	□	□
Sweater	□	□	□	□	□	□	□	□	□	□	□	□	□	□	□
Long Hair	□	□	□	□	■	■	■	■	□	■	□	■	■	□	□

Notes: This graph depicts the share over time of discrete combination 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 style 999999 to all unique styles that comprise less than 5% of the sample for all years. (In our analysis of style novelty, we use all unique styles as they are defined, and not 999999.)

Figure 4: Sample Images and Individualism Comparisons

1966

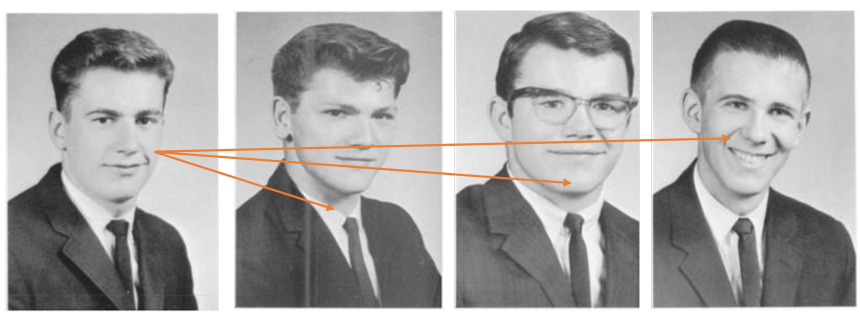

Notes: This figure illustrates the calculation of the individualism score. For individual i, it is calculated as (1-mean cosine similarity) when compared with all the images of other seniors with the same gender in the same high school cohort. The score for the individual on the left is 0.098, indicating a low level of individualism, since most style choices (suit, hair, tie) are similar.

Figure 5: Sample Images and Persistence Calculations

(a) Low persistence example

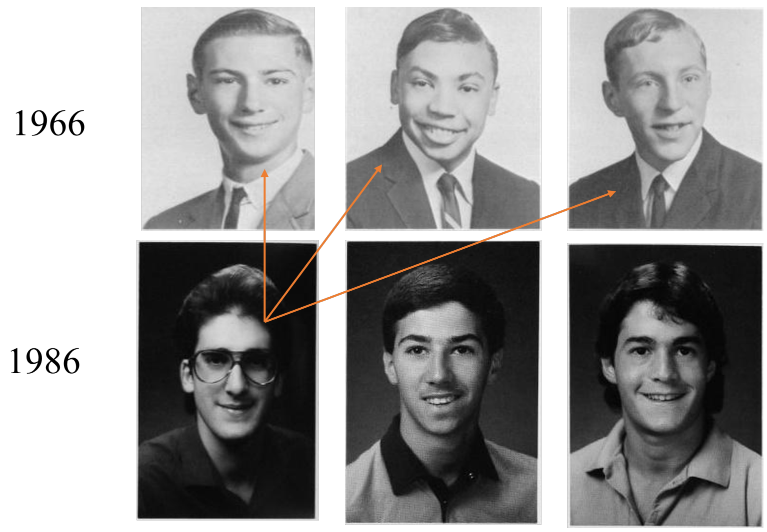

(b) High persistence example

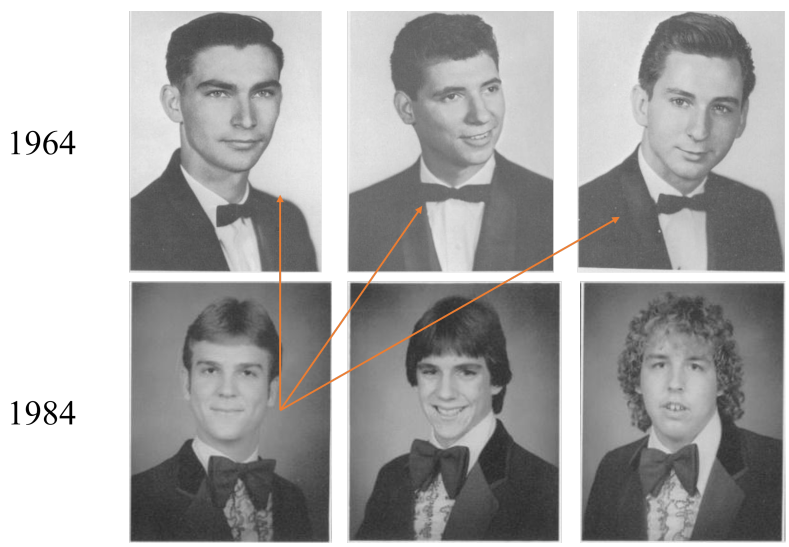

Notes: The figure illustrates the calculation of persistence scores. We calculate the similarity of everyone in a cohort, comparing each of them with the style of graduates 20 years previous (of the same gender). We then average this score for the cohort. Panel (a) is an example of low persistence (0.056). Panel (b) is an example of high persistence (0.83).

Figure 6: Sample Images and Style Novelty Example

(a) Males

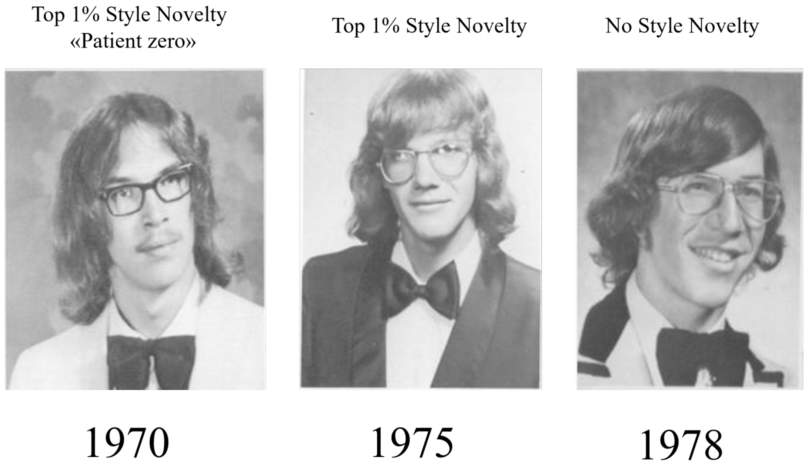

(b) Females

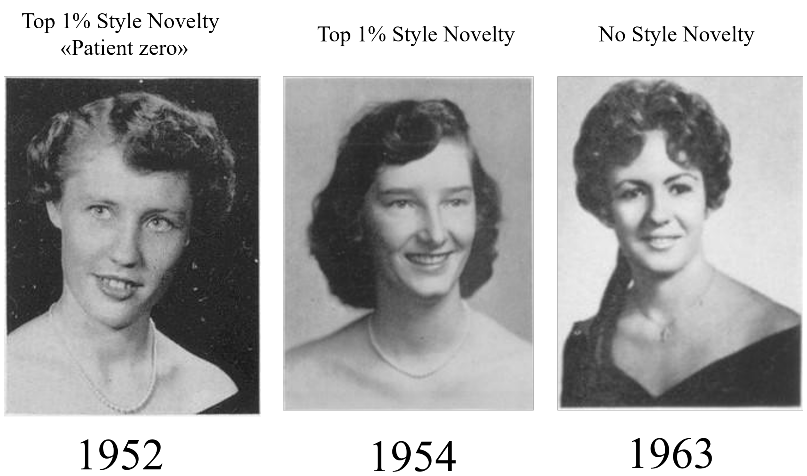

Notes: The figure illustrates our style novelty measure. The male style combination consists of long and wavy hair, shirt with collar, suit, bowtie, and glasses. It first appeared in Burlington, North Carolina in 1970. The left-most picture in Panel (a) shows the earliest example of this style in our data. The same style appears in subsequent years in other places. Take the middle images in panel (a), for example, from Georgia in 1975. Because take-up has been limited, it is still in the top 1% of style novelty. By 1978, after it has been copied many times, it is no longer a novel style in our data. The female style combination shown in Panel (b) consists of curtained and curly hair, low neckline and necklace first appears in Tennessee in 1952. When we observe the same style in Georgia in 1954, it is still in top 1% of style novelty. However, by 1963, it has spread sufficiently so that the image from Indiana on the right no longer qualifies as novel.

Figure 7: Individualism and Persistence Trends over Time, by Gender

Panel A: Individualism

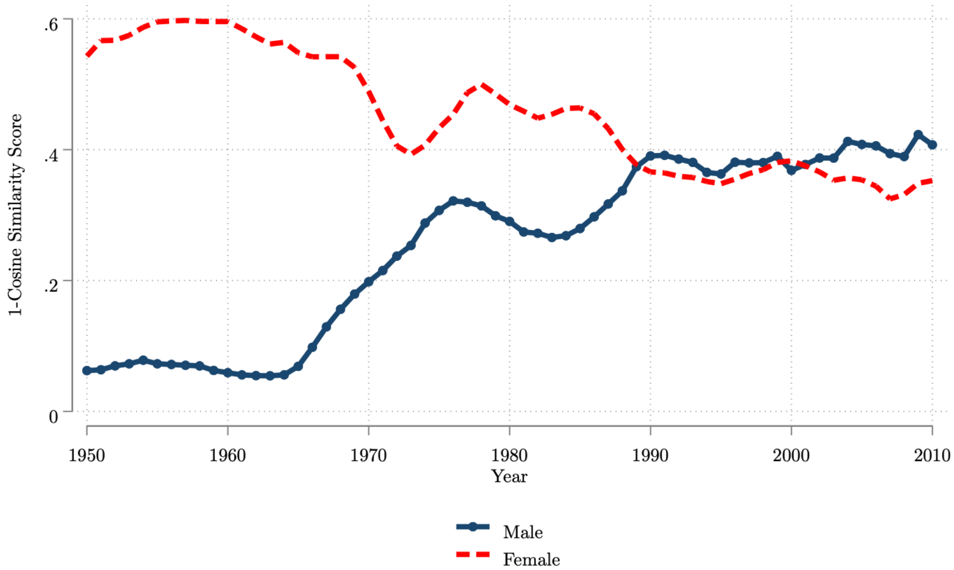

Panel B: Persistence

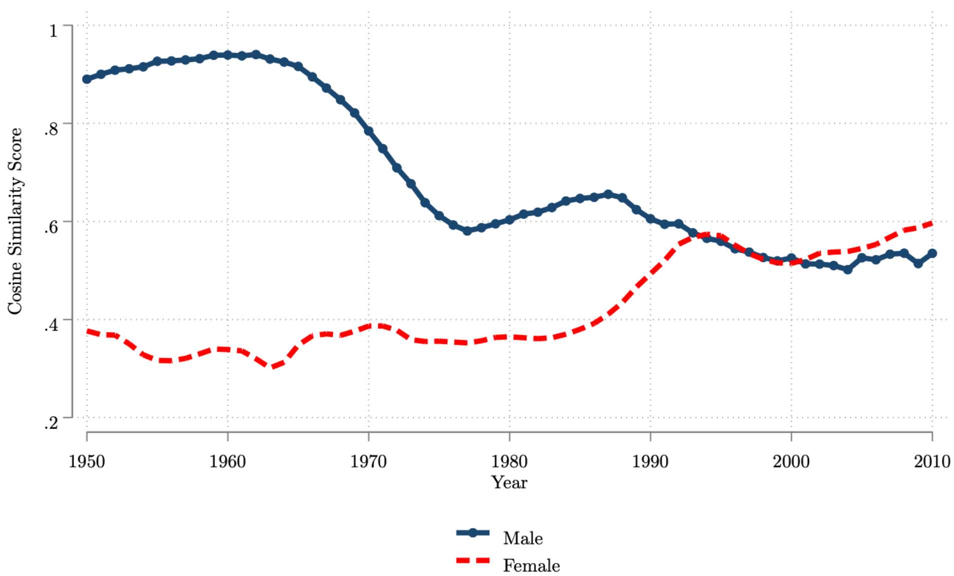

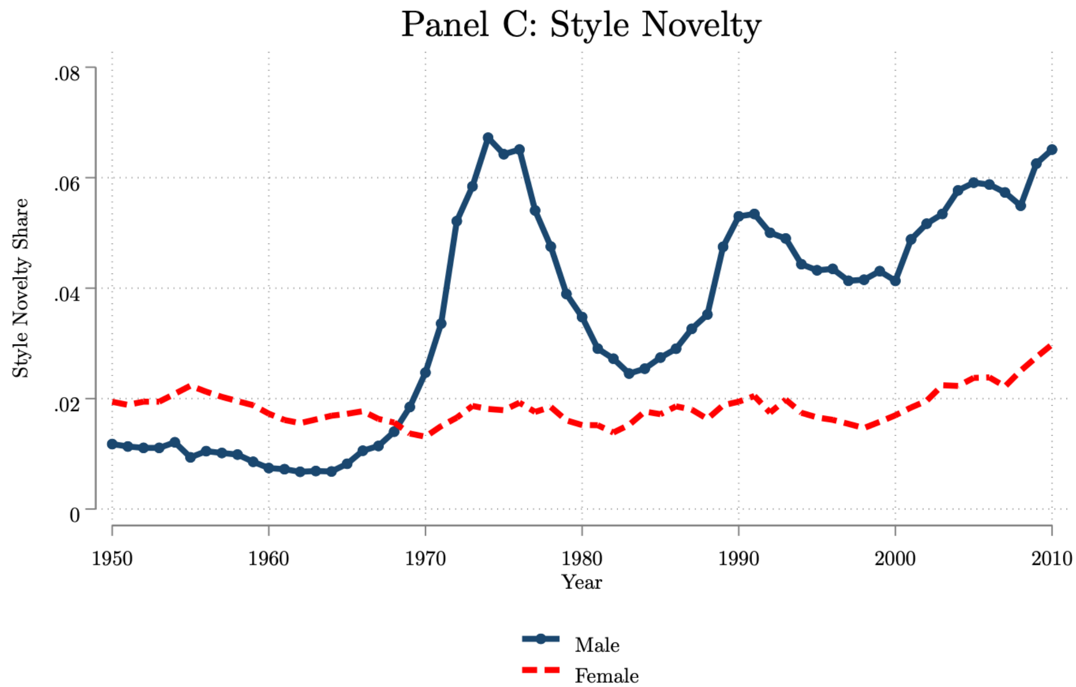

Notes: This figure plots yearly average cosine (for persistence) and 1-cosine (for individualism) scores for individualism (Panel A) and persistence (Panel B); Panel C plots the share of style innovators. Averages come from our image level dataset (14.5 million observations), split by gender.

Figure 8: Geographical Distribution over Time

Panel A: Individualism

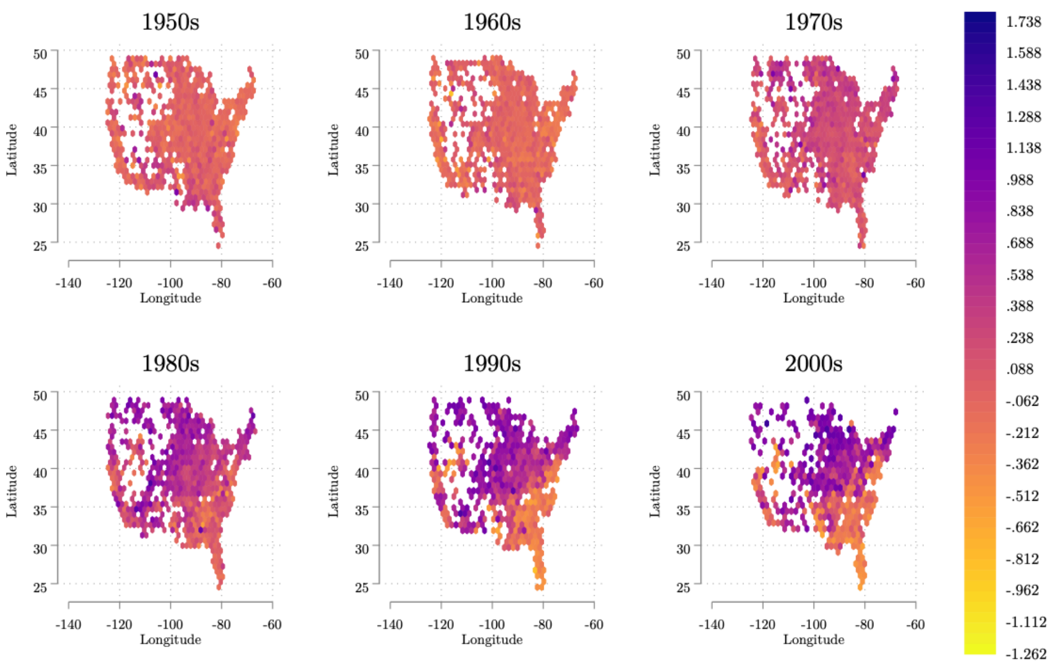

Panel B: Persistence

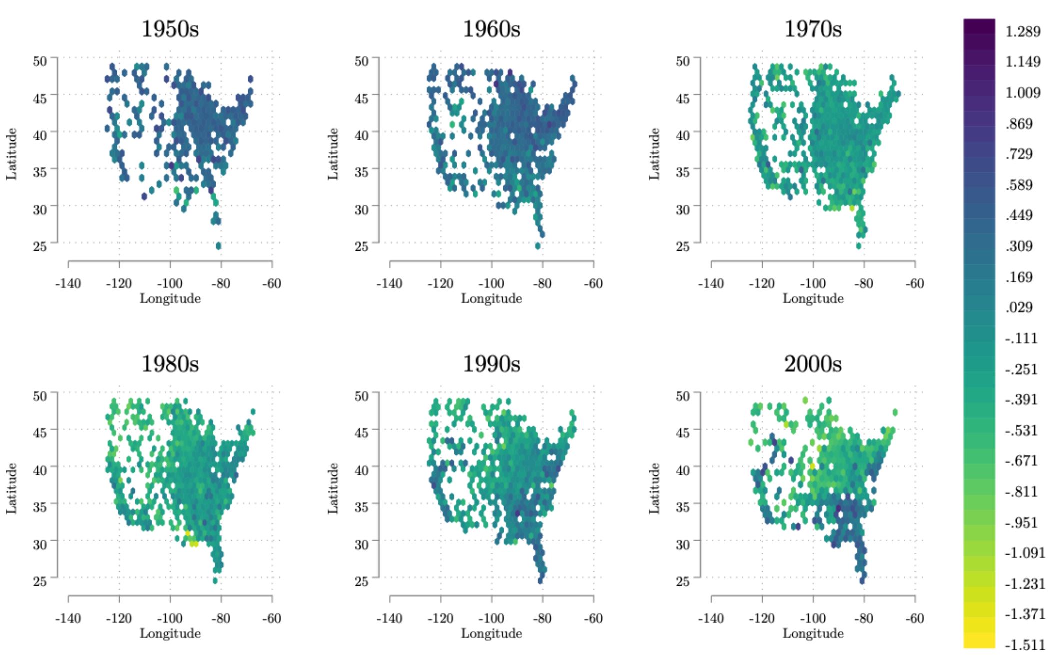

Panel C: Style Novelty

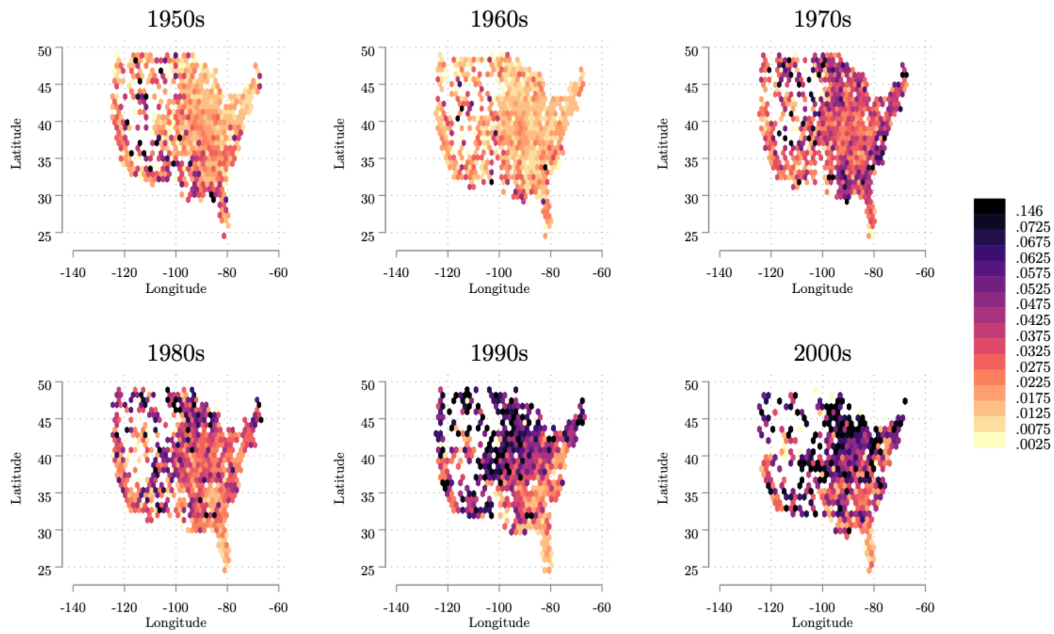

Notes: These plots display the evolution of the three style measures across decades and space. Each observation is a decade/HS geolocation average across females and males for individualism, persistence and style novelty, averaged at the high school level. High schools are geolocated by latitude and longitude. Values for Panel A and B represent z-scores; Panel C plots the share of style innovators per high school geolocation. Individualism and persistence plots range from global minimum to global maximum across all decades; style novelty ranges from 5th to 95th global percentiles across decades, to avoid noise.

Figure 9: Variation across Commuting Zones

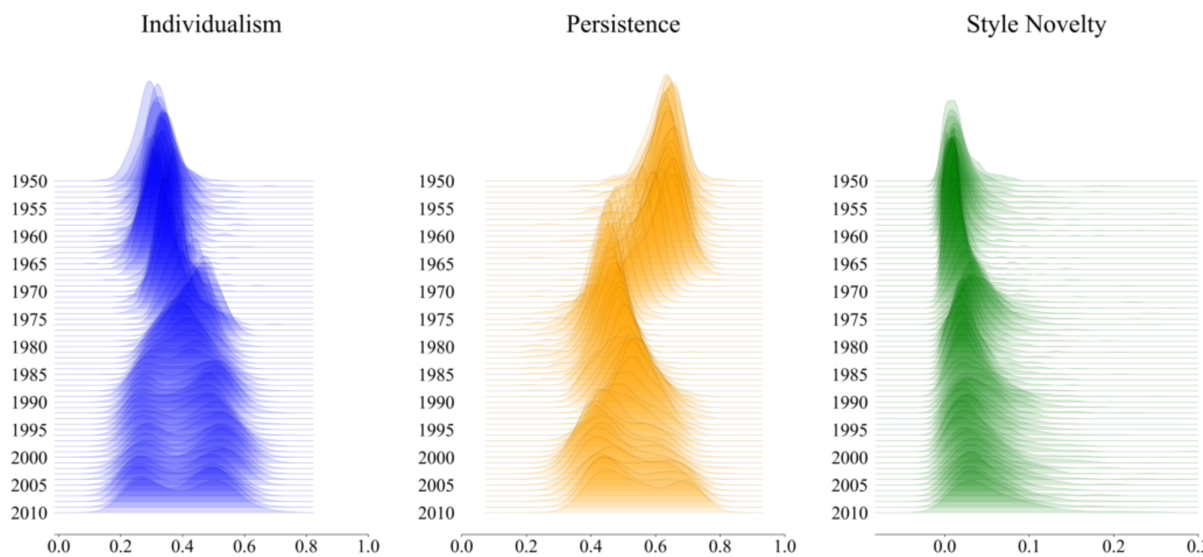

Notes: This figure shows ridge plots of the distributions of persistence, individualism, and style novelty across commuting zones, pooling both genders, for the 1950-2010 senior high school year cohorts. For each commuting zone and cohort, the average value across the images is first calculated. We then display the range of commuting-zone average. For individualism and persistence, the values on the x-axis are cosine similarity averages within commuting zone-year couplings; for style novelty, x-axis values are the share of style innovators across commuting zone-year couplings.

Figure 10: Drivers of Individualism

(a) Males

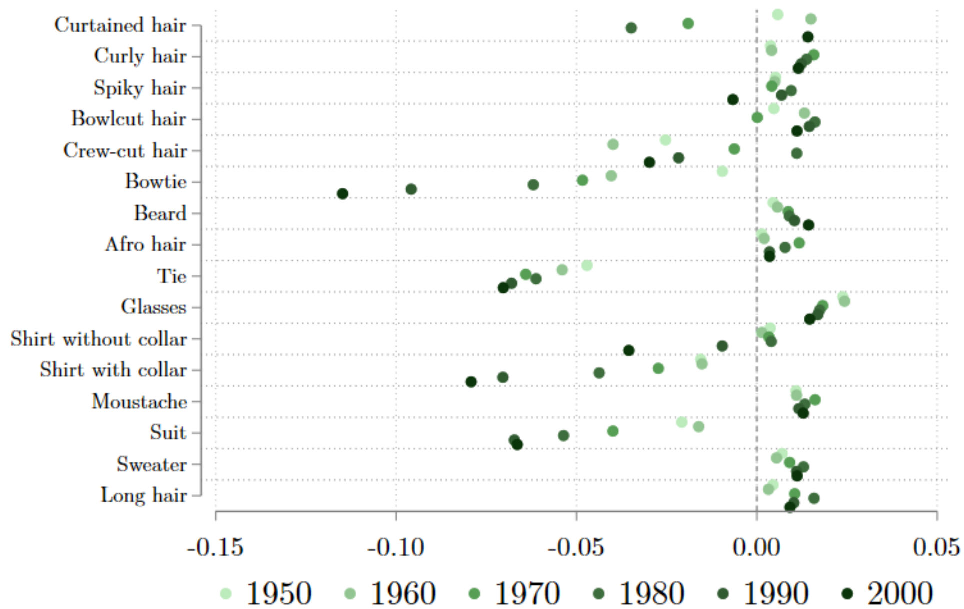

(b) Females

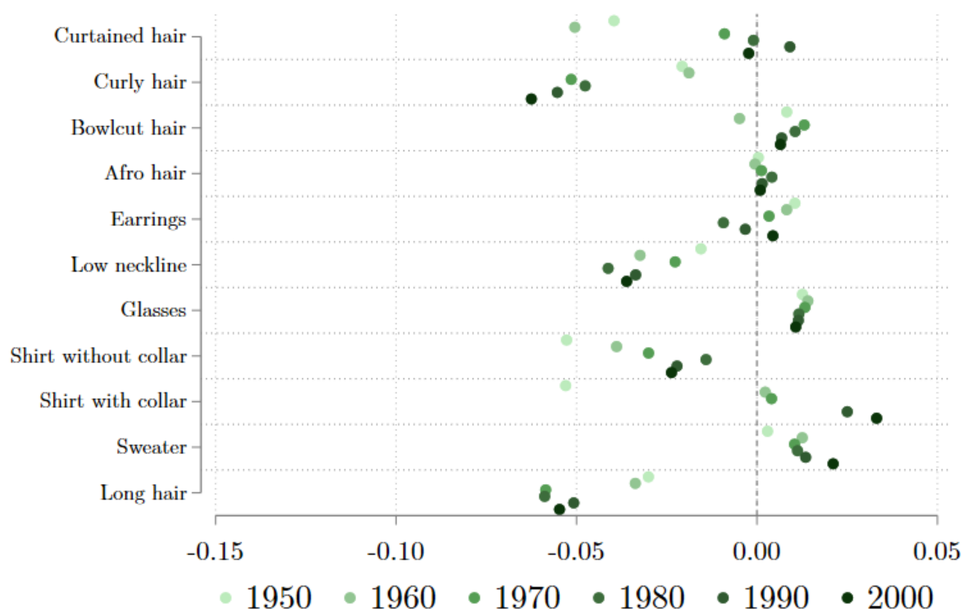

Notes: Dots represent decade-level coefficient estimates from an OLS Lasso regression of individualism scores on gender-specific style attributes. Each dot represents a single coefficient for a style feature, derived from bi-variate regressions. Values greater zero indicate a positive contribution to students' individualism scores. Panel (a) is for male images only, Panel (b) for females.

Figure 11: Steve Jobs Example

(a) Senior Portrait (1972)

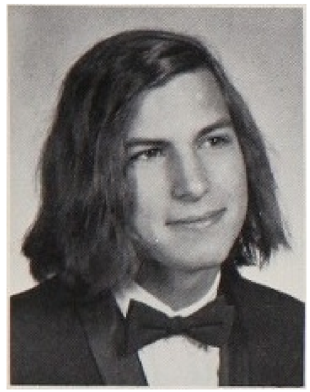

(b) Novelty of Steve Jobs' Senior Portrait over Time

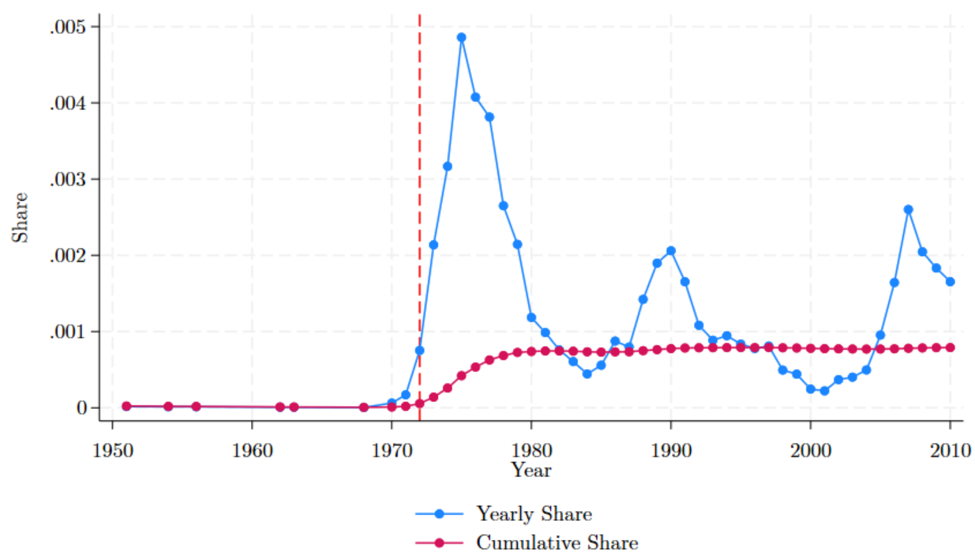

Notes: In 1972, Steve Jobs was a senior in Cupertino's Homestead High School, California. Panel (a) shows his senior portrait. The associated scores for Job's style are Individualism = 0.56; Yearly style share = 0.075%; Cumulative style share = 0.005%; Style Innovator Dummy = 1. Panel (b) shows the fraction of High School portraits showing the same style as Steve Jobs in each year, across all high schools in our sample. The red dashed line indicates Steve Jobs' senior year. At this point in time, this style had only ever been used by 0.005% of all graduating seniors ever in our data (and 0.075% of 1972 seniors)

Figure 12: Patent Applications, by Year since Graduation

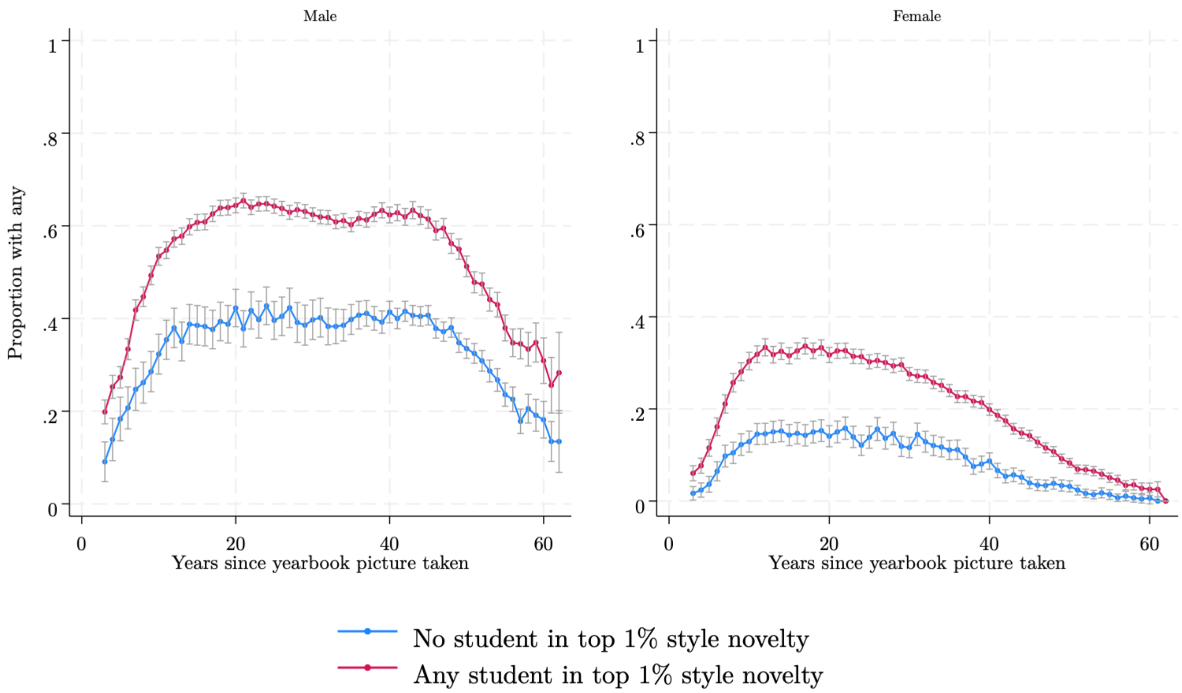

Notes: This figure shows patent applications over by year since graduation, for two groups of high school graduates – style innovators and non-style innovators. Panel A shows rates of patenting for men, Panel B, for females. Year of birth and CZ of an innovator was merged with high school and an imputed year of birth of students, calculated as year of the yearbook – 17. Applicant behavior is measured by a dummy equaling one if a CZ-year coupling exhibits at least one applicant. Red lines indicate the presence of at least one top 1% style innovator in a commuting zone-cohort. The time period ranges from 1950 to 2000.

Figure 13: Patent Applications by Decade and Style Innovation

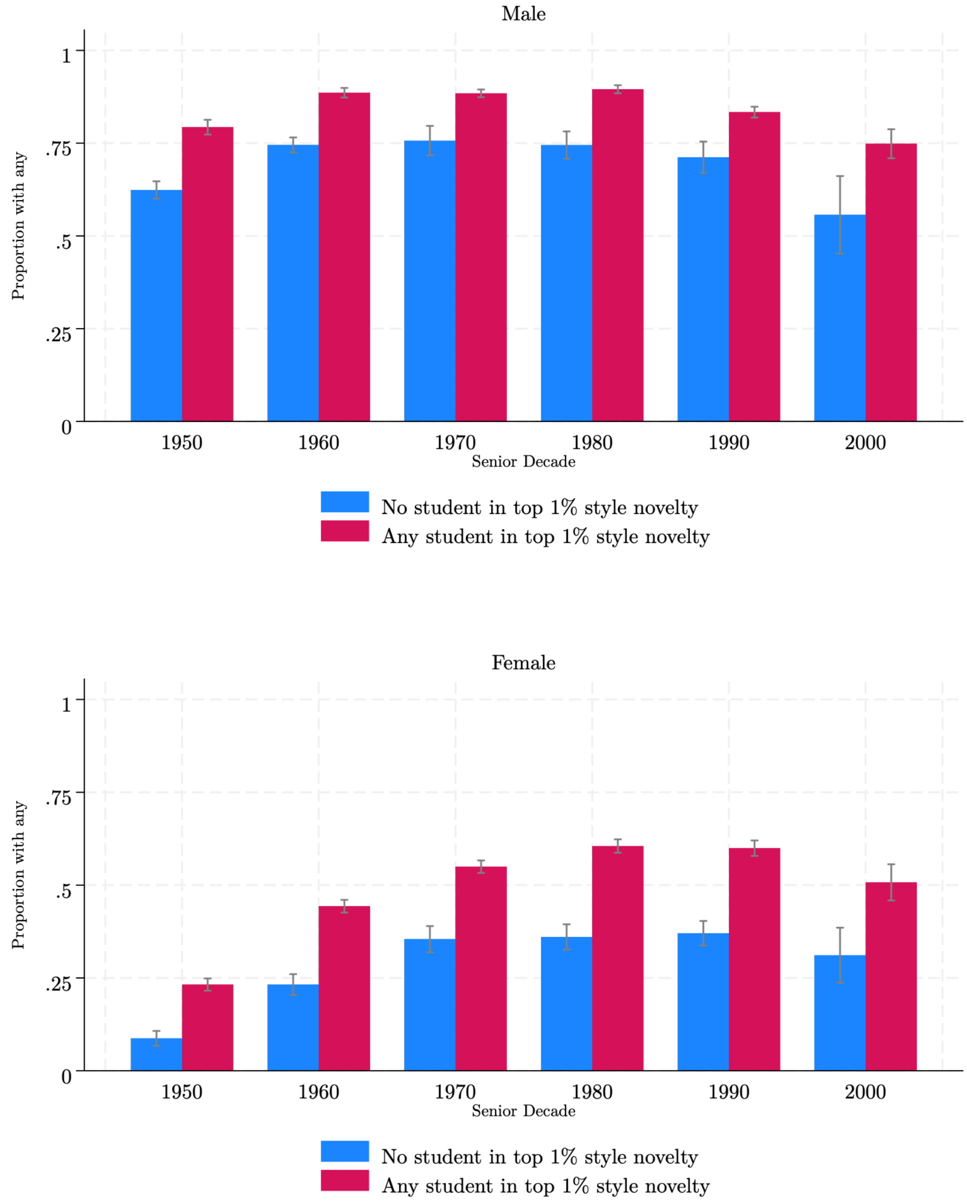

Notes: This figure displays the relationship between style novelty and patent applications (any vs none) at the CZ-decade level. Year of birth and CZ of an innovator was merged with high school data, using an imputed year of birth of students, calculated as year of the yearbook – 17. Red bars show rates of patenting applications in style innovative CZs. Applicant behavior is measured by a dummy equaling one if a CZ-year coupling contains at least one applicant. Red (blue) bars indicate the presence a least one style innovator (in the top 1% of new styles). The time period ranges from 1950 to 2000.

Tables

Table 1: Summary Statistics, Image Data

	Image Level Panel			CZ Level Panel		
	N	Mean	St. Dev.	N	Mean	St. Dev.
Style Attributes						
Male	14,403,792	0.472	0.499	23,086	0.479	0.066
Curtained hair	14,403,792	0.281	0.450	23,086	0.250	0.132
Curly hair	14,403,792	0.296	0.457	23,086	0.280	0.137
Spiky hair	14,403,792	0.017	0.130	23,086	0.017	0.035
Bowlcut hair	14,403,792	0.086	0.281	23,086	0.090	0.084
Crew cut hair	14,403,792	0.260	0.439	23,086	0.300	0.228
Bowtie	14,403,792	0.072	0.258	23,086	0.070	0.132
Facial hair	14,403,792	0.006	0.078	23,086	0.006	0.013
Afro hair	14,403,792	0.007	0.084	23,086	0.006	0.022
Low neckline	14,403,792	0.151	0.358	23,086	0.134	0.145
Earrings	14,403,792	0.084	0.278	23,086	0.080	0.073
Tie	14,403,792	0.368	0.482	23,086	0.363	0.164
Glasses	14,403,792	0.089	0.285	23,086	0.107	0.072
Shirt without collar	14,403,792	0.264	0.441	23,086	0.242	0.104
Shirt with collar	14,403,792	0.629	0.483	23,086	0.652	0.126
Moustache	14,403,792	0.036	0.185	23,086	0.033	0.040
Suit	14,403,792	0.435	0.496	23,086	0.430	0.125
Necklace	14,403,792	0.305	0.460	23,086	0.278	0.111
Sweater	14,403,792	0.086	0.281	23,086	0.095	0.066
Long hair	14,403,792	0.388	0.487	23,086	0.355	0.167
Style Measures						
Individualism - Pooled	14,403,211	0.349	0.216	23,086	0.361	0.086
Individualism - M2M	6,797,984	0.208	0.179	23,086	0.219	0.173
Individualism - F2F	7,605,227	0.475	0.160	23,086	0.498	0.105
Persistence - Pooled	6,861,328	0.543	0.237	15,105	0.533	0.092
Persistence - M2M	3,208,013	0.700	0.221	15,105	0.687	0.177
Persistence - F2F	3,653,315	0.405	0.149	15,105	0.387	0.093
Style Novelty - Pooled	14,358,447	0.023	0.149	23,086	0.026	0.025
Style Novelty - M2M	6,779,877	0.028	0.166	23,086	0.032	0.045
Style Novelty - F2F	7,578,570	0.018	0.131	23,086	0.021	0.023

Notes: This table reports summary statistics for style features, individualism, persistence, and style innovation at the individual level (image level) and at the CZ level for the sample period 1930 to 2010. Style attributes are discretized to be dummies 0-1. Similarity scores are calculated as cosine similarity relative to all other individuals in the same year/gender/high school. Persistence scores are cosine similarity scores compared with individuals in the same high school 20 years earlier. Style novelty is a top 1% dummy variable based on cumulative occurrences of style combinations.

Table 2: Summary Statistics, Patent Data

	N	Mean	St. Dev.	Min	Med
Average N of patents granted	23,086	0.527	1.008	0.000	20.005
Applicants share	22,337	0.440	0.716	0.000	11.917
Grantee share	23,086	0.348	0.546	0.000	9.985
Average N of patents granted - Males	21,160	1.051	1.924	0.000	35.111
Applicants share - Males	20,308	0.865	1.328	0.000	20.766
Grantee share - Males	21,160	0.687	1.031	0.000	17.368
Average N of patents granted - Females	21,228	0.078	0.256	0.000	21.149
Applicants share - Females	20,490	0.084	0.198	0.000	3.846
Grantee share - Females	21,228	0.056	0.132	0.000	3.846

Notes: This table shows indicators of patenting activity. Values are per 1,000 individuals, averaged across commuting zones. We match patent data by commuting zone and year to high school images through the year of birth of patentees/by taking the year of graduation-17.

Table 3: Correlations between Demographics & Style Measures

	Individualism	Persistence	Style Novelty
Share of Non-White	-.079	-.109	.055
Employment Rate	-.132	.226	-.098
Manufactury Employment Rate	-.045	.118	-.036
Share of High School completion	.248	-.09	.19
Total Population (county)	-.051	.069	-.018
Population Density (per sq. mile)	.057	-.044	.05
Median Family Income Adj.	.202	-.205	.147

Notes: This table shows the correlation between the three style dimensions (individualism, persistence, style novelty) and a set of socio-demographic controls. The data are a FIPS/Year panel ranging from 1950 to 2000 where the style indicators are merged with a FIPS/Year demographic dataset from ICSPR Study 2896. Data is available at the decade/year level and is then interpolated linearly. We merge approximately 90% of county-year images data with the ICSPR data. All variables in these regressions are standardized z-scores; FEs are at the FIPS and year level, std. errors are clustered at the FIPS level.

Table 4: Style Choice and Covariates, Panel Estimates

	Individualism			Persistence			Style Novelty		
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)
Median Familty Income Adj.	-0.318***	-0.293***		0.321***	0.317***		-0.195***	-0.183***	
	(0.041)	(0.043)		(0.042)	(0.042)		(0.030)	(0.032)	
Share of Non-White		0.0928			0.220**			0.00422	
		(0.060)			(0.069)			(0.041)	
Employment Rate		0.0981***			-0.0449*			0.0655***	
		(0.018)			(0.019)			(0.014)	
Manufactury Employment Rate		0.0612*			-0.0316			0.0461*	
		(0.024)			(0.026)			(0.019)	
Share of High School completion		-0.0443			-0.0352			-0.0195	
		(0.063)			(0.078)			(0.048)	
Total Population (county)		-0.323***			0.126*			-0.182***	
		(0.077)			(0.054)			(0.045)	
Population Density (per sq. mile)		0.0513**			0.000523			-0.00247	
		(0.018)			(0.018)			(0.016)	
Socio-economic PCA			-0.229***			0.210***			-0.168***
			(0.039)			(0.042)			(0.026)
Observations	45,974	45,974	45,974	26,007	26,007	26,007	45,974	45,974	45,974
R^2	0.462	0.470	0.461	0.658	0.664	0.656	0.268	0.271	0.268
CZ FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes
Senior Year FEs	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table shows the relationship between the three style dimensions (individualism, persistence, style novelty) and a set of socio-demographic controls. The data are a FIPS/Year panel ranging from 1950 to 2000 where the style indicators are merged with a FIPS/Year demographic dataset from Historical, Demographic, Economic, and Social Data available on ICSPR. Data is available at the decade/year level and is then linearly interpolated. We merge approximately 90% of counties-year coupling in the images data with ICSPR demographic data. All variables in these regressions are standardized; FEs are at the FIPS and year level, std. errors are clustered at the FIPS level.

Table 5: Predicting Patents

Panel A: LHS: Any applicant	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Style Novelty, Pooled, Dummy	0.162***	0.172***	0.0414***	0.0408***	0.0268*	0.0346***	0.0217*	0.303**
	(0.012)	(0.012)	(0.011)	(0.011)	(0.011)	(0.010)	(0.010)	(0.100)
Individuality, Pooled		-0.392***	0.0198	-0.105*	0.0739	-0.111*	0.0305	0.0326
		(0.047)	(0.041)	(0.044)	(0.045)	(0.050)	(0.050)	(0.050)
Persistence, Pooled, 20y lag (-2,+2)		-0.259***	-0.217***	-0.368***	-0.0581	-0.234***	0.0298	0.0286
		(0.042)	(0.037)	(0.041)	(0.047)	(0.043)	(0.048)	(0.048)
Population, Log				0.211***	0.247***	-0.0640	-0.0658	-0.0561
				(0.036)	(0.036)	(0.090)	(0.089)	(0.089)
Share of white people, CZ mean				0.329***	0.277***	0.551***	0.280**	0.277**
				(0.026)	(0.026)	(0.100)	(0.099)	(0.099)
Share of republican votes, CZ mean				0.00325	0.0382	-0.0588	-0.0462	-0.0466
				(0.027)	(0.030)	(0.031)	(0.038)	(0.038)
Senior Year				-0.000185	0.00258**	0.00179***	0.00440***	0.00442***
				(0.000)	(0.001)	(0.000)	(0.001)	(0.001)
Share of urban population in CZ				-0.0635***	-0.105***	0.385***	-0.0727	-0.0741
				(0.019)	(0.019)	(0.054)	(0.057)	(0.057)
Style Novelty, Pooled, Dummy=1 × Population, Log								-0.0232**
								(0.008)
Observations	13,867	13,867	13,867	13,867	13,867	13,867	13,867	13,867
R^2	0.013	0.018	0.253	0.263	0.290	0.421	0.443	0.443

Panel B: LHS: Any grantee	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Style Novelty, Pooled, Dummy	0.126***	0.142***	0.0302**	0.0215*	0.00784	0.0214*	0.00857	0.326***
	(0.011)	(0.011)	(0.010)	(0.010)	(0.010)	(0.010)	(0.010)	(0.095)
Individuality, Pooled		-0.347***	-0.00442	-0.00236	0.172***	-0.00393	0.129**	0.131**
		(0.043)	(0.039)	(0.041)	(0.041)	(0.048)	(0.047)	(0.047)
Persistence, Pooled, 20y lag (-2,+2)		-0.0896*	-0.0538	-0.309***	0.0528	-0.244***	0.0862	0.0849
		(0.039)	(0.035)	(0.038)	(0.044)	(0.041)	(0.045)	(0.045)
Population, Log				0.249***	0.284***	-0.0182	-0.0168	-0.00583
				(0.034)	(0.033)	(0.086)	(0.084)	(0.084)
Share of white people, CZ mean				0.230***	0.184***	0.253**	0.0283	0.0249
				(0.024)	(0.024)	(0.095)	(0.094)	(0.094)
Share of republican votes, CZ mean				0.0198	0.00469	-0.0298	-0.107**	-0.108**
				(0.025)	(0.028)	(0.030)	(0.036)	(0.036)
Senior Year				-0.00293***	-0.000821	-0.00176***	0.000338	0.000359
				(0.000)	(0.001)	(0.000)	(0.001)	(0.001)
Share of urban population in CZ				-0.0755***	-0.114***	0.276***	-0.179***	-0.181***
				(0.018)	(0.017)	(0.051)	(0.054)	(0.054)
Style Novelty, Pooled, Dummy=1 × Population, Log								-0.0261***
								(0.008)
Observations	13,867	13,867	13,867	13,867	13,867	13,867	13,867	13,867
R^2	0.009	0.015	0.222	0.242	0.269	0.381	0.406	0.407
CZ FEs	No	No	No	No	No	Yes	Yes	Yes
Senior Decade FEs	No	No	No	No	Yes	No	Yes	Yes
Pop Deciles FEs	No	No	Yes	Yes	Yes	Yes	Yes	Yes
Pop Deciles X Log Pop Controls	No	No	Yes	Yes	Yes	Yes	Yes	Yes

Notes: This table reports horse race regression results for applicant behavior and grantees on the three style dimensions, pooled gender. Panel A shows results for LHS variable any applicant, a dummy on whether a commuting zone/year exhibits at least one applicant; Panel B shows results for LHS any grantee, a dummy on whether a CZ/year exhibits at least one granted application. Different columns plug in different combinations of controls and fixed effects. Population controls and FEs are flexibly chosen, always including, from Column 3, log population, population deciles FEs and their interaction. Other fixed effects are CZ level, senior decade level FEs. The time period ranges from 1950 to 2000.

A Appendix Figures

Figure A.1: Example of Yearbook Images

Notes: The figure shows a sample page from a high school yearbook. Students in the example are 1959 graduating seniors from Tift High School in Tifton, Georgia.

Figure A.2: Sample Descriptives over Time

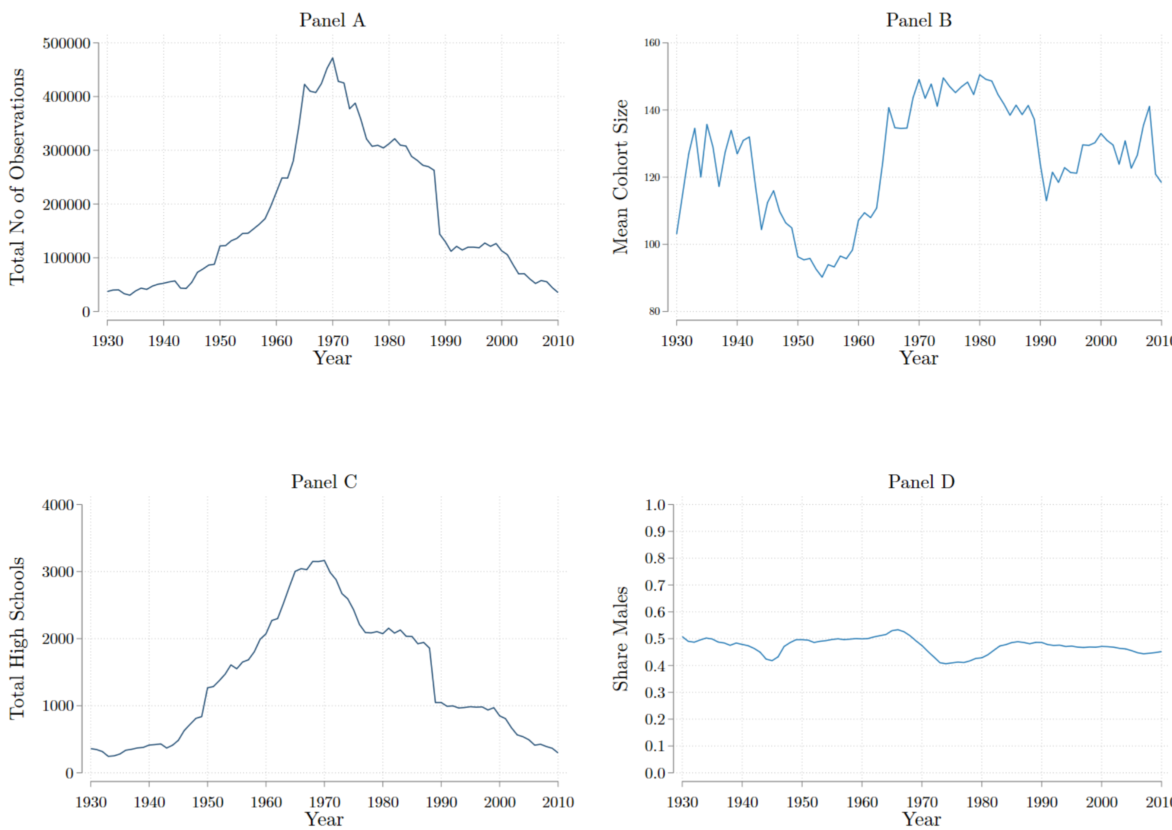

Notes: This figure shows sample descriptives over time. Panel A shows the total number of students that are sampled over time; Panel B plot plots the mean cohort size for a given high school/year pairing over time; Panel C shows the total number of high schools in observation in our sample; Panel D plots the share of male students in our sample.

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

(a) Males

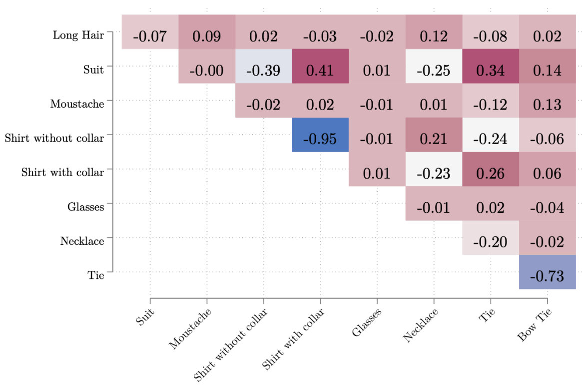

(b) Females

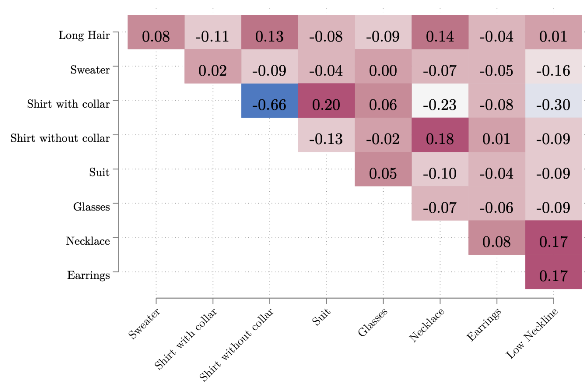

Notes: These matrix plots show correlations between a set of male and female attributes coming from multilabel predictions and the long/short hair classifiers. Attributes are selected as mostly male and female prevalent, where each style has to account for an average of at least 2% (for each gender) over the time period 1930-2010. Each observation is a single image with its corresponding style vector. Panel (a) shows male attributes correlations, Panel (b) shows female attributes correlations.

Figure A.4: Trends over Time

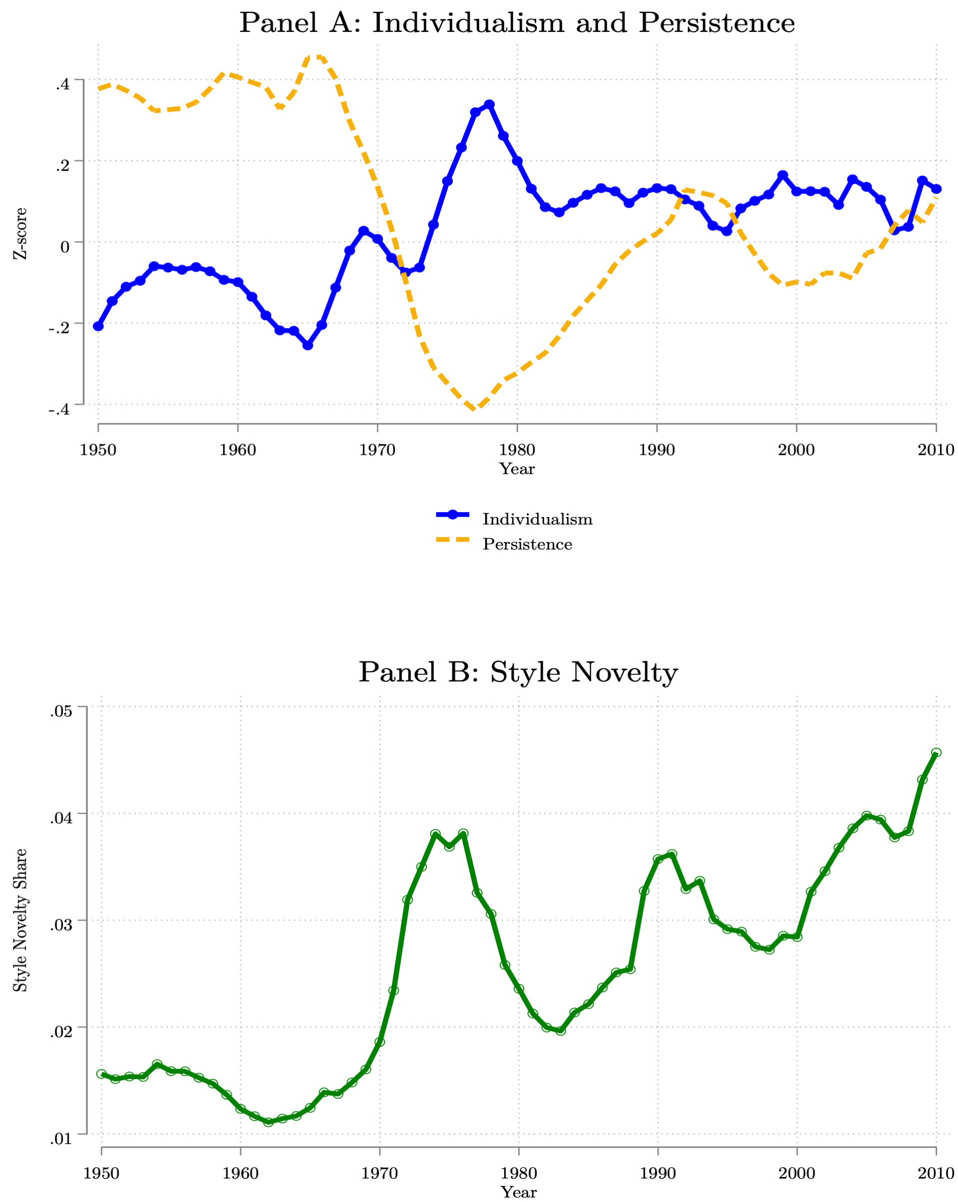

Notes: This figure plots yearly annual average scores for three style measures – individualism, persistence, and style novelty from our image level dataset (14.5 million observations). Panel A plots individualism and persistence as z-scores, measured on the left axis; Panel B shows style novelty as the share of innovators per year.

Figure A.5: Variation across Commuting Zones, by Gender

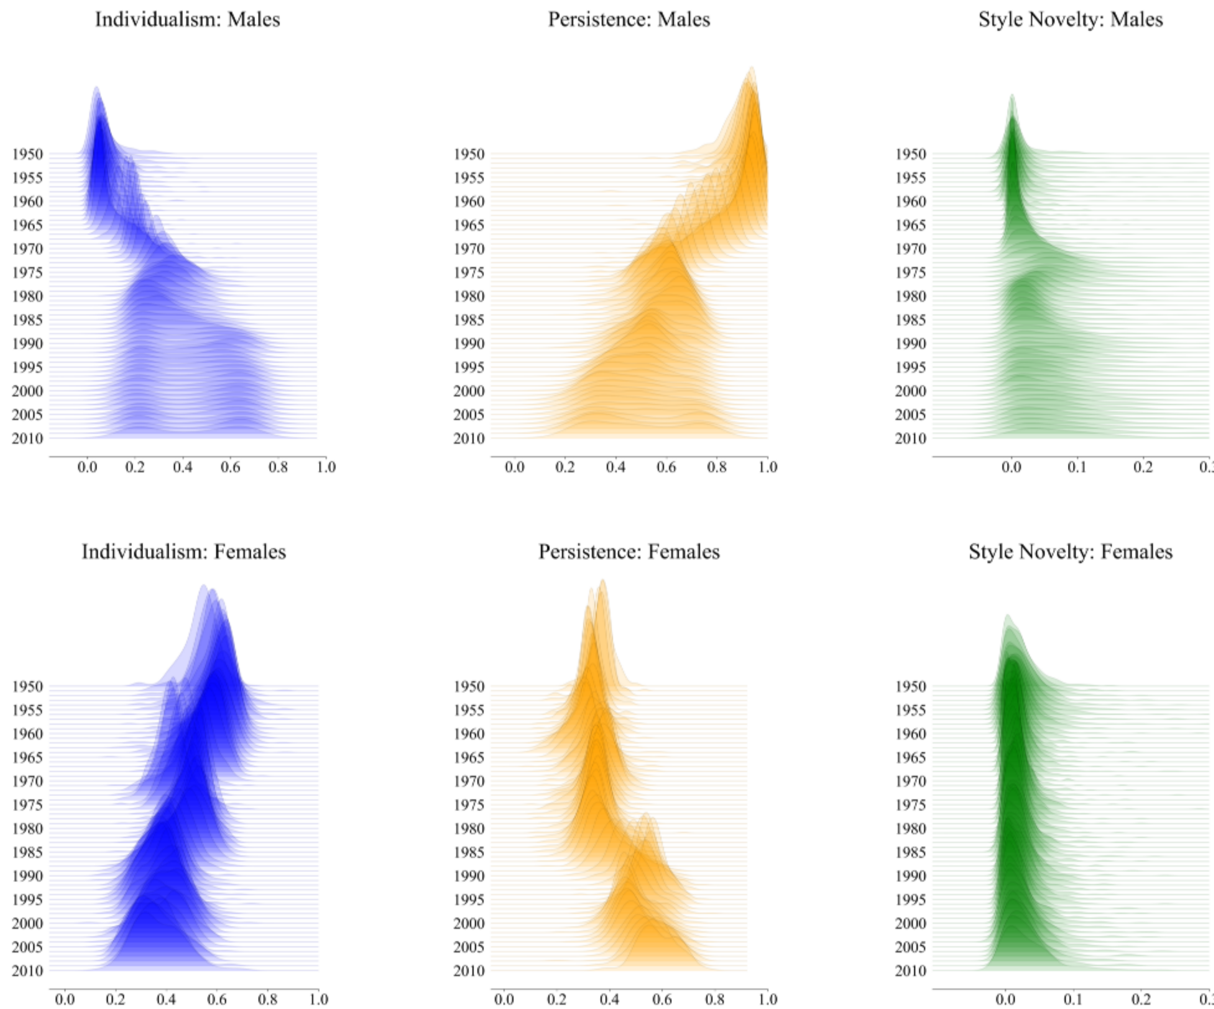

Notes: This figure shows ridge plots of the distributions of persistence, individualism, and style novelty across commuting zones, for the 1950-2010 senior high school year cohorts. For each commuting zone and cohort, the average value across the images is first calculated. We then display the range of commuting-zone average. For individualism and persistence, the values on the x-axis are cosine similarity averages within commuting zone-year couplings; for style novelty, x-axis values are the share of style innovators across commuting zone-year couplings. The top row shows results for males only. The second row shows results for females only.

Figure A.6: Persistence over Time

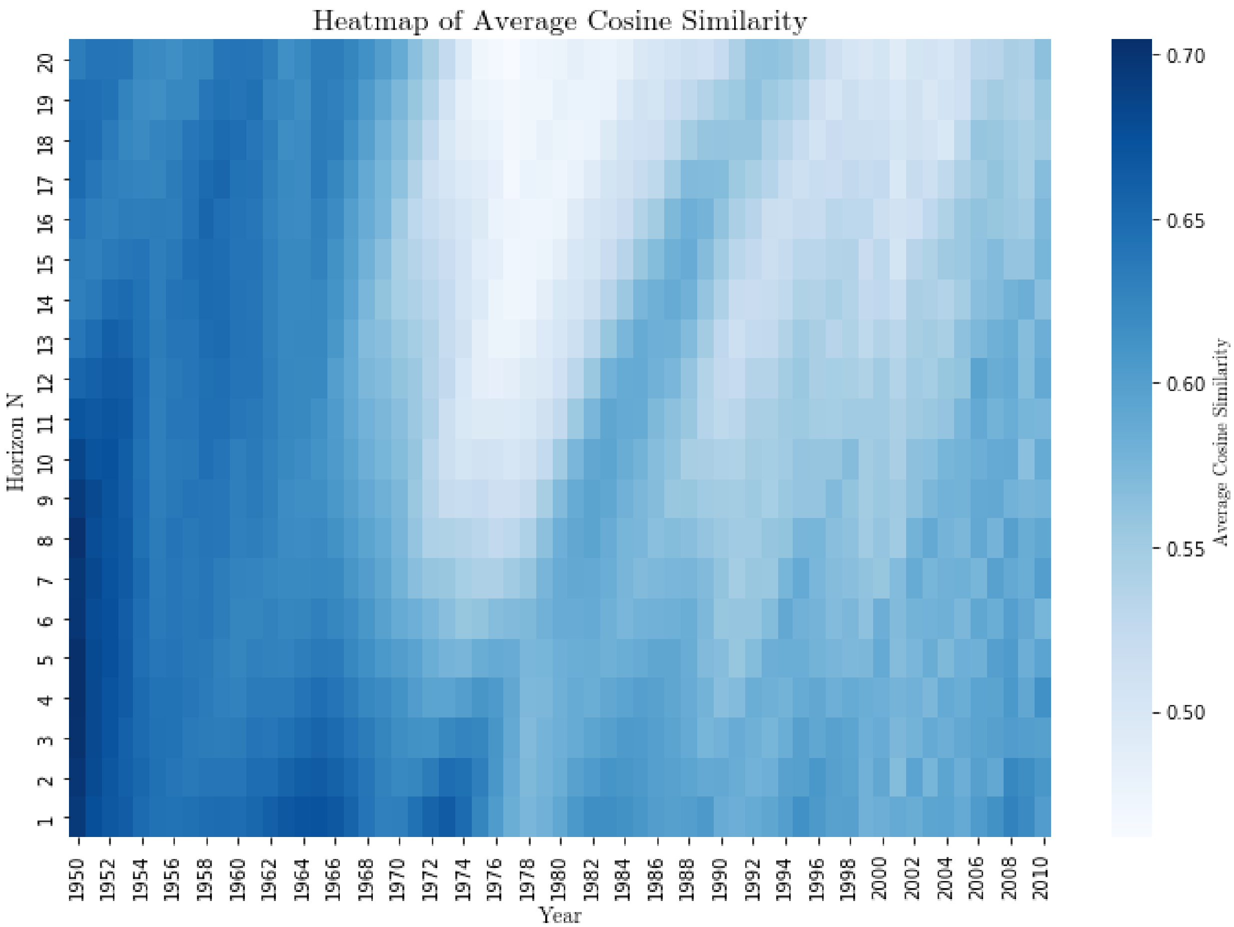

Notes: The figure shows the degree of cosine similarity between high school students in each year (on the x-axis) and cohorts in the same high school N years before (y-axis). Darker colors indicate higher similarity. Similarity is calculated separately for men and women, and then averaged. Dramatic style change – and markedly lower persistence at any horizon – is visible from 1970 onwards.

Figure A.7: Drivers of Persistence

(a) Males

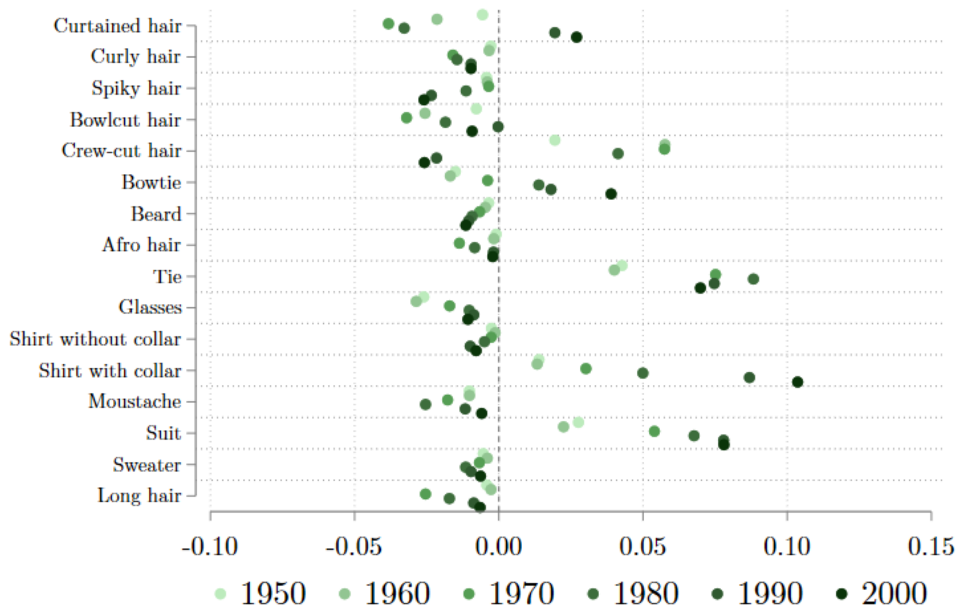

(b) Females

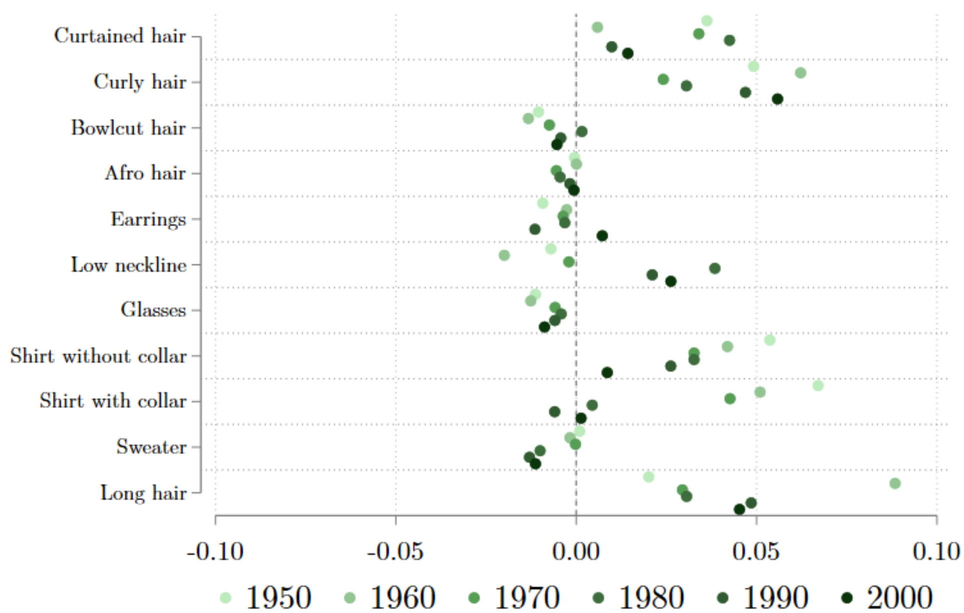

Notes: Dots represent decade-level coefficient estimates from an OLS Lasso regression of persistence score on gender selected relevant style attributes. Each dot represents a single coefficient for a style, derived from bi-variate regressions. Values greater zero indicate a positive contribution to students' persistence scores. Panel (a) is for male images only, Panel (b) for females.

Figure A.8: Drivers of Style Novelty

(a) Males

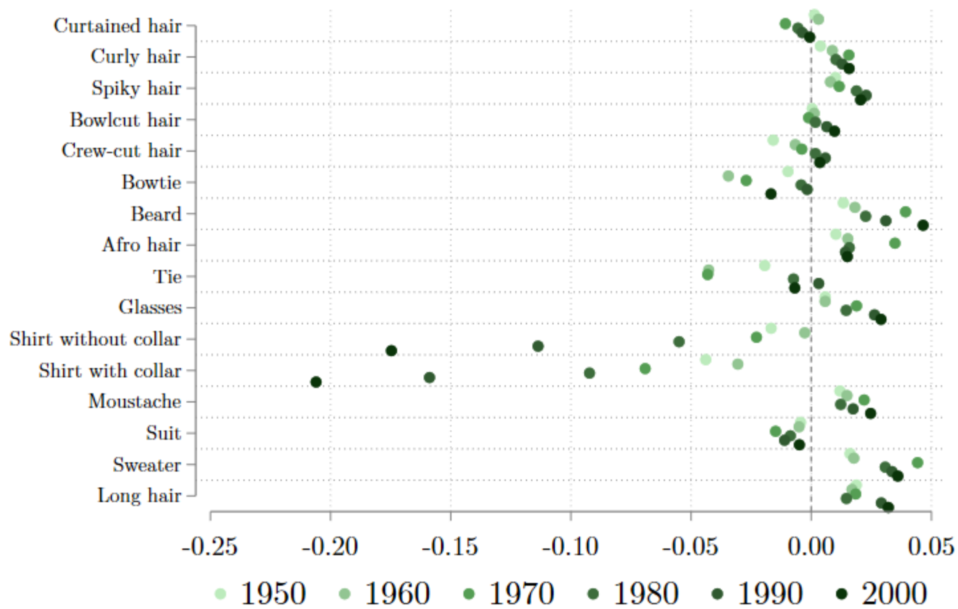

(b) Females

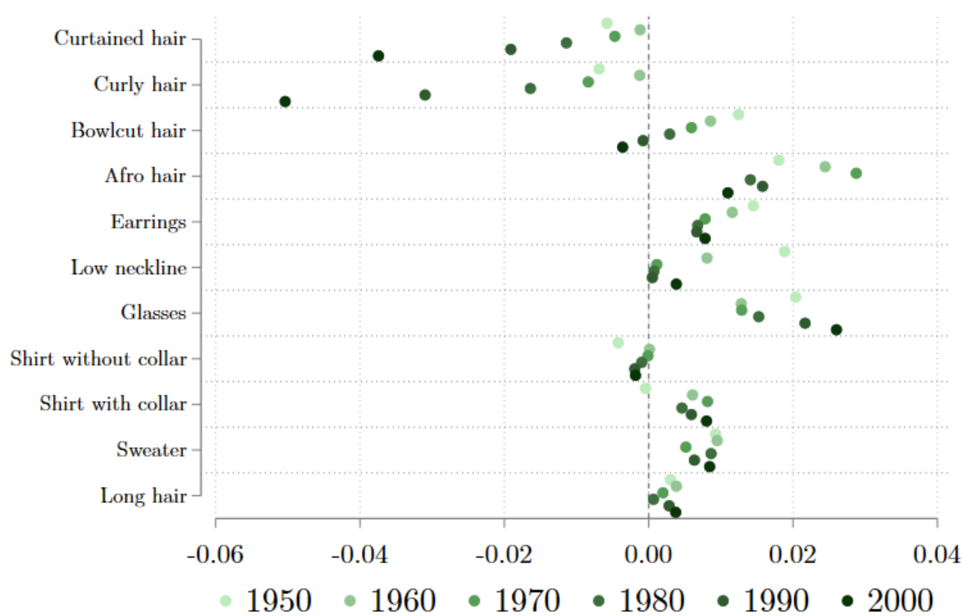

Notes: Dots represent decade-level coefficient estimates from an OLS Lasso regression of top 1% style innovation dummies on gender selected relevant style attributes. Each dot represents a single coefficient for a style, derived from bi-variate regressions. Values greater zero indicate a positive contribution to students' style innovation. Panel A is for male images only, Panel B for females.

Figure A.9: Name vs Style-based Individualsim

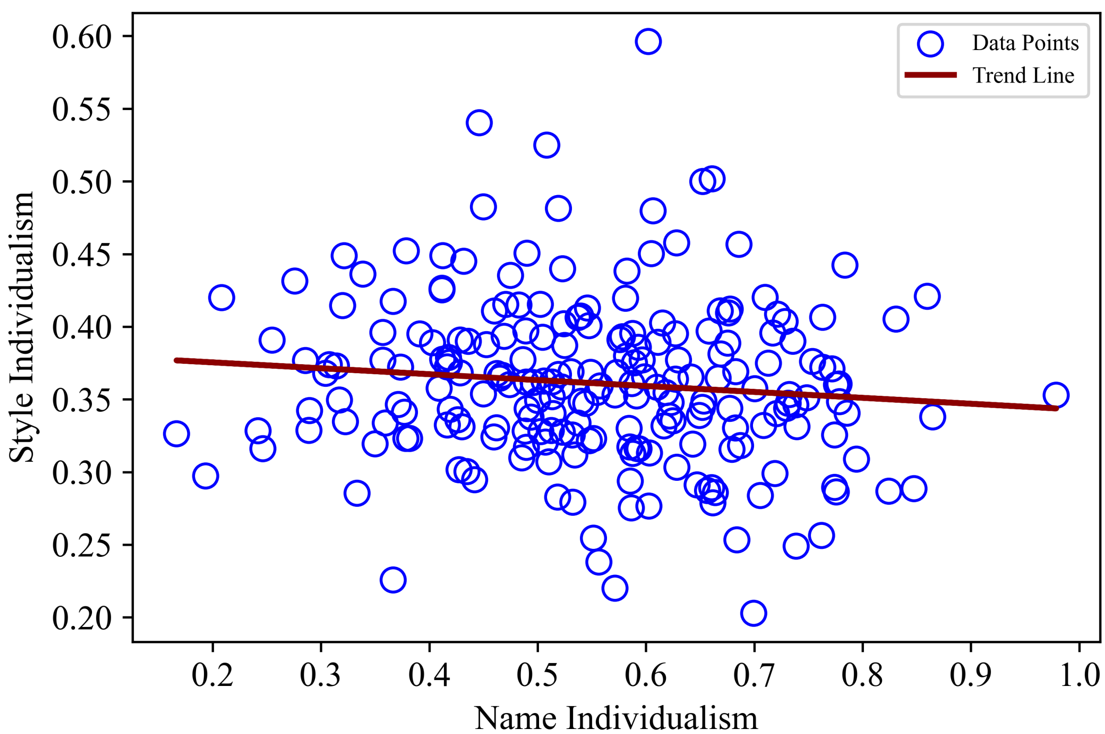

Notes: We plot 1-share of common names, in 220 high schools during the 1960s, and compare this measure with style-based individualism in the same high schools. The beta coefficient for the relationship between these two measures is -0.04 and the $R^2 = 0.01$.

Figure A.10: Correlations between Style Measures

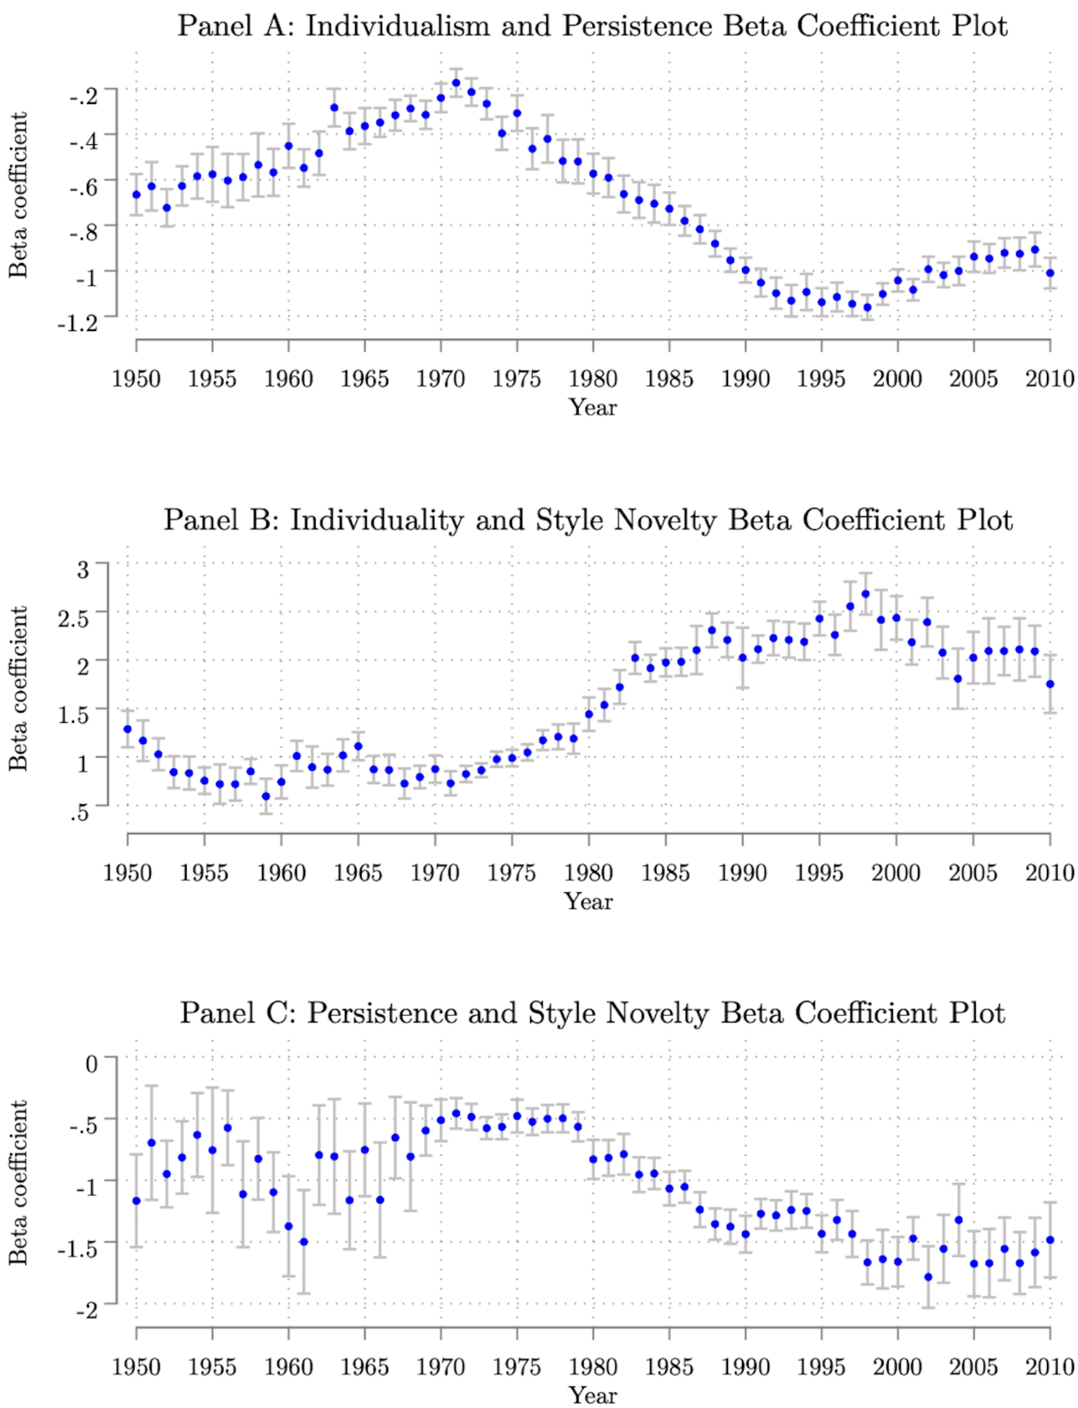

Notes: These three plots show the coefficients from regressing, year by year, individualism and persistence (Panel A), individualism and style novelty (Panel B), and persistence and style novelty (Panel C), in the cross-section of High Schools. Confidence intervals reflect robust standard errors.

Figure A.11: Granted Patents, by Year since Graduation

Panel A: Share of Patents Granted by Style Novelty and Gender

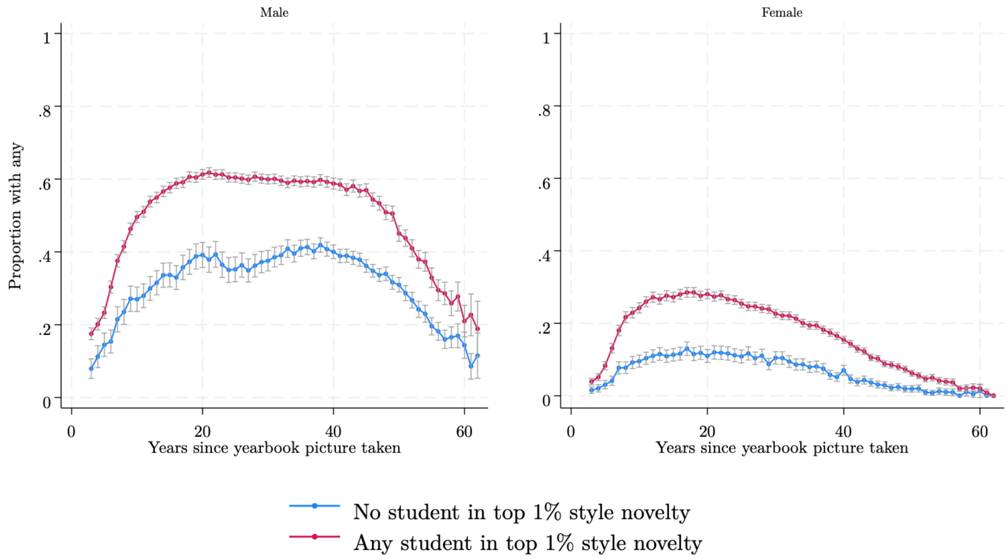

Panel B: Number of Granted Patents per Individual by Style Novelty and Gender

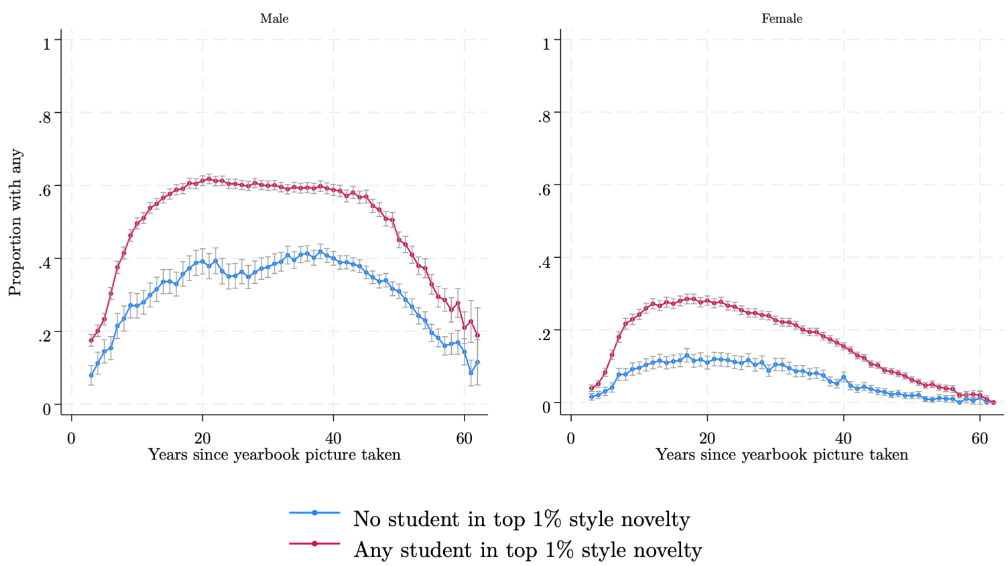

Notes: This figure shows granted applications by year since graduation, for two groups of high school graduates – style innovators and non-style innovators, split by gender. Panel A shows the share of patents subsequently granted (any vs none, described by a dummy variable) in a CZ-year; Panel B plots the number of granted patents per individual (any vs none, indicated by a dummy variable) in a CZ-year. Year of birth and CZ of an innovator was merged with high school and an imputed year of birth of students, calculated as year of the yearbook – 17. Red lines indicate the presence of at least one top 1% style innovator in a commuting zone-cohort. Time period is 1950 to 2000.

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 out 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).

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 and do not contain any identifying information concerning the individuals depicted. 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 Selecting senior portrait pages

First, we turn yearbook pages into series of images by identifying consecutive runs of at least 4 pages that share similar image size and consistent use of color/no color. Second, we use a number of identifiers to give each run a score of the likelihood of being a senior page. Yearbook appearances are highly heterogeneous, but we can use a few key generalizations about yearbook characteristics to identify senior portrait pages. We use optical character recognition and convolutional neural network assessment of the size, shape and color of the images to evaluate the following:

- pages with seniors will likely be identified with the word "senior" and do not contain other role names such as "faculty", "teachers", or "juniors".
- senior pictures are more likely to be in color.
- senior pictures are often larger than the rest of those in a yearbook;
- pages with seniors on them will often say "Class of [year]" where [year] matches the year of the yearbook's publication;

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

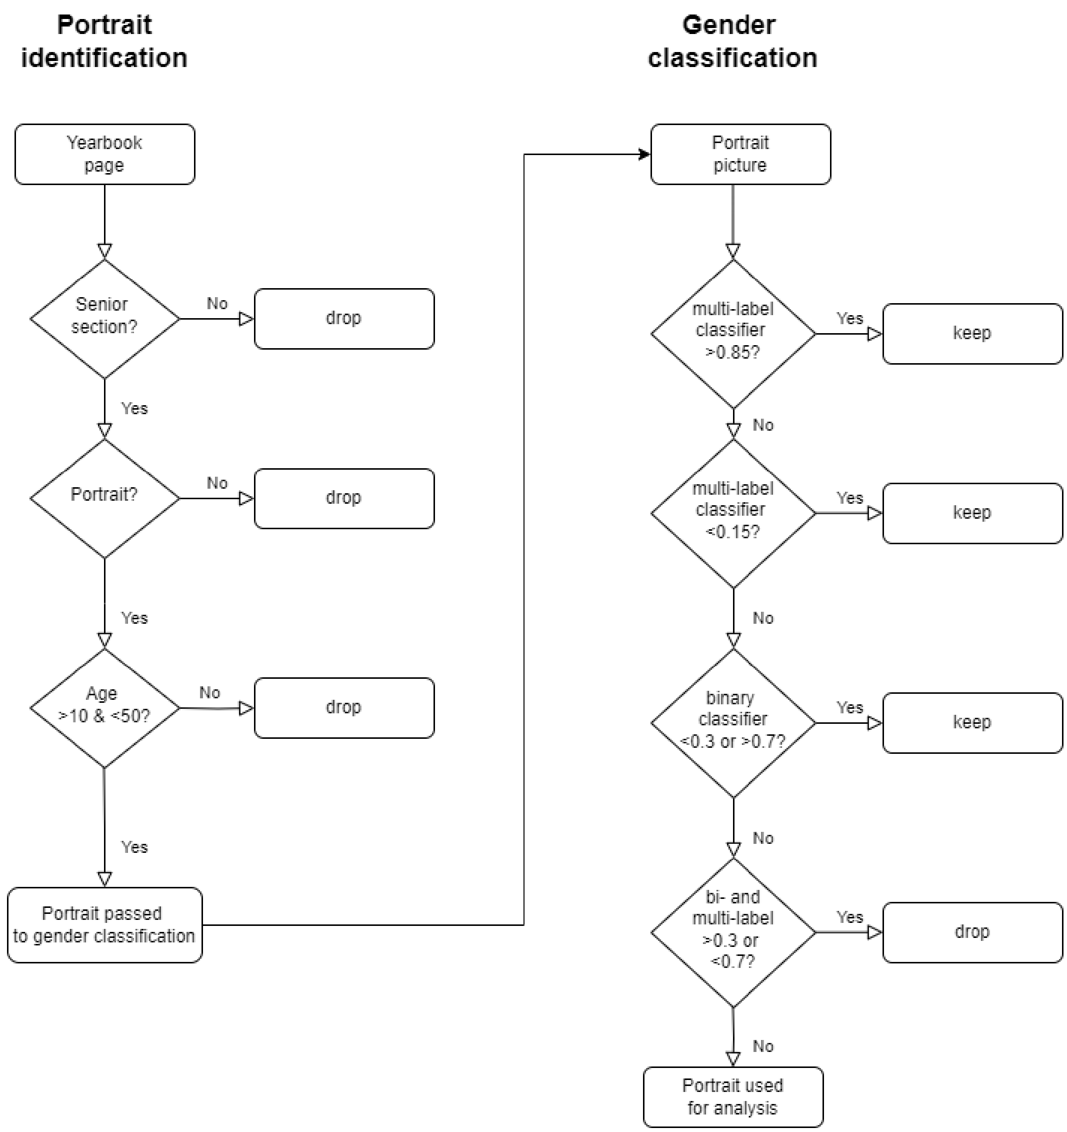

Notes: The flowcharts depict the steps from raw yearbook page to portrait picture used for further analysis. For the gender classification, 0 is equivalent to the classifier being very certain of a female portrait, and 1, of a male one.

Scoring these features creates a candidate list of senior portrait pages. Next, we train a Google Vertex AI image classification model to assign actual senior page status, using a hand-labelled training dataset. Table B.1 gives the results. This classifier correctly identified 263 out of 286 (92%) true senior pages, and correctly rejected 2,998 of "not senior" 3,126 pages (96%). The classifier missed slightly less than a third of true positives (128/391). We deliberately use a high confidence threshold to minimize the number of false positives. 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	2,998	25	3,023
- senior page	1	128	263	391
	total	3,126	288	3,414

B.3 Identifying portraits

The yearbook images are transformed into grayscale and the border around the image is whitened. We define the background area of the image as the area characterized by the brightest color point. We convert all background pixels without portraits that lie within a brightness interval to white. Next, we turn all portraits into black rectanges. Our portrait recognition algorithm checks whether these rectangles satisfy threshold criteria such as whether their area covers 1-33% of a page or the height and width cover 10-50% of the page. We crop the drawn rectangles in a page if three criteria are met:

- the total number of faces (recognized by the cv2 face detection model in Python) recognized on a yearbook page vs is not greater than 1.2 times the number of images on the same page;
- at least 50% images only contain one face, and the face covers more than 30% of that image;
- there are at least two images on the same page.

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 with suit and tie will not be considered "innovative´´ in 1950s America, but a woman would have been very unusual, the example of style icons like Marlene Dietrich not withstanding. To this end, we adopted a decision path that avoids false positives as much as possible – even at the risk of dropping more images from our sample. We combined predictions from the multi-label and binary models (see Appendix C for details). We first classified all images through the multi-label classifier, obtaining a masculinity score (0 to 1). If the model prediction showed high confidence (score for male ≤ 0.15 or ≥ 0.85), we kept the label. For images which scored in the interval 0.15-0.85, we used the binary model prediction unless both models were unsure (predictions in the range 0.3 and 0.7); these observations were dropped (Figure B.1).

B.6 Geocoding high school locations

To geocode the high schools in the nationally representative 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: Age Distribution by Decade

age interval decade	<3	3 – 9	10 – 19	20 – 29	30 – 39	40 – 49	50 – 59	60 – 69	>69	Total
1930s	6 (0.0%)	9,508 (2.0%)	110,254 (23.0%)	301,616 (62.9%)	28,613 (6.0%)	19,068 (4.0%)	8,898 (1.9%)	1,380 (0.3%)	105 (0.0%)	479,448 (100.0%)
1940s	10 (0.0%)	12,959 (1.7%)	149,605 (19.1%)	524,990 (66.9%)	45,337 (5.8%)	30,218 (3.8%)	18,813 (2.4%)	2,756 (0.4%)	239 (0.0%)	784,927 (100.0%)
1950s	91 (0.0%)	34,345 (1.8%)	486,457 (25.5%)	1,147,483 (60.2%)	116,182 (6.1%)	67,550 (3.5%)	45,710 (2.4%)	7,526 (0.4%)	954 (0.1%)	1,906,298 (100.0%)
1960s	136 (0.0%)	137,762 (3.2%)	1,247,651 (29.0%)	2,430,737 (56.5%)	294,402 (6.8%)	132,237 (3.1%)	53,581 (1.2%)	7,395 (0.2%)	1,451 (0.0%)	4,305,352 (100.0%)
1970s	320 (0.0%)	71,258 (1.6%)	931,178 (21.0%)	3,118,976 (70.3%)	169,051 (3.8%)	85,689 (1.9%)	51,320 (1.2%)	8,379 (0.2%)	1,502 (0.0%)	4,437,673 (100.0%)
1980s	544 (0.0%)	105,609 (3.2%)	875,914 (26.3%)	1,985,385 (59.5%)	199,716 (6.0%)	91,525 (2.7%)	62,144 (1.9%)	11,943 (0.4%)	2,605 (0.1%)	3,335,385 (100.0%)
1990s	710 (0.1%)	32,214 (2.3%)	311,770 (22.4%)	873,151 (62.7%)	106,190 (7.6%)	47,984 (3.4%)	17,935 (1.3%)	2,137 (0.2%)	424 (0.0%)	1,392,515 (100.0%)
2000s	1,556 (0.2%)	12,548 (1.5%)	175,041 (21.1%)	536,318 (64.5%)	62,410 (7.5%)	30,016 (3.6%)	11,703 (1.4%)	1,527 (0.2%)	164 (0.0%)	831,283 (100.0%)
2010s	143 (0.3%)	762 (1.8%)	8,383 (19.3%)	28,071 (64.7%)	3,267 (7.5%)	1,800 (4.2%)	789 (1.8%)	131 (0.3%)	22 (0.1%)	43,368 (100.0%)
Total	3,516 (0.0%)	416,965 (2.4%)	4,296,253 (24.5%)	10,946,727 (62.5%)	1,025,168 (5.9%)	506,087 (2.9%)	270,893 (1.5%)	43,174 (0.2%)	7,466 (0.0%)	17,516,249 (100.0%)

Note: The table gives the distribution of age classifications by decade in the original 17.5 million observations dataset prior to age and gender filtering.

Table B.3: Yearbooks by State

State	**Yearbook Count**	**State**	**Yearbook Count**
Alabama	2822	Nebraska	1547
Arizona	1265	Nevada	525
Arkansas	2762	New Hampshire	706
California	11863	New Jersey	5416
Colorado	1613	New Mexico	934
Connecticut	2728	New York	1145
Delaware	436	North Carolina	2742
Florida	3896	North Dakota	728
Georgia	3318	Ohio	1045
Idaho	894	Oklahoma	701
Illinois	10710	Oregon	722
Indiana	8367	Pennsylvania	1459
Iowa	3782	Rhode Island	309
Kansas	3043	South Carolina	361
Kentucky	2457	South Dakota	421
Louisiana	1490	Tennessee	627
Maine	1818	Texas	681
Maryland	2155	Utah	248
Massachusetts	3402	Vermont	239
Michigan	8808	Virginia	559
Minnesota	3412	Washington	908
Mississippi	1465	West Virginia	650
Missouri	4517	Wisconsin	1007
Montana	900	Wyoming	188

Notes: This table shows the count of high school yearbooks used by state. Data includes yearbooks from 1930 to 2010.

C Appendix: Style Feature Classification

Images were manually labeled to create a training set. We attached labels to each image as detailed in Table B.4:

Table B.4: 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
Moustache	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 and gender. To do so, we use four separate classifiers: a) gender (male/female), b) hair length (long/short), c) multi-label hairstyle classifier (6 styles),[9] and c) a multi-label classifier for other style items. We train these classifiers on hand-labelled data. For the gender classifier, we use 2,188 images, of which 1,751 are for training, 219 for validation, and 218 for testing. For the hair classifier, we use 2,127 images – 1,702 for training, 213 for validation, and 212 for testing. Fr 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 four 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

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

glasses". In the case of single-label classifiers (hair, gender) we use a cut-off of 0.5; for the hairstyle multi-label classifier, we used a cut-off of 0.5 as well; for the style item multi-label, we use a 0.3 threshold. The decision process for gender is described above in Appendix B. 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.

Figure C.1: Precision-Recall Curves

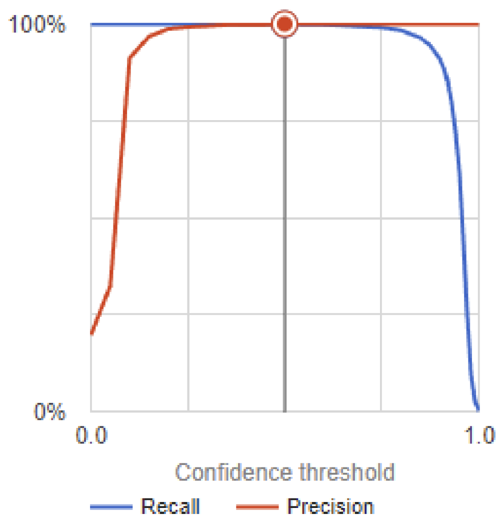

(a) **Multi-label model:** predictions on bowtie, tie, beard, afro, low neckline, male, earrings, tie, glasses, shirt w/o collar, shirt with collar, mortarboard

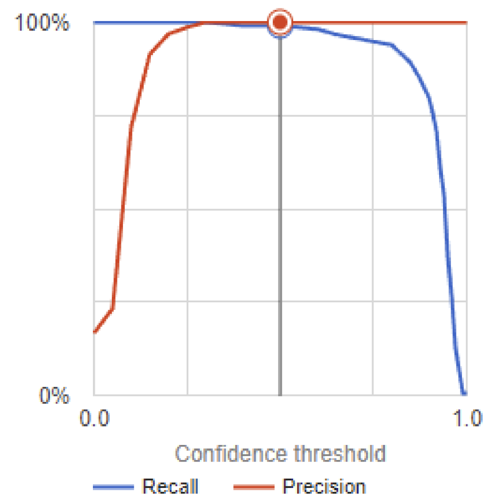

(b) **Hair classifier:** predictions on afro, curtained, curly, spiky, bowl cut, crew cut

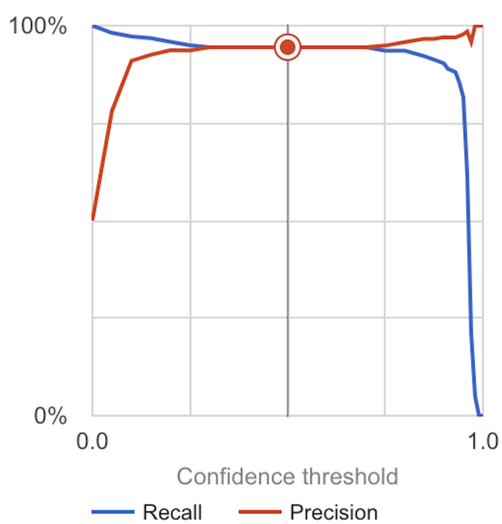

(c) **Gender classifier:** predictions on male, female

Notes: These graphs show precision and recall scores on the y-axis, for the three main classifiers used to categorize images. Panel (a) shows results for the multi-label model, Panel (b) for the hairstyle classifier, Panel (c) for gender.

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.[10] We find a high level of accuracy overall, ranging from 79-98% by category. Agreement among the human evaluators is higher than the agreement algorithm-human, but the differences range from less than 20% (necklace) to less than 1% (moustache).

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

column	(1)	(2)	(3)	(4)	(5)
Attribute	**Sub-categories**	**Accuracy (Specific)**	**% of Human Agreement**	**Algo Accuracy (Aggregated)**	**Training**
Glasses		95.98%	98.78%	96%	2500 images per category (i.e. glasses yes/no)
Ties	Tie	85.20%	98.32%	90%	1338 images per category (i.e. tie/bow tie/no tie)
	Bowtie	93.87%	98.99%		
Jewelry	Necklace	76.96%	96.55%	84%	2500 images per category (i.e. necklace yes/no)
	Earrings	90.27%	97.77%		
Hair	Curtained	76.74%	81.47%	83%	2500 images per category (i.e. short/long hair)
	Curly	67.44%	86.64%		
	Spiky	95.98%	96.01%		
	Bowl cut	94.29%	93.91%		
	Crew cut	71.46%	90.76%		
	Long hair	76.96%	87.10%		
	Afro	97.89%	97.77%		
Clothes	Shirt without collar	70.61%	87.06%	79%	1250 images per category (i.e. suit yes/no)
	Shirt with collar	72.52%	90.17%		
	Suit	79.70%	97.61%		
	Sweater	86.89%	93.53%		
	Low neckline	85.41%	90.04%		
Facial Hair	Beard	97.89%	99.50%	98%	525 images per category (i.e. beard yes/no)
	Moustache only	98.52%	98.91%		

Notes: 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.

Figure D.1, Panel (a), shows the confusion matrix for the first-pass gender classification, using the multi-label classifier predictions; Panel (b) does so for the second pass, after using a specialized gender classifier in addition.[11] We improve accuracy from 95.8% to 98.75% in our human audit exercise.

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

[11] Note that the accuracy in Figure C.1 (c) concerns the single-label gender classifier alone, not the 2-stage process as outlined here.

Figure D.1: Gender Prediction Performance - Confusion Matrices

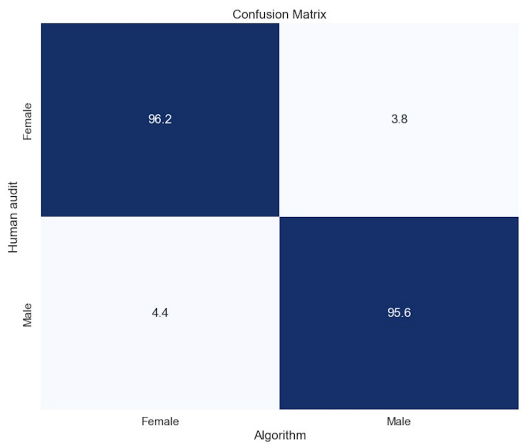

(a) **Stage 1 performance**

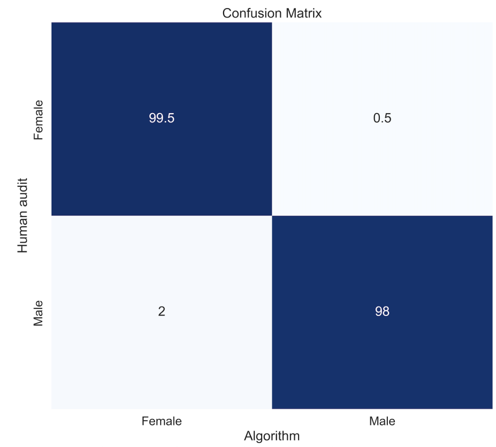

(b) **Stage 2 performance**

Notes: The figures compare prediction accuracy (out of sample) for two classifier solutions – a) the output of the multi-label algorithm and b) the two-stage filtering approach described in the text. We survey 4 people through Prolific for 2 different batches of images. The first batch contains 485 images; the second 463.

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.2). 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.3 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.2. For each comparison set, we collected responses from 60 distinct human evaluators. Initially, we determined the mode of the human responses

Figure D.2: Sample Screen - Prolific

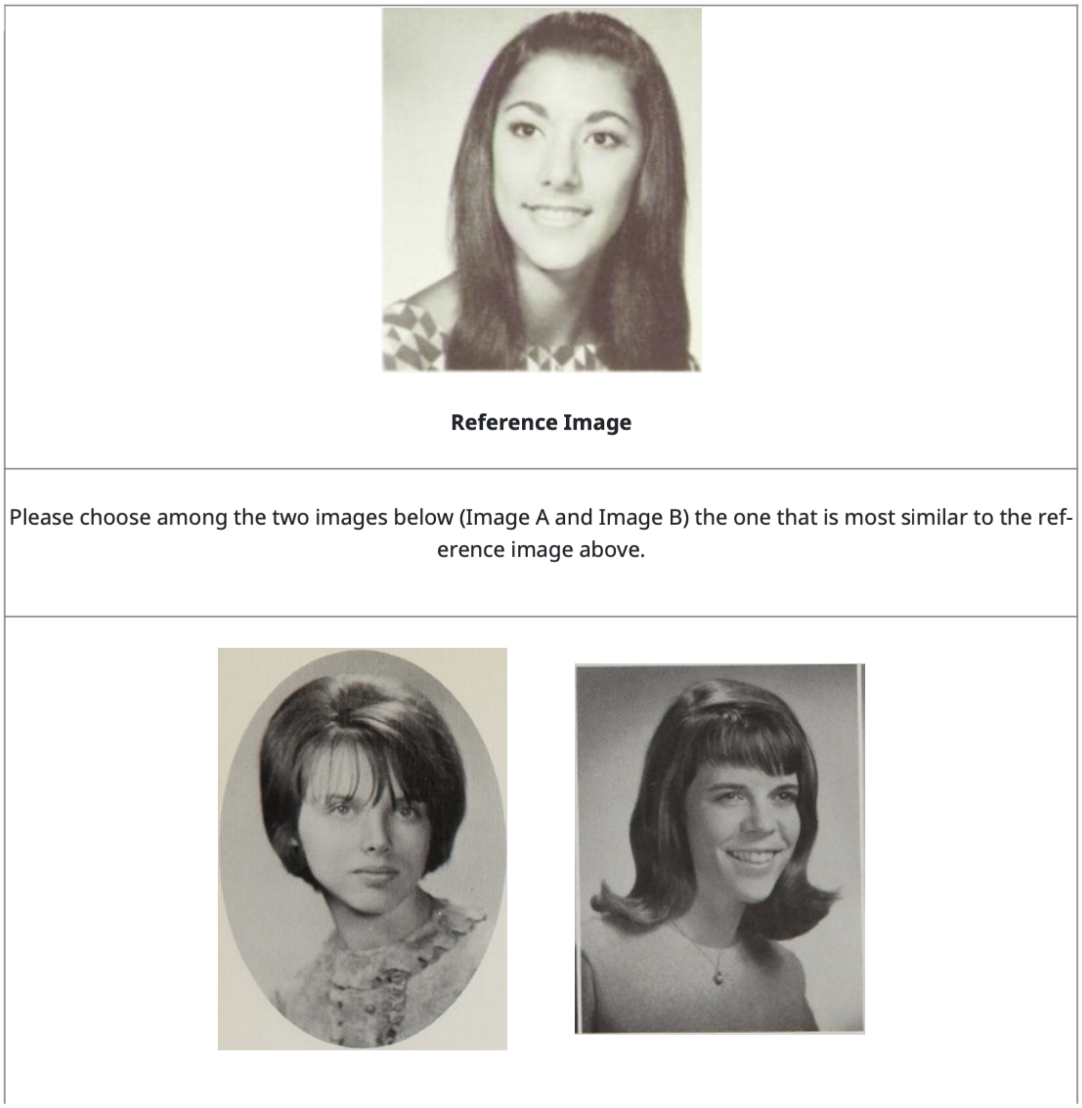

Notes: The figure shows examples of reference and comparison pictures. Through the online survey provider Prolific, we recruited 60 evaluators, who were asked to judge similarity of style choice. In particular, they were asked to decide whether picture A or B is more similar to the reference image.

Figure D.3: Agreement between Humans and Humans/Algorithm

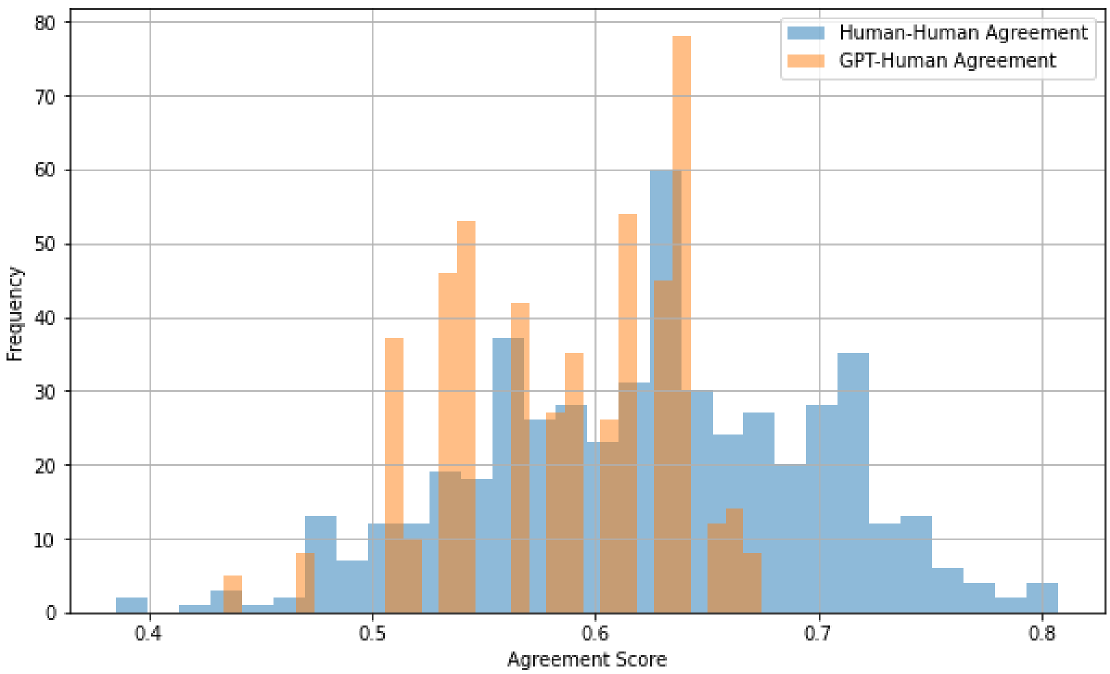

Notes: Each value on the x-axis indicated the % of answers where human-human or human-algorithm agree with each other. The height of the bars represents the frequency of agreement. The data are comprised of 500 human-human and 500 human-algorithm comparisons.

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.

Table D.2: Agreement between Human Coders and Algorithm

Question code	**Human Coder**				**Modal Human Choice**	**% Human Agreement**	**Algorithm Choice**
	(1)	(2)	...	(60)			
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

In the spirit of the famous "Turing test", we examine whether the cosine similarity calculations perform similarly to human evaluators in terms of accuracy. Table D.3 gives the range of agreement between humans, and the extent to which the algorithm agrees with the modal human choice. Humans also 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 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