IMAGE(S)*

Hans-Joachim Voth
U Zurich

David Yanagizawa-Drott
U Zurich

April 13, 2023

Abstract. From clothes and hairstyles to fashion accessories, humans use a great range of stylistic elements to express themselves, impress others, demonstrate their individuality, or show that they belong to a group. Despite its central importance as a form of social interaction and self-expression, and a rich body of theoretical work, empirical work on fashion and style choices is rare. We present new methods to use images as a high-frequency, granular source for the analysis of cultural change. 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 11 million high school yearbook pictures of graduating US seniors to analyze persistence and change in style. We use the arrival of the Beatles in the United States in 1964 as a case study to demonstrate the potential of image analysis to detect cultural innovation and diffusion.

JEL codes: A12, Z13, N34
Keywords: Culture, images, persistence, cultural change, similarity, counter-culture, hippies.

* We gratefully acknowledge financial support via SNF grant 100018-208092. We thank Marianne Bertrandt, Leonardo Bursztyn and Emir Kamenica for advice, and Igli Bajo, Emilia Bullano, James Godlonton, Ewan Rawcliffe, Max Read, Rosella de Sabatta, and Vanessa Stichler for outstanding research assistance.

The image we project, the style we chose, the clothes we wear, the haircut we have, and the glasses we pick reflect personal preferences. Do I express my individuality, and decide to look completely different from anyone around me? Do I wear the same suit and tie as everyone else? Every one of us, every day, makes choices of personal style. But the clothes we wear and the styles we pick are not only driven by personal preferences. They also reflect society-wide trends (Hancock, Johnson-Woods, and Karaminas 2013). Firms launch products (Pesendorfer 1995), images circulate in the media, and millions of individual, everyday choices jointly determine the visual culture of an age. At any given time, the choice-set is limited. Some of these limits are formal: Many religions closely circumscribe what is considered acceptable dress, especially in the case of women; sumptuary laws routinely determined which groups in society were allowed to wear what (Riello and Rublack 2019). Other restrictions reflect social norms: Is it permissible to deviate from what others wear? Can a man have long hair, a woman short hair? Do men wear wigs or tights?

Stylistic choices are part of everyday culture. Culture often persists over long periods, with everyday practices reflecting economic incentives and cultural factors centuries ago (Guiso, Sapienza, and Zingales 2016; Alesina, Giuliano, and Nunn 2011; S. O. Becker et al. 2016). At the same time, it can also change rapidly. Attitudes towards pre-marital sex, smoking, homophobia, and racial views have been transformed in many Western countries in recent decades (Fernández, Parsa, and Viarengo 2019; Giuliano and Nunn 2021). One key challenge in the analysis of cultural change is a lack of high-frequency, granular data. Survey measures often use relatively small samples, and disaggregation to the state or local level reduces cell size still further. Another challenge is the lack of consistent measurements.[1] In many models of group behavior (Kuran 1989; Chamley 2004), waves of change can be initiated by a few individuals. What is missing in the analysis of cultural *change* is a consistently measured, high resolution indicator that is observed with sufficient

[1] For example, in the GSS, some questions are added from wave to wave, while others are dropped.

frequency over time and space. While recent research has made important progress analyzing themes in folk tales (Michalopoulos and Xue 2021) or using surveys to measure rates of time preference (Sunde et al. 2022), we still lack good indicators of cultural change for most periods and countries.

In this paper, we introduce a new source to measure cultural innovations at high frequency and at the granular level, exploiting a rich dataset of High School senior images. We develop methods to analyze both persistence and conformity in portraits, and apply them to the case of the US, 1940-2010. Looking at a single country in the more recent past is motivated by several considerations: In the more distant past, sumptuary laws were regularly used to regulate who is allowed to wear what; it is only in recent centuries that individual choices without formal restrictions become a dominant factor in fashion. This leads us to focus on the more recent past. The advent of photography greatly facilitates access to accurate representations of actual clothing and hair styles. Arguably, the US was at the forefront of cultural change in the second half of the 20th century. Its size and diversity allow for meaningful cross-sectional analysis. We focus on high school senior images for reasons of data availability, and because they capture how students chose to portray themselves to others (and posterity) at an important inflection point of their lives – when their secondary education comes to an end, and they either go off to college or join the labor force. US high school yearbooks are typically sold to the entire student body and have high sentimental value.

To demonstrate the usefulness of this new data source and of the methods we develop, we present three case studies. First, we examine the rise and fall of conformity and persistence in styles over time in the US. Within each graduating high school class, we examine how different styles are, and document the timing, speed, and location of the decisive changes that began in the late sixties, following the Woodstock festival and the "summer of love" in 1968. We show that counter-culture – defined as a low level of conformism and low persistence of styles vis-à-vis the parent generation – peaked in the mid-seventies. Since then, many of these trends have reversed, leading to more conformity and growing persistence.

What does our new data say about the arrival of new trends? Which ones "go viral" and which ones flop? We draw on recent work in the economics of innovation (Kelly et al. 2021) to create a measure of influential innovation. For the case of California, we examine where such innovation occurred, and show that innovation is most pronounced in major cities. We observe San Francisco as an epicenter of style innovation in the early 70s, in keeping with its role as a focal point of counterculture in the late 60s.

In our third application we explore whether our methods and source can be used to document and trace the impact of a single, major event that is widely credited with changing US visual culture – the arrival of the Beatles in 1964. While still largely unknown in 1963, the Beatles appeared four times on the popular CBS Ed Sullivan show in the spring of 1964. They immediately vaulted to the top of the charts, and "Beatle-mania" swept the US. Popular commentary focused as much on the group's hair style as on their music. We train a classifier to detect the particular hair style in our images. In the same year as the broadcast, the areas close to CBS stations, the share of men who sport the Beatles-style look in their senior portrait jumps, but stays constant elsewhere. Areas served by other TV stations show no significant jump. After 1964, the mean probability of a Beatles-style "mop-top" starts to rise everywhere; within a few years, the mean probability of a "mop-top" hairstyle increased by 25 pp, suggesting that lots of high school seniors had adopted the hair style of the Fabulous Four from Liverpool.

Interest in using photographs and regularities in appearance as a source for social science analysis dates back to the 19th century. A few decades after the invention of photography, Francis Galton (1878) famously sought to identify the typical facial characteristics of criminals, mentally ill, and prostitutes by superimposing multiple portrait photographs. Systematic analysis of images in computer science has taken off in the last 20 years.[2] For example, Hidayati et al. (2014) analyze fashion trends based on New York fashion week couture based on a

[2] Ginosar et al. (2015) apply a similar approach to contemporary student portraits, and examine the display of emotions over time.

classification algorithm, and Lee et al. (2015) extract information on agricultural trends from images. 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. In economics, image analysis has recently begun to be used, where researchers have evaluated the effect of facial features on judges' bail decisions or on lending behavior (Athey et al. 2022; Ludwig and Mullainathan 2023).

Our paper is the first to use a large sample of images as a source for analyzing cultural change. We do so through the lens of fashion. Analyzing fashion as an indicator of social change has a long lineage. Theodor Veblen (1973) famously analyzed conspicuous consumption as a marker of social distinction. Georg Simmel's (1957) essay about the "Philosophy of Fashion" already emphasized how adopting particular styles was mostly about conveying one's own position, standing, and taste: "Judging from the ugly and repugnant things that are sometimes en vogue, it would seem as though fashion were desirous of exhibiting its power by getting us to adopt the most atrocious things for its sake alone." Economic analysis of fashion goes back to the classic work by 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 microfounded model of such consumption externalities. Matsuyama (1991) explicitly models groups of conformists and non-conformists, and shows that demand for goods can fluctuate cyclically ("fashion cycles"). 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 similiarities of style and fashion over time and space.

1. History and Background

Portraits and photography. Clothing and jewelry are probably 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, Michel, and Harlow 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, change in dress is probably as a human universal (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, Brewer and Plumb (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.

Forms of presenting oneself to others and posterity date back to 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 favourite pets. Many famous artists painted self-portraits (Carbon 2017), from Dürer to Picasso, presenting everything in every style from darkly realistic images to idealized versions of themselves (Beyer 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 individuality, from the style of photograph chosen to the manner in which one dressed and presented oneself to the world. The very first portrait pictures date back to 1839; by the 1840s, daguerrotypes 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 describe events as well as depicting teachers and sports teams.

Post-war American culture and 1968. 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. While creating frictions in civic society, and between old and young, some observers argue that it ultimately led to a more tolerant and diverse society.

The counter-culture of the 1960s 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 counter-culture 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 came to diffuse 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 mohair, leather and spike style of punks. Nonetheless, influential innovation in the 1980s and 1990s never reached the levels of the 1960s and 1970s.

2. 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 portraits of students. 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.[3] 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

[3] To avoid this influencing our results, we convert all pictures to black-and-white.

1930 – 2010, it contains a total of over 350,000 yearbooks. Classmates.com already covers thousands of high schools in 44 states from 1930 onwards. This number increases further into the 1980s.

We first run a portrait recognition algorithm to identify where in a yearbook pictures of students are displayed. It 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". We also use information on color and size to identify senior images (which are more likely to be in color and often larger). Human audit samples suggest that we have about 5% false positives, and identify about 70% of all available senior portraits correctly.

We use two main datasets for our analysis. First, we create a balanced panel without sacrificing broad geographic coverage (US Sample). To do so, we rank yearbooks in each city by coverage – the number of yearbooks between 1930 and 2010. Within each city, we also rank high schools by the coverage of the corresponding yearbooks. For each of the top 25 cities acc. to this ranking, we use the three yearbooks with best coverage. The benefit of this sampling approach is that less of the variation over time is the result of new schools coming in or dropping out of the sample. Figure A.2 shows the number of high schools covered by the data as well as the total number of available images. In total, we have some 8 million images in this national dataset.

For the second sample, we draw from the entire universe of yearbooks; for the purposes of this working paper, we only utilize this data for California. Here, we use all images from high school yearbooks – some 4.2 million portraits in total.

To identify individual style elements in our pictures, we first define a vector of 7 style elements, for men and women. We asked a total of 155 MTurk workers to make 18,878 images-category decisions, using a set of 6,000 randomly chosen images. We then trained a convolutional network (ImageNet Inception V3, using Keras under Python), based on transfer learning (Szegedy et al. 2016). We trained the network on 10,613 labelled images-category observations from the MTurk exercise as well as the 800,000 annotated images from the DeepFashion database (Liu et al. 2016), and then evaluated the predictive performance. We then predict style characteristics for 8,265 image-category observations in our national high school image database. Testing accuracy was high, with a range of 0.69-0.93. The network achieved the highest precision when it comes to classifying ties; it did least-well with facial hair. Table A. gives an overview of the training and prediction results.

3. Methods

Humans can typically judge the similarity of images instinctively and quickly (Ginosar et al. 2015). Consider Figure 1. On the first row is a selection of 1964 yearbook images from Natick High School, Massachusetts. All the young men wear jackets, white shirts, and black ties; all have short hair and no moustache. Their facial expressions, while not identical, are broadly similar – somber, confident, serious. Now consider the second row, from the same high school in 1984. Two of the men have long hair, two short hair. One wears suit and tie, two a t-shirt and one an open shit. One sports a moustache, the others not. Looks range from defiant to absent-minded or gazing into a far distance.

While the human eye processes these differences rapidly and half-consciously, a systematic approach to image similarity requires an algorithm that can capture most of these differences quickly and reliably. Analyzing and classifying images is an important challenge in computer science. To compare styles, we need a metric that allows us to measure the style attributes of individual portraits. There are two broad approaches – analysis of features captured by the image, or direct comparisons of image features and attributes like the color or brightness of pixels, individually and collectively, which is then often translated into a quality score. In the second approach, quality measures of two images are then compared through measures such as the structural similarity index measure, SSIM (Wang et al. 2004).

We focus on the analysis of style vectors – combinations of individual, identifiable features. To this end, we first define a style vector of N characteristics, specific to each portrait, A = {a,b,c,...N}, where a, b, c... are style attributes chosen. Figure 2 gives an overview of style attributes in our data. Men with ties is extremely common in the 1930s, and stays above 70% of our sample into the late 1960s. It then falls precipitously, reaching 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 60s, 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 fewer long-term trends. Long hair is en vogue twice in our sample, the 1940s and early 50s, and then again from the 1970s onwards; jewellery is very rare in the 1940s and then becomes common from the 50s onwards.

Styles are not independent of each other. In Figure 3 shows the correlation pattern of style attributes. Suits are positively correlated with ties and bow ties, as one would expect for the case of men, but long hair is remarkably uncorrelated with indicators of formal wear, such as ties and suits. A similarly low correlation is visible for women's jewelry across style attributes.

Next, we combine these style characteristics into unique combinations and examine their evolution over time. In figure 4, we show the composition of styles in our sample for men. In the 1940s, around 70% of men were in the same group – suit, tie, shirt with a collar, no glasses, short hair, no jewelry, and were clean-shaven (style 221). Another tenth had glasses and facial hair, but shared all other attributes (style 222 and 223). Another 10-15% had short hair, no suit, no tie, shirt, no facial hair (style 241). Four styles alone account for 90% of the sample in most years. Style 221 declines in the 1950s already, and men with suits and bow ties become more common. By the late sixties, the influence of the most traditional style (221) falls sharply, from a share of 60% in 1966 to less than 10% by the mid-seventies. Some its use is taken over by pre-existing styles, but the share of "other" (style 999) increases sharply, accounting for more than 30% of the sample by the late seventies. Suit and tie, combined with a moustache, enjoys a brief moment as the dominant style in the 1980s. The more informal shirt without a collar, short hair, no tie style (265) becomes the most common combination by the 1990s. It is striking that "other" (style 999) is the most common type in our sample by the 1970s already, and outranks the proportion of any single style – a sharp change from the 1940s and 1950s, when less than 10% 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 the largest single group.

We are interested in analyzing conformity and persistence – the extent to which individuals in a group "stand out" relative to their peers at any one moment in time, and the degree to which they look like their parents. Complete conformity would imply that everyone picks the same action, for example A=B={2,2,2,2...2}. By contrast, individualism implies a high degree of individual expression: A={2,2,2,2...2} and B={52, 2, 932, 1, ... , 177}.

When persistence is low, there will be little overlap in the styles chosen in the two periods – for example, A_t={99, 88, 77, ... 99} and A_t+1={33, 21, 1, ... 911}. A highly conformist society can also show a great deal of persistence and stability over time: A_t={2,2,8,2...2} and A_t+1={1,2,2,2...2}. We use cosine similarity as the main indicator of style differences

$$\cos(\theta)=\rho_{AB}=\frac{A\cdot B}{||A||||B||} \tag{1}$$

where A and B are vectors of style attributes. Cosine similarity ranges from 0 to 1; two images that share no characteristics score 0, and identical ones, unity. Cosine similarity is useful for our purposes because it is designed to capture *relative* differences between pairs rather than making absolute comparisons.

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 this end, we focus on a) tie (no tie, tie, bow tie), b) glasses, c) facial hair, d) jewelry, e) long hair, f) shirt with or without collar, and g) suit as style attributes. This results in a total of 191 possible combinations (for men) and 192 for women. Human audit samples suggest that we achieve accuracy in the 70-99% range for most characteristics (facial hair=70%; glasses=99%), with a majority of features in the 90+% range (gender, hair, tie, glasses, jewelry).

4. Applications

Measuring cultural change and persistence. We first present broad patterns of change over time and space, using the dataset of cities with the best data coverage during the period 1930-2010. We examine how similar a randomly chosen High School student is to other individuals from his or her class ("individualism"). We also repeat the exercise by comparing each student in year t with people graduating from the same high school 20 years earlier ("persistence"). Figure 7 plots the cosine similarity scores for both types of analysis over time.

We find a high degree of persistence, and low levels of individualism, for most of the 1950s and 1960s. Individualism levels show an inflection point in 1963/4 and begin to rise at an accelerating rate thereafter; from the early to mid-70s, they increase rapidly. Persistence stays relatively high until the late 1960s, before plummeting from 1970 onwards. Since the mid-1970s, individualism is gradually declining, and persistence has been creeping up. While levels of persistence are still far from those seen in the 1950s and 60s, individualism by 2010 had declined to levels close to those seen in the mid-1950s.

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 (Mulhainathan and Spiess 2017). It is especially useful when dealing with high-dimensional datasets with many correlated predictors. Mathematically, it chooses β so as to minimize the loss function given by:

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

where the first term penalizes the lack of accuracy and the second term penalizes the introduction of new covariables. 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) onto the vectors of styles 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 (for example, the contribution of facial hair to persistence in Figure A.5). 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, in most of the cases all styles survive (i.e. 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.

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 6 plots the coefficients of LASSO regressions explaining the contribution of individual style characteristics to the overall individualism score of an image. Take the case of ties, for example. In the 1950s, 60s, and 70s, wearing a tie reduced a young man's individualism score – it made him more similar to another, randomly chosen classmate. By the 1980s, this effect had declined markedly, and then reversed in the 1990s – US students graduating in that decade, and having their portrait taken with a tie, were *more* individualistic, meaning their choice of this particular style attribute made them stand out more. Other factors underwent similarly striking changes over time. Long hair, for example, contributed the most to individualism scores in the 1970s. In contrast, glasses are always positively associated with greater individualism, and their effect is fairly constant across the decades.

In Figure 7, we examine persistence over time more closely. Instead of the (arbitrary) 20-year horizon used in Figure 5, 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 cultural revolution of the early 1970s, similarity scores as far back as 15 years earlier were still positively associated with the style of seniors graduating from high school.

With the sharp rise of counter-culture in the 1970s, this changes profoundly. Similarity scores become zero or negative, indicating a radical decline in persistence. The discontinuity following the hippie movement is so profound and long-lasting that it is not before the mid-1990s that we detect substantial similarity from one generation to the next. In other words, it is only when the first cohorts affected by 1968 become the comparison group for contemporary high schoolers that we find evidence of persistence.

Innovation and influence in style. How can we tell if a style or fashion is genuinely new? And what innovations stick and propagate – and which ones wither on the vine? In our second example of using images as a source for detecting the changes in culture, we apply a simple methodology. It was recently developed for the analysis of US patents and their importance (Kelly et al. 2018). Here, we show how it can be adapted to document and measure how original and influential “innovations in style” are. Kelly et al. (2018) first calculate “backward similarity” as

$$BS_j^{\tau} = \sum_{i \in \mathcal{B}_{j,\tau}} \rho_{j,i} \tag{4}$$

Where ρ is the cosine similarity between the vector of image characteristics for individual j and individual i belonging to the set $\mathcal{B}$ of those who lived up to τ years earlier. In other words, backwards similarity captures how similar the looks of individual j are compared with all other individuals living between the present time t and t-τ.

The second step in the Kelly method is to calculate forward similarity, defined analogously as

$$FS_j^{\tau} = \sum_{i \in \mathcal{F}_{j,\tau}} \rho_{j,i} \tag{5}$$

Where $\mathcal{F}$ is the set of individuals living up to τ years in the future. The measure of image and style innovation of individual (or class) j is then simply

$$\mathrm{q}_j^{\tau} = \frac{FS_j^{\tau}}{BS_j^{\tau}} \tag{6}$$

To fix ideas, consider two cases of style innovation: A young man or woman in 1961 may chose a style that is very different from anything that we can find in yearbooks in the preceding 30 years – BS would be low because the cosine similarity of his/her style with that of other people is low. But if FS – the extent to which people in the future look like him or her – is also low, q would be low. We consider this a case of “failed innovation”, a novel style that failed to catch on. On the other hand, men with long hairs, round sunglasses and scruffy beards in 1961 yearbooks would receive a low BS rating (they are different from the past), but a high forward-looking score because they look like many high school graduates in *their own future* (in this case, the late sixties and early seventies). Their q-score would hence be high – their deviation from the 1961 “norm” would be rated as important using this method because their style was both novel, and because it was picked up in the future. As in the original Kelly et al. paper, we only use a subsample of images to reduce computational complexity.[4] The computational details of this innovation index are described in Appendix II.

In contrast to the Kelly et al. approach, which considers the *entire* future and past for patent innovation, our main measure only examines a five-year window looking backwards and the same span looking forward. The reason is conceptual: While a patent in year t might be influenced by an innovation at t-50, in terms of personal style, this is exceedingly unlikely. We demonstrate the robustness of our findings in the appendix.

Figure 8 shows the Kelly et al. measure over time. It is relatively flat and low in the fifties, and then declines during the early- to mid-sixties. Then, in the late sixties and early seventies, there is a lot of influential innovation – radically different from what came before. Backward similarity falls first, before forward similarity takes off in the mid-seventies. Influential innovation stays high throughout the 1970s, the most long-term influential period of style innovation in our data. While from the late seventies, high-school portraits start to look more like the past once more, forward similarity reaches new heights in the eighties, indicating influential innovation that is picked up in the nineties and noughties.

Next, we apply this method to analyze cultural innovation in California. Since California has long been considered "ground zero" of the Sixties hippies movement, we analyze yearbook pictures from its main metropolitan areas. We first examine where most of the important cultural innovations take place, and then analyze how they diffuse across space.

Using the q-score metric, we can determine which areas of the US were "culturally leading" – innovating and being ahead of the curve – and which areas were cultural followers, imitating the innovations of the leading areas years later. Figure 9 shows a heatmap of important innovation by area. There is substantial variation over both time and space. Panel A shows the centers of important

[4] We use 25% of all available images for California, and 10% for the rest of the country.

innovation in California during the period as a whole. Unsurprisingly, San Francisco and LA stand out as centers of new style. Panel B shows the pattern of innovation in the Bay Area in 1970; Panel C does so for 1980. Here, we can see drastic changes over time. By the late Sixties, San Francisco becomes the area with the highest scores of the California sample – coinciding with the rise of the hippie movement and the influence of the Haight-Ashbury scene (Cottrell 2015). This dominance does not last; for example, by 1980, Oakland and San Jose have become more important centers of innovation.

Where innovation occurs is one dimension of cultural change. Who follows whom, and how innovations diffuse in space, is another. Here, we apply the concept of granger causality in style adoption to map style innovation over time and space in California. In particular, we examine the effect of any one particular style in X on its prevalence in Y by testing whether:

$$\mathbb{P}[Y(t+1)\mathcal{I}(t)] \neq \mathbb{P}[Y(t+1) \mid \mathcal{I}_{-X}(t)] \tag{7}$$

where $\mathbb{P}$ is probability, and $\mathcal{I}$ is the information set at time t. In other words, we predict take-up of a particular style in t+1 from a distributed lag model with N lags in location Y, using the full information set minus X. We then examine if predictions of style takeup in Y can be improved by adding information about take-up in location X up to and including time t.

For example, style 234 – men with a suit and long hair – is first detected in San Francisco in 1963. It then appears in San Diego and LA in 1967, before taking off across California in the early 1970s. The Granger causality test suggests p-values between 0.08 (SF to LA) and 0.00006 (SF to Sacramento); there is no evidence of diffusion in the opposite direction.

Table 1 summarizes the pattern of influence across space. We focus on styles that reach a minimum of 5% adoption in our sample. This gives us a set of 25 style vectors. With six metro areas, the theoretical maximum set of significant patterns of influence is 750. In actual fact, we only find 163 pairs with significant links *in a single* *direction*.[5] LA, for example, influenced the adoption of 13 styles in Fresno in our period. At the same time, there is only one case of influence going the other way. LA and San Francisco stand out for having the highest number of styles influencing other metro areas. We can total up the number of other metros influenced by each area, and how often they themselves were influenced in turn. As the final line in Table 1 makes clear, only two areas are "net exporters" of style – LA and San Francisco. Sacramento, at the opposite end of the spectrum, is the biggest "net importer". While a total of 30 observations is insufficient to perform more detailed analysis, it appears that cultural imports and exports follow the gravity model, with larger metro areas influencing smaller ones, and this influence waning with distance.

"You say you want a revolution." Can our data detect rapid cultural change? In this application, we examine one of the most famous cultural events in recent history – the Beatles' arrival in the United States. Booked for the CBS Ed Sullivan show, the band appeared three times in February 1964, and once more in 1965. Their fame in the US had slowly started to grow in 1963; the show is widely credited with their breakthrough in the States, creating "Beatlemania". The broadcast's audience numbered 73 million viewers, a record at the time, and more than three times higher than for an ordinary Sullivan show. Within a month of their appearance, the Beatles sold 2.5 million records in the US alone; by April, their singles occupied all the top 5 spots in terms of record sales.[6]

The press commented almost as much on the band's hairstyle as on their music. All four musicians sported similar front-combed "mop top" hair. Beatles wigs copying the hairstyle were selling in their thousands per week (Mulligan 2010). In the US, just as in England, "the 'mop top' haircut became the new fashion style for teenage boys" (Loker 2009). In retrospect, it is hard to appreciate how much of the group's

[5] We also have 54 cases of bi-directional Granger causality, where innovations to a style frequency in one place predict innovations in frequency in the other, and vice versa.

[6] We thank Leonardo Bursztyn for inspiring this analysis.

novelty centered on their hair style, not their music. Paul McCarthy observed (Roylance 2000):

> "We came out of nowhere with funny hair, looking like marionettes or something. That was very influential. I think that was really one of the big things that broke us – the hairdo more than the music, originally. A lot of people's fathers had wanted to turn us off. They told their kids, 'Don't be fooled, they're wearing wigs.'
>
> A lot of fathers did turn it off, but a lot of mothers and children made them keep it on. All these kids are now grown-up, and telling us they remember it…. I get people like Dan Akroyd saying, 'Oh man, I remember that Sunday night; we didn't know what had hit us – just sitting there watching Ed Sullivan's show.' Up until then there were jugglers and comedians like Jerry Lewis, and then, suddenly, The Beatles!"

While some believe that the Beatles' success owes more to broad societal trends and chance than to their unique style of music and fashion (Sunstein 2022), we can examine how much of a discontinuity their arrival on the American musical scene made.

We first train a model to identify the fringe/"mop top" hairstyle, using 4,061 hand-coded images for training. The model achieves an overall accuracy of 83%. Looking at the sample of cities with long-coverage yearbooks, we show the evolution of Beatles-inspired hairstyles over time in Figure 10, panel A. We take the last year before the Beatles' arrival on the Ed Sullivan show (1963) as our baseline, and plot coefficients of the likelihood of a mop-top in each year, distinguishing areas by distance to the nearest CBS TV station.[7] There is a clear trend break after the Beatles' appearance on the Sullivan show – prior to 1964, the trend is flat, and two years after their show, aggregate probabilities of Beatles-style hair have risen by 10 pp.

Interestingly, we find suggestive evidence that the CBS appearance was behind this aggregate trend break. The only areas showing a significant increase in the year

[7] In the quantitative analysis below, we use data from an ITM model in the spirit of Olken (2009). Here, we use distance to illustrate the basic patterns in our data.

of the show are those where CBS was broadcasting. Panel B explores this link further, plotting the probability of a mop top against distance to the nearest CBS station. For 1963, there is not link between fringe-style hairdos and distance to CBS; for 1964, this changes dramatically, with much higher rates near the CBS stations. In panel C, we show that in 1963, overall distributions of Beatles-style hair are identical in the CBS and non-CBS areas; in 1964, this is no longer the case, with low probabilities declining sharply and medium-to-high probabilities (40%+) rising sharply.

In Table 2, we examine the statistical significance of our finding. We use our data in a simple difference-in-difference framework, comparing areas exposed to CBS-programming to those that are not exposed. We do so at two time horizons – short-term (Panel A) and long-term (Panel b) – and for two measures of our outcome – the average probability of Beatles-style hair, and the likelihood of an individual image having at least a 75% probability showing a Beatles-style hairdo.

When we compare two years before treatment, 1962 and 1963, with the one when the Sullivan show hit the airwaves (1964) we find a highly significant jump in the share of images classified as having “mop-top” hair. The effect is large. Going from the 25th to the 75th percentile of the CBS signal strength distribution increased mop-top uptake by almost 2 percentage points, or 10% of the baseline rate. The share of images that are Beatles-style with more than 75% probability increases by almost 60% of the pre-Sullivan show level. These results are robust to controlling for state fixed effects, year fixed effects, and a third-order polynomial in the distance to the nearest CBS transmitter. Even when we control for the strength of other TV stations do our results hold. The same pattern is visible for the long-term sample, using a five-year window around 1963. As take-up of the Beatles hair-style increases in the country as a whole, the sample mean post-treatment increases. Nonetheless, we find a 2-2.5 pp increase in the likelihood of adoption across specifications.

Figure 9 presents a heatplot of the share of Beatles-style hairdos by state over time. While there is an occasional efflorescence of a similar style in some states in earlier years, it never becomes dominant or even an important minority "taste". The discontinuity after 1964 is visible across states. Some places see a more rapid rise in Beatles-style hairstyle – California, for example, has relatively rapid take-up in the three years after the show screened. Other states like Delaware and South Dakota show limited adoption initially, but eventually converge to high or very high levels of Beatles-style hair.

5. Conclusions

The Cambridge Dictionary defines 'culture' as "a. the way of life, especially the general customs and beliefs, of a particular group of people at a particular time, b. the attitudes, behaviour, opinions, etc. of a particular group of people within society…". In recent years, interest in cultural economics has increased sharply (Akerlof and Kranton 2000). Empirical studies of culture have focused almost exclusively on attitudes and beliefs, articulated either in surveys and written text, or reflected in actions. Customs and beliefs, and the extent to which they are specific to a particular group, however, comprise a much wider range of activities. In particular, visual culture and self-representation through style and fashion choices are arguably important but largely neglected aspects of cultural economics.

In this paper, we make two contributions. First, we introduce a set of methods and tools that allow rigorous and precise 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. The vector representation in turn facilitates the use of standard measures of similarity. These can be used to map into two key dimensions of culture, homogeneity ("conformity") and persistence. We also show how methods from the measurement of innovation – previously applied to patenting – can be applied to examine influential style changes.

Second, we apply our new methods to a large dataset of US high school senior images. We trace the decline and fall of image 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. Using the Kelly et al. (2021) method for identifying influential innovation, we show how some areas of the US contributed markedly more to important new trends than others, and demonstrate that periods of peak creativity and influence are relatively brief.

The combination of new methods and new data also allows us to identify "insertion points" of new cultural trends, leveraging the granularity and frequency of our image data. We analyze changes in men's hairstyles in the US during the Sixties, following the Beatles' hallmark appearance on the "Ed Sullivan" show. A highly distinctive hair style, the "mop top", attracted almost as much attention as their music. While this style was rare and not growing as a proportion of senior portraits before 1964, it quickly gained popularity afterwards. In the year of the broadcast, the effect is largest in areas where CBS, the TV network that broadcast the Sullivan show, had affiliated stations – and the further from a CBS station an area was, the lower the likelihood of a jump in mop top hairstyles among men in 1964. This in turn lead to wider adoption in the (youth) population at large, even outside areas with CBS coverage – resulting in a spectacular upward trajectory of fringe-style hairstyles. Within three years of the broadcast, an additional 15% of American teenagers sported Beatles-style hair.

The German philosopher Georg Simmel (1957) famously defined fashion as "a form of imitation…, paradoxically, in changing it differentiates one time from another." How such imitation can lead to fashions that define a period is challenging to document empirically. The spread of "Beatlemania" in the US after 1963 as reflected in their hair style allows us to demonstrate where innovations come from, and how imitation drives changes in style. The pattern we document is in line with interpretations emphasizing informational cascades behind rapid cultural, social, and political changes (Kuran 1989; Sunstein 2022).

References

Akerlof, George A., and Rachel E. Kranton. 2000. 'Economics and Identity'. *The Quarterly Journal of Economics* 115 (3): 715–53.

Alesina, Alberto, Paola Giuliano, and Nathan Nunn. 2011. 'Fertility and the Plough'. *American Economic Review* 101 (3): 499–503.

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'. National Bureau of Economic Research.

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

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): 941–64.

Becker, Sascha O., Katrin Boeckh, Christa Hainz, and Ludger Woessmann. 2016. 'The Empire Is Dead, Long Live the Empire! Long-run Persistence of Trust and Corruption in the Bureaucracy'. *The Economic Journal* 126 (590): 40–74.

Beyer, Andreas. 2003. *Portraits: A History*. Harry N. Abrams.

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

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

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

Chamley, Christophe. 2004. *Rational Herds: Economic Models of Social Learning*. Cambridge University Press.

Cottrell, Robert C. 2015. *Sex, Drugs, and Rock'n'Roll: The Rise of America's 1960s Counterculture*. Rowman & Littlefield.

Crafts, Nicholas F. 1985. 'British Economic Growth during the Industrial Revolution.' *Oxford Univ. Press.*

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

Fernández, Raquel, Sahar Parsa, and Martina Viarengo. 2019. 'Coming out in America: AIDS, Politics, and Cultural Change'. National Bureau of Economic Research.

Galton, FJ. 1878. 'Composite Portraits'. *Nature* 18: 97–100.

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

Giuliano, Paola, and Nathan Nunn. 2021. 'Understanding Cultural Persistence and Change'. *The Review of Economic Studies* 88 (4): 1541–81.

Guiso, Luigi, Paola Sapienza, and Luigi Zingales. 2016. 'Long-Term Persistence'. *Journal of the European Economic Association* 14 (6): 1401–36.

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

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.

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

Karni, Edi, and David Schmeidler. 1990. 'Fixed Preferences and Changing Tastes'. *The American Economic Review* 80 (2): 262–67.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy. 2021. 'Measuring Technological Innovation over the Long Run'. *American Economic Review: Insights* 3 (3): 303–20.

Kiapour, M Hadi, Kota Yamaguchi, Alexander C Berg, and Tamara L Berg. 2014. 'Hipster Wars: Discovering Elements of Fashion Styles'. In *European Conference on Computer Vision*, 472–88. Springer.

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

Kuran, T. 1989. 'Sparks and Prairie Fires: A Theory of Unanticipated Political Revolution'. *Public Choice* 61 (1): 41–74.

Lee, Stefan, Nicolas Maisonneuve, David Crandall, Alexei A Efros, and Josef Sivic. 2015. 'Linking Past to Present: Discovering Style in Two Centuries of Architecture'. In .

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

Liu, Ziwei, Ping Luo, Shi Qiu, Xiaogang Wang, and Xiaoou Tang. 2016. 'Large-Scale Fashion (DeepFashion) Database'. *Xiaoou TangMultimedia Laboratory, The Chinese University of Hong Kong, Category and Attribute Prediction Benchmark. Https://Url. Kr/DfQWlV*.

Loker, Bradford E. 2009. *History with the Beatles*. Dog Ear Publishing.

Ludwig, Jens, and Sendhil Mullainathan. 2023. 'Machine Learning as a Tool for Hypothesis Generation'. National Bureau of Economic Research.

Matsuyama, Kiminori. 1991. 'Custom versus Fashion: Path-Dependence and Limit Cycles in a Random Matching Game'. Discussion paper.

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 Melanie Meng Xue. 2021. 'Folklore'. *The Quarterly Journal of Economics* 136 (4): 1993–2046.

Mulligan, Kate Siobhan. 2010. *The Beatles: A Musical Biography*. ABC-CLIO.

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

Olken, Benjamin A. 2009. 'Do Television and Radio Destroy Social Capital? Evidence from Indonesian Villages'. *American Economic Journal: Applied Economics* 1 (4): 1–33.

Pesendorfer, Wolfgang. 1995. 'Design Innovation and Fashion Cycles'. *The American Economic Review*, 771–92.

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

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

Roylance, Brian. 2000. *The Beatles Anthology*. Chronicle Books.

Simmel, Georg. 1957. 'Fashion'. *American Journal of Sociology* 62 (6): 541–58.

Sunde, Uwe, Thomas Dohmen, Benjamin Enke, Armin Falk, David Huffman, and Gerrit Meyerheim. 2022. 'Patience and Comparative Development'. *The Review of Economic Studies* 89 (5): 2806–40.

Sunstein, Cass R. 2022. 'Beatlemania'. *Journal of Beatles Studies*.

Szegedy, Christian, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. 'Rethinking the Inception Architecture for Computer Vision'. In , 2818–26.

Veblen, Thorstein. 1973. *The Theory of the Leisure Class*. Boston, MA: Houghton Mifflin.

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 from High School Yearbooks

Note: First row is from 1964, second row is from 1984.

Figure 2: Style Attributes in High School Senior Yearbook Pictures.

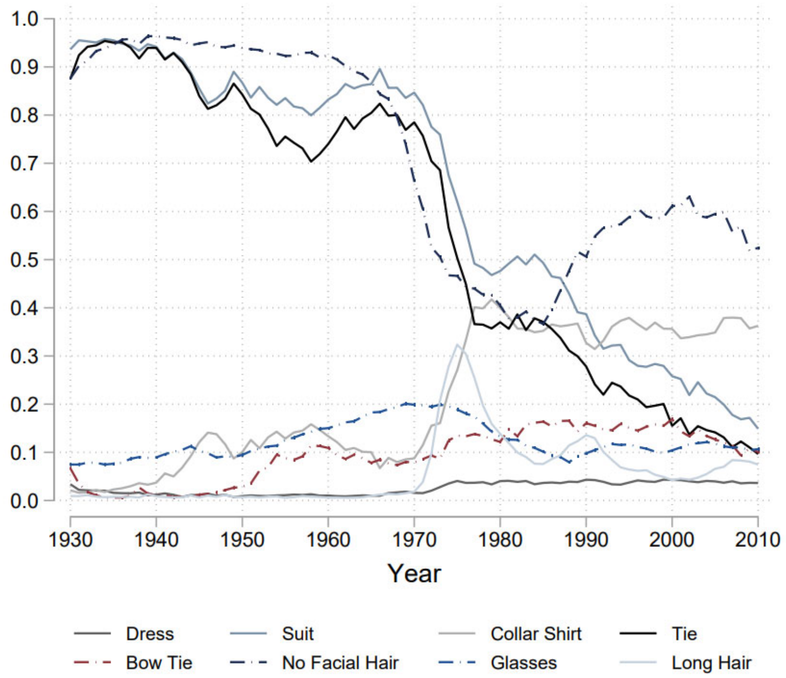

Panel A: Men

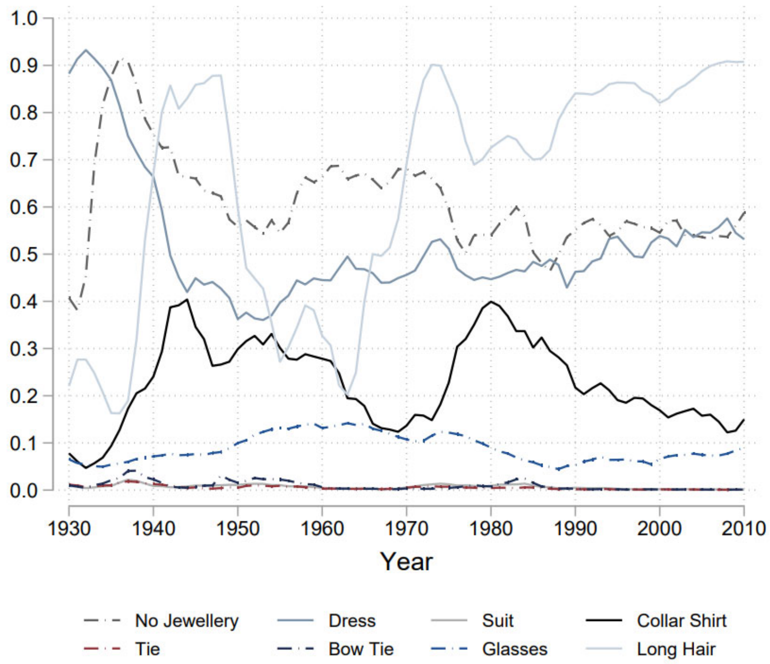

Panel B: Women

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

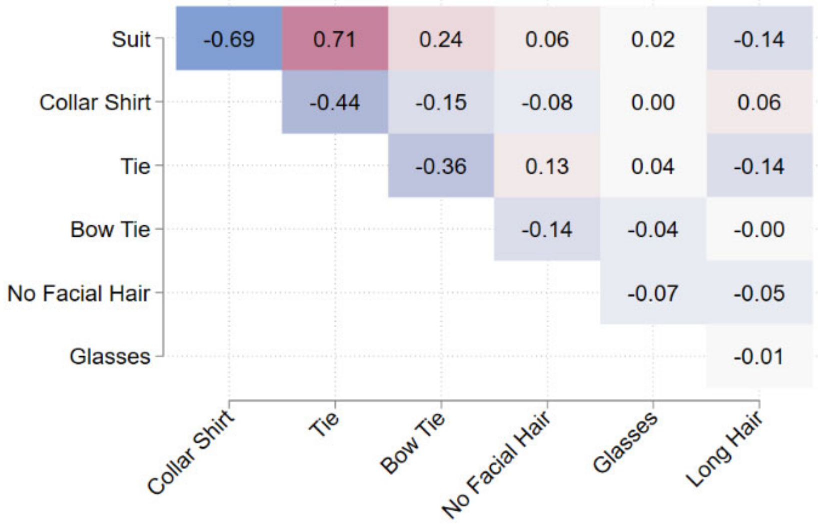

Panel A: Men

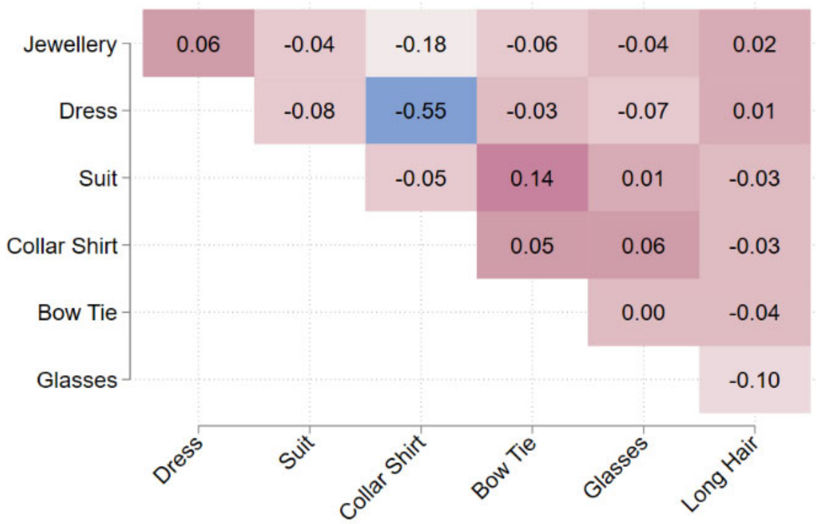

Panel B: Women

Figure 4: Styles in High School Senior Yearbook Pictures.

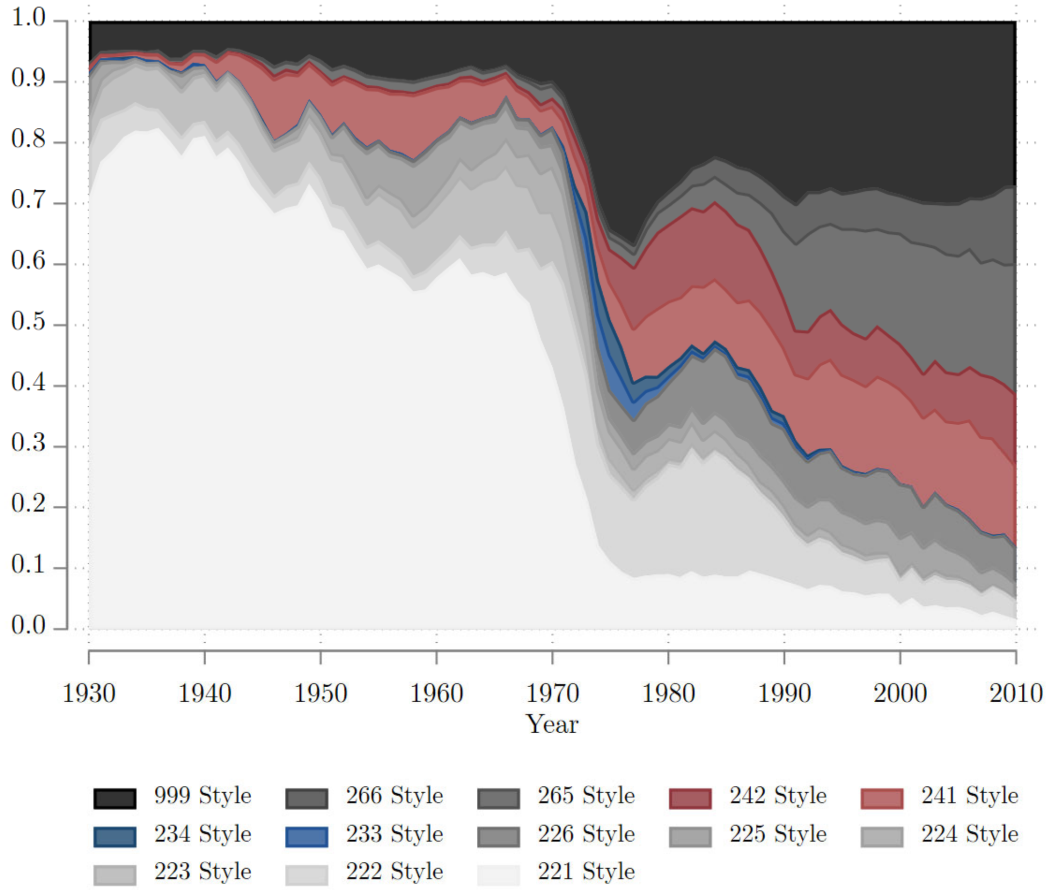

Note: This graph depicts the share over time for each discrete combination of style markers. Each image is assigned a style by collapsing the probability of each style marker. Styles that comprise less than 5% of the sample for all years are assigned Style 999 – Other. A description of styles follows:
Style 221: suit/short hair/normal tie/no glasses/no facial hair/no jewelry
Style 222: suit/short hair/normal tie/no glasses/yes facial hair/no jewelry
Style 223: suit/short hair/normal tie/yes glasses/no facial hair/no jewelry
Style 224: suit/short hair/normal tie/yes glasses/yes facial hair/no jewelry
Style 225: suit/short hair/bow tie/no glasses/no facial hair/no jewelry
Style 226: suit/short hair/bow tie/no glasses/yes facial hair/no jewelry
Style 233: suit/long hair/normal tie/no glasses/no facial hair/no jewelry
Style 234: suit/long hair/normal tie/no glasses/yes facial hair/no jewelry
Style 241: shirt with collar/short hair/no tie/no glasses/no facial hair/no jewelry
Style 242: shirt with collar/short hair/no tie/no glasses/yes facial hair/no jewelry
Style 265: shirt without collar/short hair/no tie/no glasses/no facial hair/no jewelry
Style 266: shirt without collar/short hair/no tie/no glasses/yes facial hair/no jewelry
Style 999: Others

Figure 5: Individualism and Persistence in US High School Yearbook Images, 1950-2010.

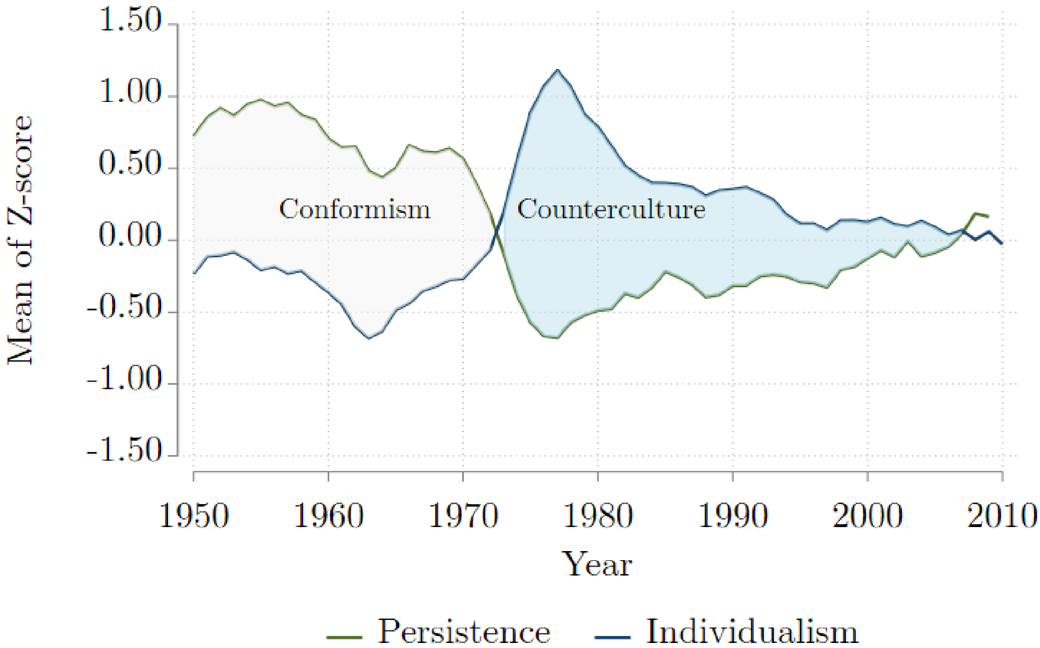

Note: All individuals across US, each compared with all individuals from same high school from their own year (individualism) and 20 years before (persistence). Cosine similarity is inverted for individualism, both are Z-scored.

Figure 6: Drivers of Individualism in US High School Yearbook Images, 1950-2010.

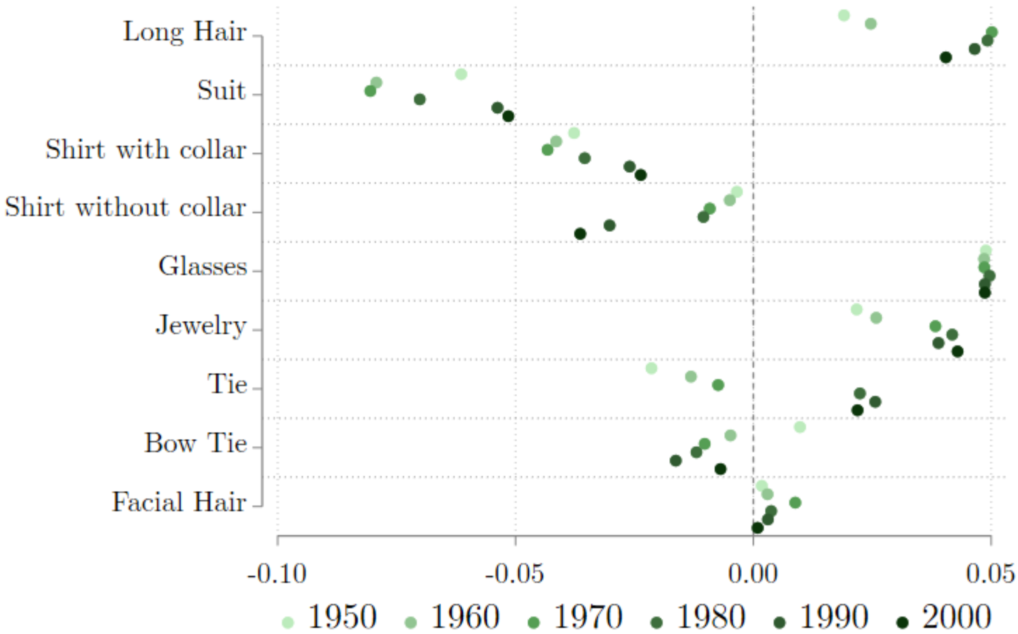

Note: Dots represent coefficient estimates from a Lasso regression of individualism score on style characteristics. Each dot represents a single coefficient for a style, derived from bi-variate regressions. Values greater zero indicate a positive contribution to students' individualism score.

Figure 7: Persistence over Time.

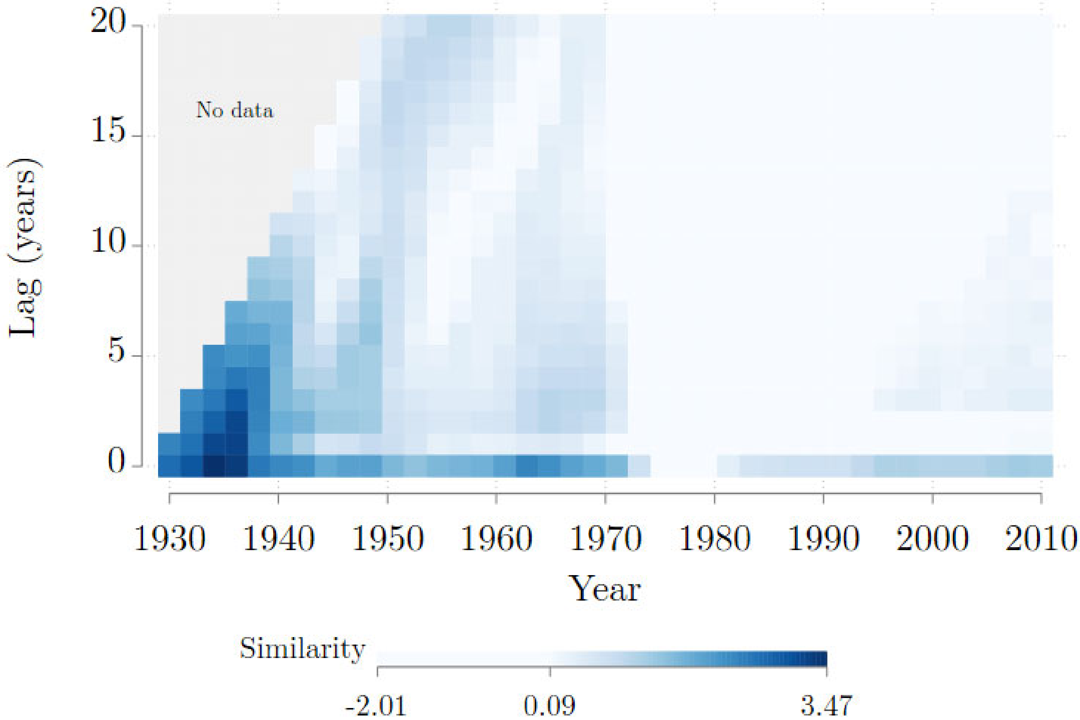

Note: Individuals, drawn against 100 random people who attended their high school up to twenty years before them. Similarity scores normalised and are censored below the median z-score.

Figure 8: Influential Innovation in Style over Time, Decomposed into Forward and Backward Similarity

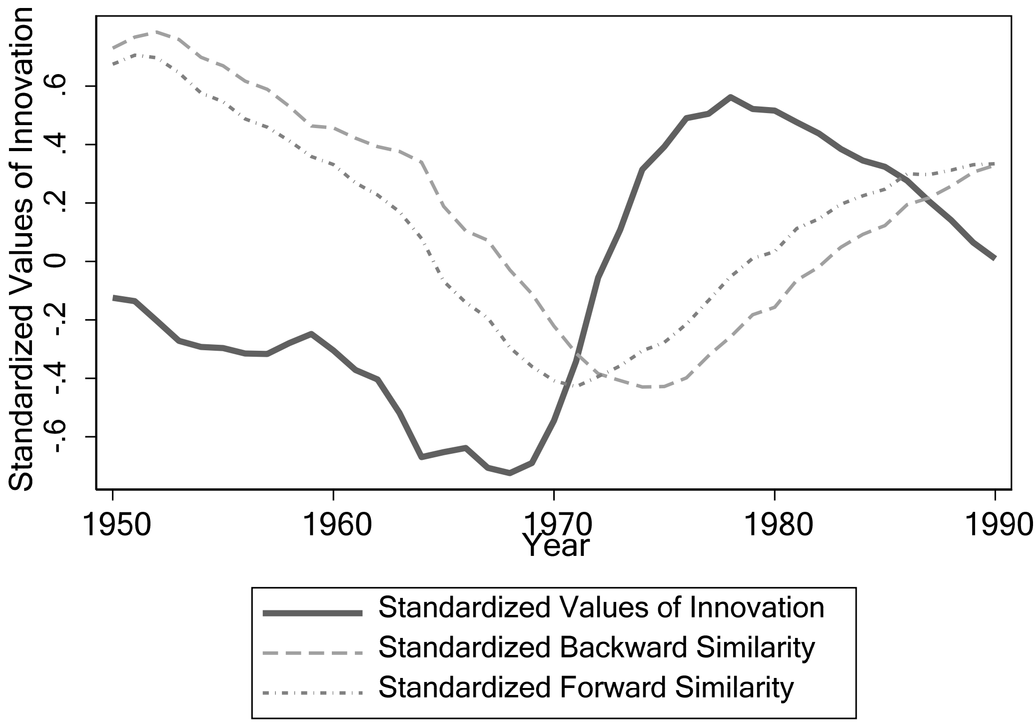

Note: For 10% of the national sample, we calculate the forwards and backwards similarities using the five years ahead and behind of each year for each image. We z-score the results for these three values and collapse to year means.

Figure 9: Kelly et al. Measure of Innovation in California, 1960-1990

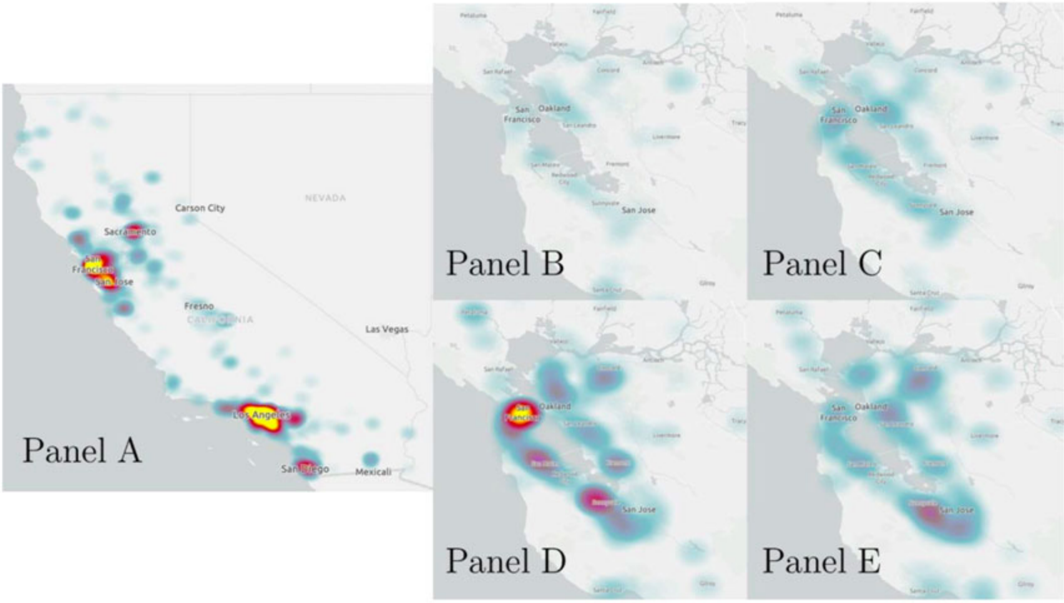

Note: For 25% of the images in California, we calculate the innovation index using a five year window forwards and backwards. We calculate city averages of innovation for each year from 1960 to 1980. Panel A shows a heat map of innovation according to the Kelly et al. measure for the period as a whole; Panel B is for 1950, C for 1960, D for 1970, and E for 1980.

Figure 10: Beatles-style "Mop-Top" Hair in High School Senior Portraits

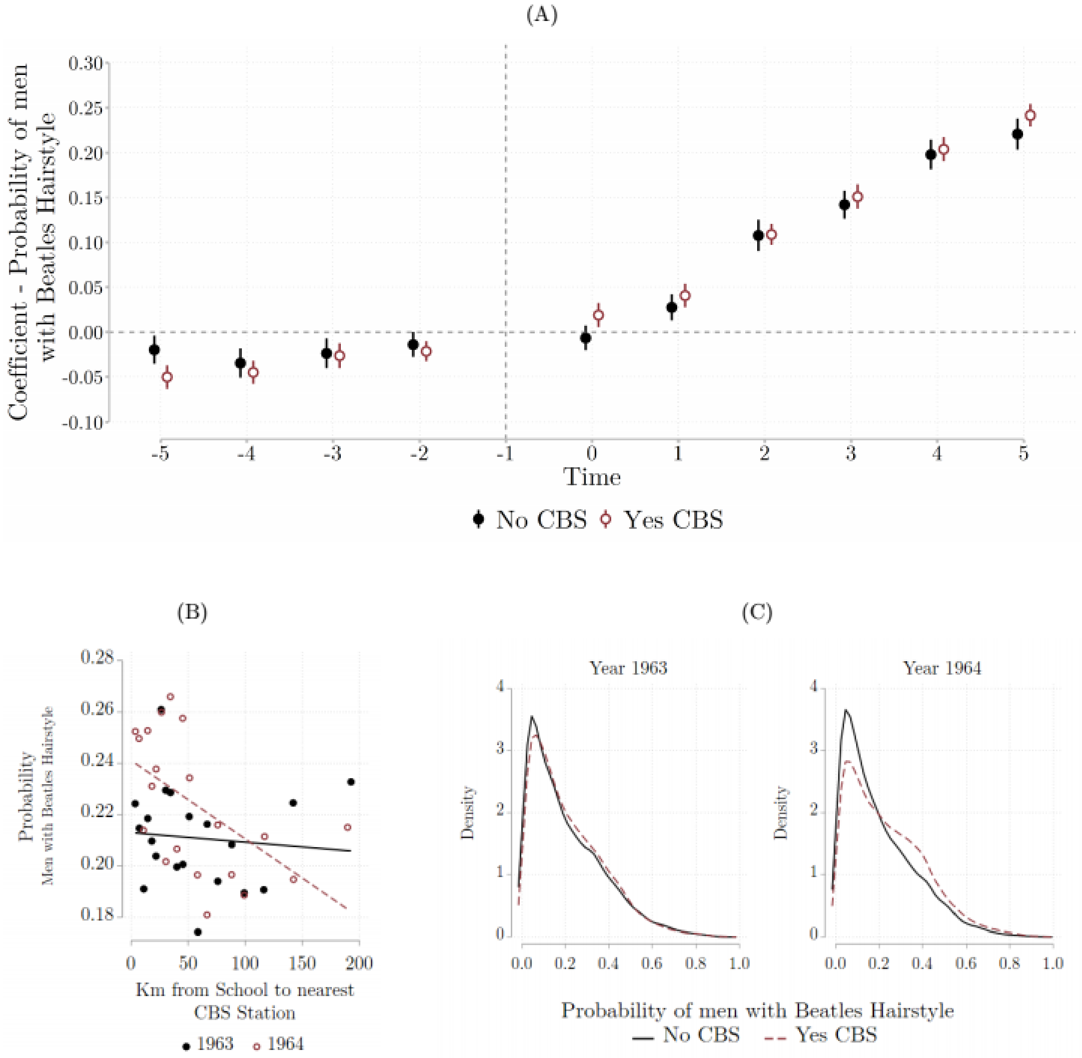

Note: Panel A. Probability of Beatles-style "Mop-Top" Hair in High School Senior Portraits (Event study - Average probability over time CBS vs non-CBS, standard errors are clustered at school level and includes school FE.); Panel B. Binscatter of probability of Beatles-style hair, by distance to nearest CBS station, 1963 vs 1964, in km; Panel C. Kernel density function, probability of Beatles-style hair, CBS vs non-CBS, 1963 and 1964.
CBS = 1 if CBS Station within 60 km from school.

Figure 11: Probability of Beatles-style "Mop-Top" Hair in High School Senior Portraits, Heatmap by State

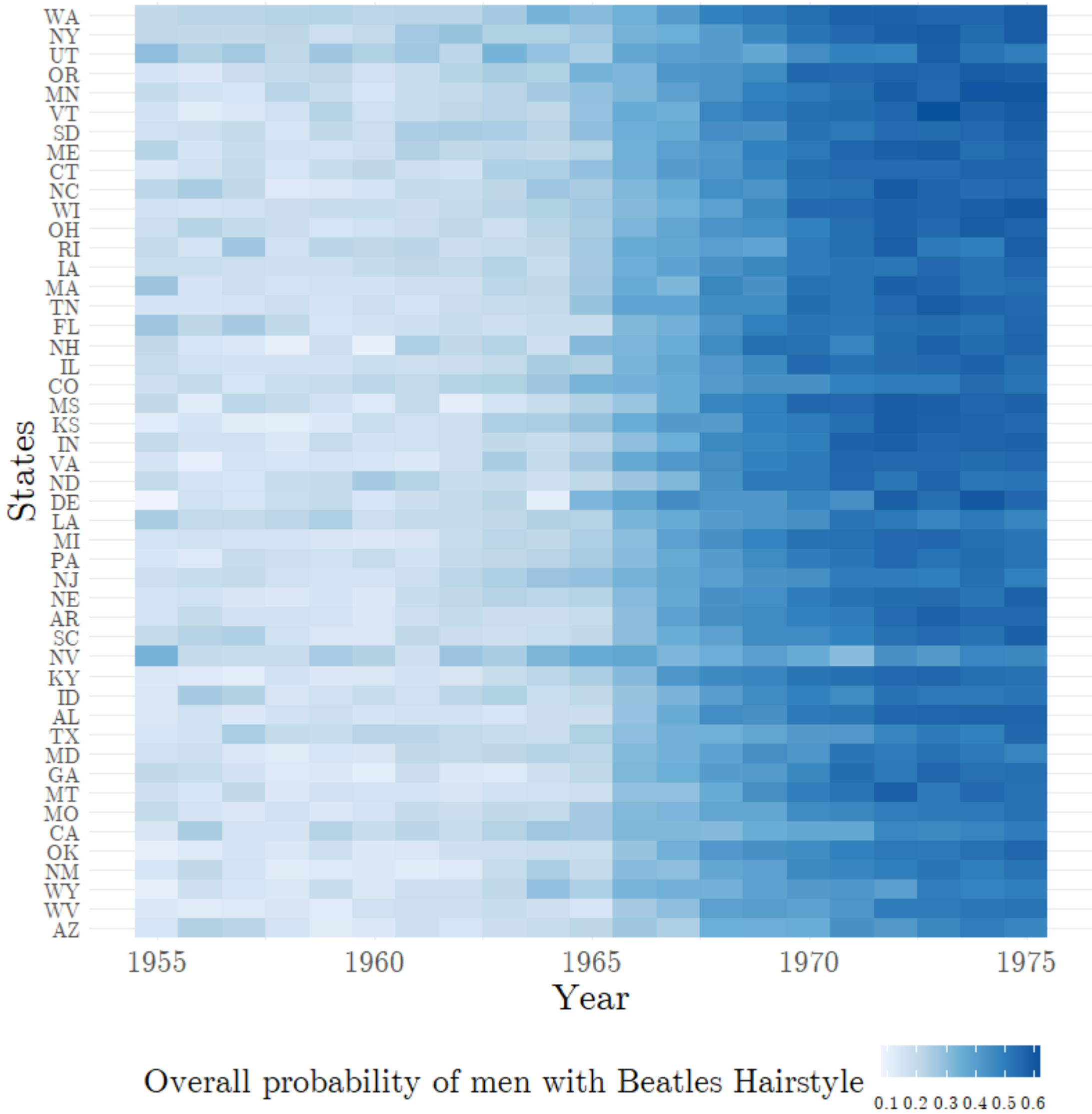

Note: Average probability of Beatles Hairstyle by state and year

Tables

Table 1: Exports and Imports of Style: Granger Causality Matrix - California

from	to							
	Bakersfield	*Fresno*	*LA*	*Sacramento*	*San Diego*	*San Francisco*	***Grand Total***	***Net Diff***
Bakersfield		5	5	6	7	5	28	**-6**
Fresno	6		1	5	6	4	22	**-5**
LA	5	13		7	6	4	35	**13**
Sacramento	8	2	2		4	3	19	**-11**
San Diego	6	2	5	5		4	22	**-8**
San Francisco	9	5	9	7	7		37	**17**
Grand Total	34	27	22	30	30	20	163	

Each entry gives the number of significant Granger effects from metro area X (listed top to bottom) to metro area Y (listed across), where significance is at least at the **5**% level. The net difference indicates the difference between the number of styles of a city that it influenced other cities on minus the number of styles a city received influence by others.

Table 2: Results Beatles Regression

TABLE 2

	Prob(Beatles Hairstyle)			Prob(Beatles Hairstyle) > 75th		
	(1)	(2)	(3)	(4)	(5)	(6)
PANEL A: SHORT-RUN						
CBS	0.002	0.004	0.006	-0.002	0.001	0.001
	(0.01)	(0.02)	(0.02)	(0.00)	(0.00)	(0.00)
Post	-0.113	0.000	0.000	-0.281**	0.000	0.000
	(0.86)	(.)	(.)	(0.13)	(.)	(.)
Post x CBS	0.044**	0.042**	0.043**	0.008***	0.008***	0.008**
	(0.02)	(0.02)	(0.02)	(0.00)	(0.00)	(0.00)
Non-CBS			-0.023			-0.004
			(0.02)			(0.00)
Post x Non-CBS			-0.001			-0.000
			(0.02)			(0.00)
R^2	0.006	0.042	0.043	0.000	0.004	0.004
Third-order Polyn. Distance	No	Yes	Yes	No	Yes	Yes
State-Year FEs	No	Yes	Yes	No	Yes	Yes
Observations	192,669	192,669	192,669	192,669	192,669	192,669
PANEL B: LONG-RUN						
CBS	-0.0006	-0.0061	-0.0007	-0.0024*	-0.0044	-0.0022
	(0.01)	(0.01)	(0.01)	(0.00)	(0.00)	(0.00)
Post	11.5024***	0.0000	0.0000	2.8819***	0.0000	0.0000
	(0.73)	(.)	(.)	(0.36)	(.)	(.)
Post x CBS	0.0312**	0.0341***	0.0280**	0.0136**	0.0146**	0.0105
	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)	(0.01)
Non-CBS			-0.0245			-0.0044
			(0.02)			(0.00)
Post x Non-CBS			0.0129			0.0090
			(0.01)			(0.01)
R^2	0.103	0.183	0.183	0.014	0.038	0.038
Third-order Polyn. Distance	No	Yes	Yes	No	Yes	Yes
State-Year FEs	No	Yes	Yes	No	Yes	Yes
Observations	735,408	735,408	735,408	735,408	735,408	735,408

The table presents coefficients capturing changes in the probability of male high school seniors sporting Beatles-style hair. CBS is the signal strength of the nearest CBS station (or CBS affiliate), derived from an ITM model based on transmitter location and terrain features. Post is set to 1 for the period 1964 and following years. Specifications # also control for a third-order polynomial in distance to the nearest CBS station, so that only terrain features influence signal strength.

Appendix

Appendix I: Database Construction

a. *Download yearbook pages from classmates.com*

The data from classmates.com is acquired in two phases. 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. With 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.

b. *Crop all 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 that lie within a brightness threshold to white or black, to then draw rectangles around black areas. 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 10-50% of the page. We crop the drawn rectangles in a page if three criteria are met:

- the number of faces (recognized by the cv2 model) vs the number of images in a page is close;
- at least 50% images only portray one face, and the face covers more than 30% of the image;
- there are at least two images on the page.

The portrait recognition algorithm can still be improved in these critical cases:

- The background color is similar to the image color

- The portraits have a darker connecting border compared to the background of the page

- There is text overlaid and connecting different portraits

c. Select senior portraits

First, we block out the pages in the yearbooks into series of images by identifying consecutive runs of 4 pages that share similar image color and size. 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 seniors. We use optical character recognition and information about 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", for example.
- 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 listed on classmates.

With the crop and select pipeline, we estimate that we are able to select up to 70% of all senior portraits. The number of false positives is within 5%.

d. Building the classifiers

Images are manually labeled to create a training set. For each style marker we want to analyze, we label around 5,000 images. We encode each image with the following markers:

- Gender: female/male
- Hair: long/short
- Tie: no tie/normal tie/bow tie
- Clothing: dress/suit/shirt without collar/shirt with collar
- Glasses: yes/no
- Jewelry: yes/no
- Facial hair: yes/no

For each image, the classifier returns the probability of each tag occurring for every option above – each image has a probability p of "glasses" and a probability 1-p "no glasses". Up to this point we have a collection of individual portraits, along with information on the year and the high school it belongs to. Next step is to convert each image into information that can be analyzed. In order to achieve this we train a classifier for each style, using the manually labelled images as our training sample. We pre-process the images before applying any machine learning algorithm, transforming the images from RGB to greyscale and resizing them. We train 7 classifiers (one for each style dimension), 5 of them performing a binary classification, and 2 (tie and clothing) a multi-class classification. The human-audited accuracies are estimated as follows:

- Gender: 93%
- Hair: 91%
- Tie: 97%
- Clothing: 76%
- Glasses: 99%

- Jewelry: 91%
- Facial hair: 70%

e. Geocoding the High Schools

To geocode the high schools in the nationally representative sample, we use OpenStreetMaps API, and fill in missing information with manual searches. The California sample utilizes a slightly different geocoding approach than the national sample, utilizing the availability of high-quality data from the California State Government's Geodata Portal. This dataset is advantageous in that it contains information on the locations of schools not currently open or those that have been merged with other high schools. 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 geocoded by hand.

Appendix II: Details of Kelly et al. Analysis

We map the process of Kelly et. Al (2018) to our dataset by describing each image as a patent and its style markers as its textual content. Each cohort is described by an image id, year, and a vector of style markers. Kelly et. Al constructs their vector using "term frequency backwards inverse document frequency" (TFBIDF). Since our data arrives from the classifier as a fully formed table, we use the vector of styles as outputted by the classifier model as our vector for comparison, eliminating the need for computationally intensive normalization procedures as used by Kell et. al.

Because the algorithm is extremely computationally intensive, we use two stages of sampling to make the problem feasible. First, we draw a small sample of comparisons -0.5% of the data- for the calculations of forward and backward similarity. This is done across the entire range of the dataset, ensuring buffers on either end of the study period. For each run of the algorithm, we use the same comparison sample across all partitions of the data. Second, we sample from within our year range on which the innovation algorithm is computed. We draw from 10% of our nationally representative sample and 25% of our California sample. As in Kelly et. Al, we set a threshold - in our case .25 - to be classified as a binary to create a sparser matrix. The California analysis evaluates a quarter of the available data and uses 0.5% of the data as a comparison sample. This corresponds to approximate 930,000 images analyzed against 18,500 images each. This equates to approximately 1.7 billion pairwise similarities. The data is then collapsed to school-year cohorts, taking the average of the innovation scores in the school. We select a window of 5 years, as this is a reasonable window to expect fashion innovations to catch on. As Figure A.6 shows, trends generally hold across comparison windows, though larger windows produce more muted peaks and troughs.

Appendix III: Details of Beatles Analysis

First we run a classifier to measure the probability of a man with the Beatles hairstyle. For that we trained a Random Forest algorithm with 4,061 images and validated it over 947 images. The accuracy of the model is 83%. We classified all men in the national sample between 1955-1975 (1,449,490 obs).

To construct the variable of CBS/No CBS coverage we use the Broadcast Yearbook of 1964. We consider each TV Station that had CBS net affiliation. Then, we geocode each tower using the Homeland Infrastructure Foundation-Level Database from the US Department of Homeland Security. For those stations we don’t have information about the exact position, we input the latitude and longitude of the county’s centroid as an approximation. After that, we calculate the distance from each high school to the nearest station.

Appendix IV: Figures

Figure A.1: Example of Yearbook Images: Tift High School, Tift, Georgia, 1959

Figure A.2: Descriptive graphs – US Sample

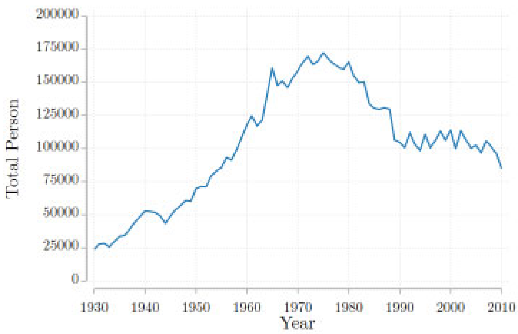

(a) People

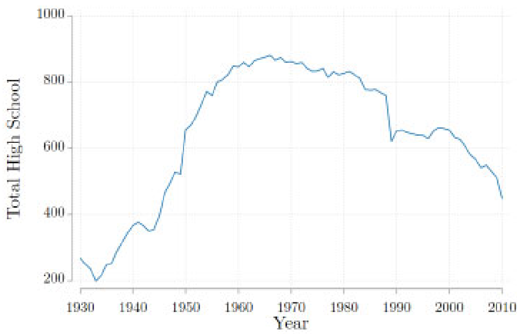

(b) High Schools

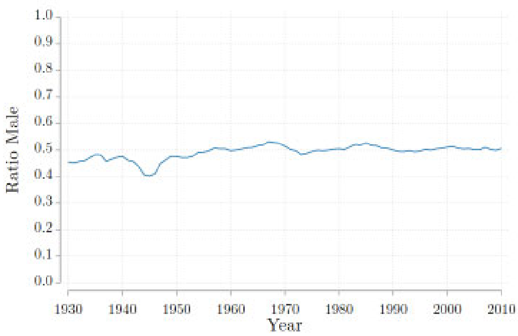

(c) Ratio Male

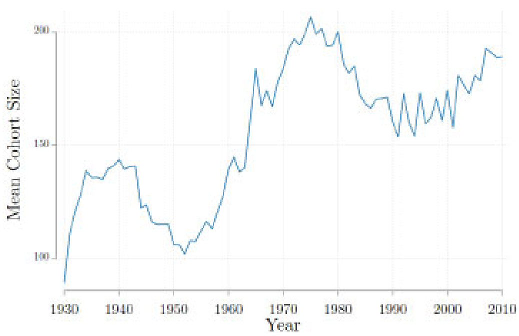

(d) Mean cohort size

Note: Panel A. Number of students by year; Panel B. High Schools by year; Panel C. Share of males by year; Panel C. Mean cohort size by year.

Figure A.3: Training and Prediction of Image Features, High School Dataset

Attribute	Training	Testing	Testing accuracy
Glasses	2500 images per category (i.e. no/yes glasses)	1250 images per category (i.e. no/yes glasses)	87%
Ties	1338 images per category (i.e. tie/bow tie/no tie)	540 images per category (i.e. tie/bow tie/no tie)	93%
Jewelry	2500 images per category (i.e. no/yes jewelry)	2500 images per category (i.e. no/yes jewelry)	80%
Hair	2500 images per category (i.e. short/long hair)	2500 images per category (i.e. short/long hair)	87%
Clothes	1250 images per category (i.e. dress/suit/shirt with collar/shirt without collar)	1250 images per category (i.e. dress/suit/shirt with collar/shirt without collar)	73%
Facial hair	525 images per category (i.e. no/yes facial hair)	225 images per category (i.e. no/yes facial hair)	69%

Figure A.4: Share of US Sample Images Relative to Population Aged 15-19

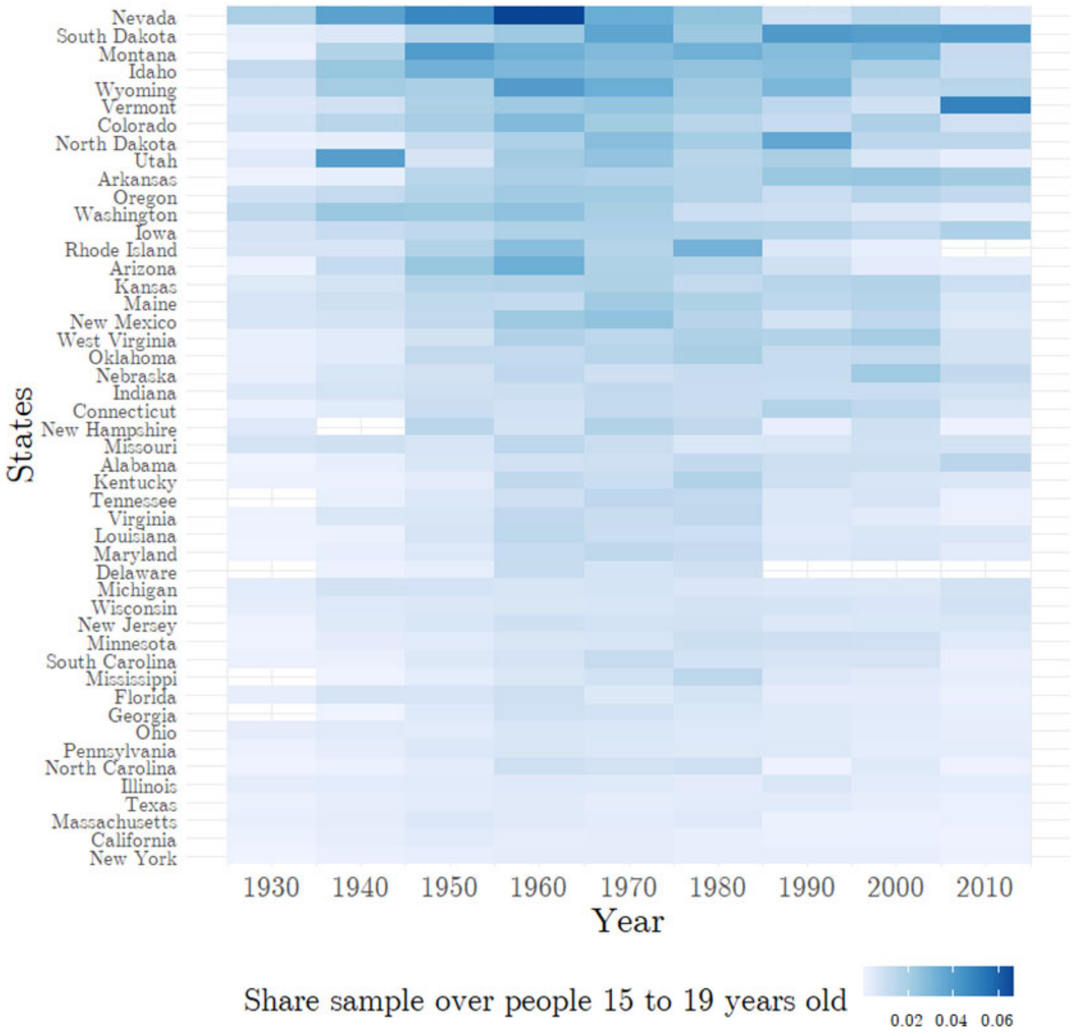

Note: Population by Age and State level from NHGIS Census Records.

Figure A.5: Drivers of Persistence in US High School Yearbook Images, 1950-2010.

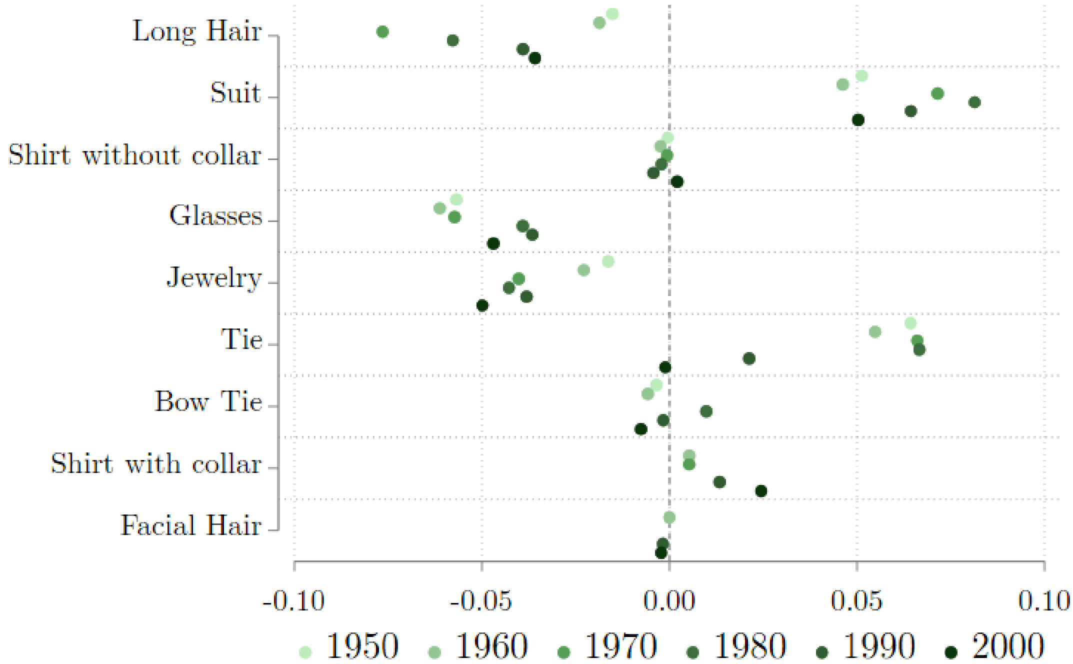

Note: Dots represent coefficient estimates from a Lasso regression of persistence score on style characteristics. Each dot represents a single coefficient for a style, derived from bi-variate regressions. Values greater zero indicate a positive contribution to students' persistence score.

Figure A.6: Kelly et al. Measures at Different Time Horizons

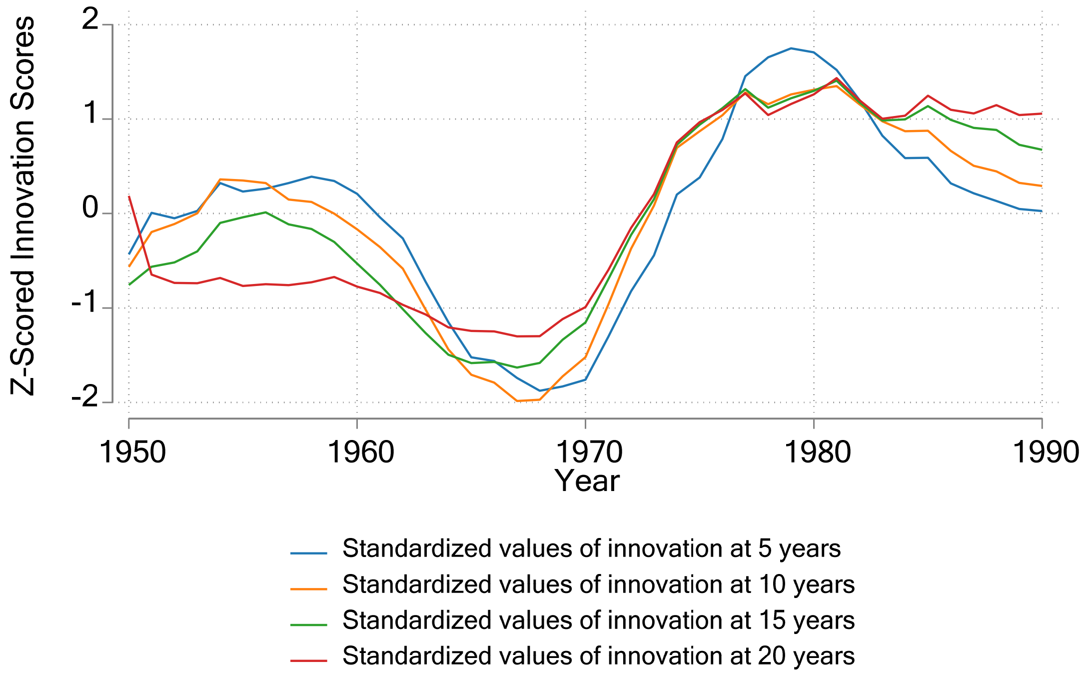

Note: We calculate our innovation metric using forwards and backwards similarities of different sizes on 1% of our nationally representative sample (83,000 images). The results are z-scored and collapsed to year means.

Figure A.7: Kelly et al. Measure of Innovation by California City, 1950-1980

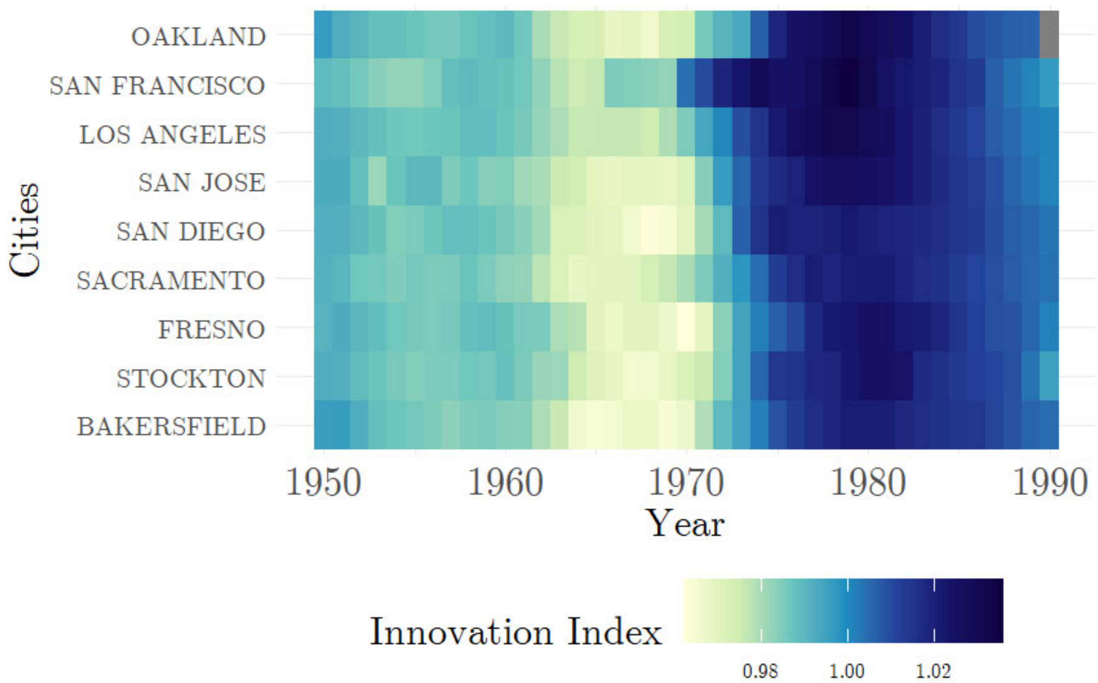

Note: For 25% of the images in California, we calculate the innovation index using a five year window forwards and backwards. We calculate city averages of innovation for each year from 1950 to 1990.

Figure A.8: The Beatles, 1964

(A)

(B)

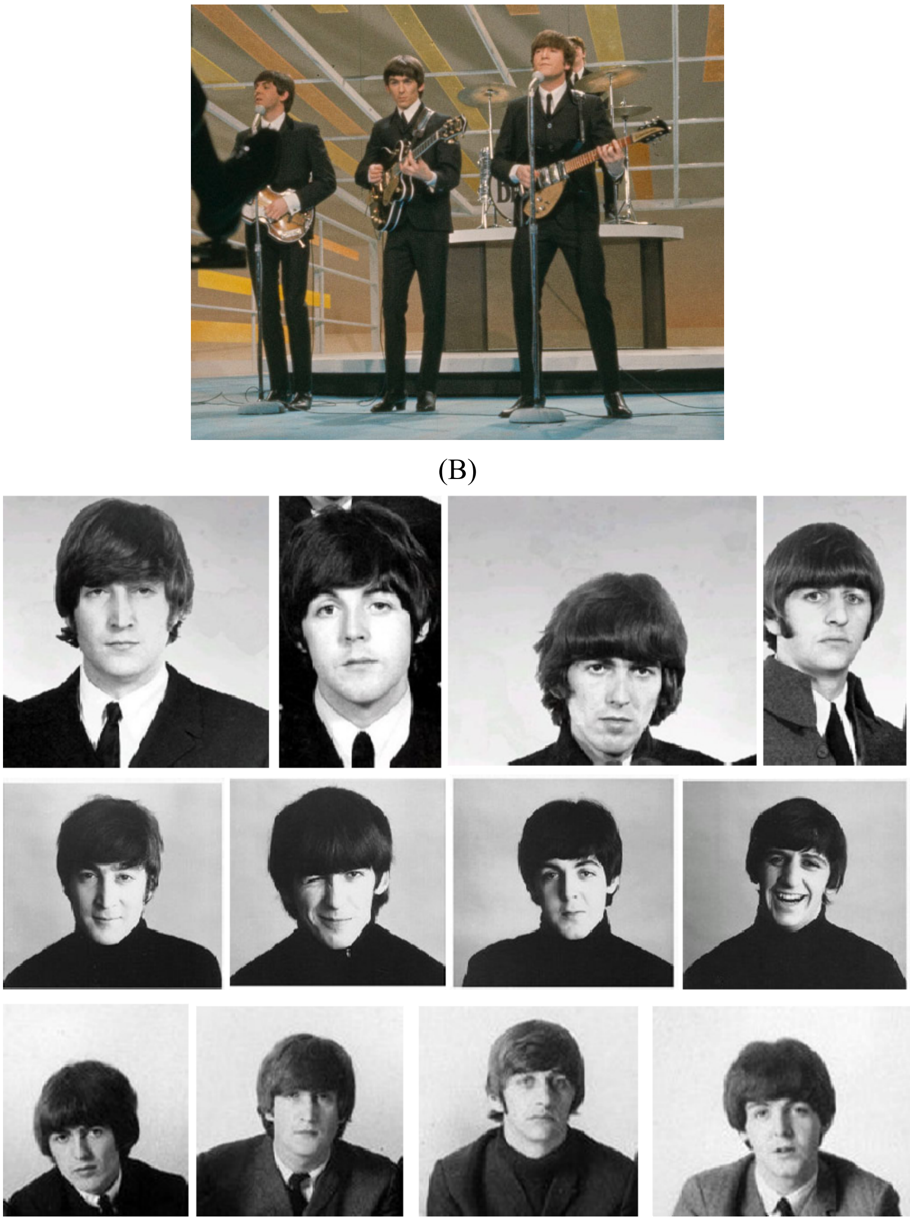

Note: Panel A. On the Sullivan show, 1964; Panel B. Individual images

Figure A.9: Hair Style Classifier – Probability of Images being “Beatles-style” Mop-Top

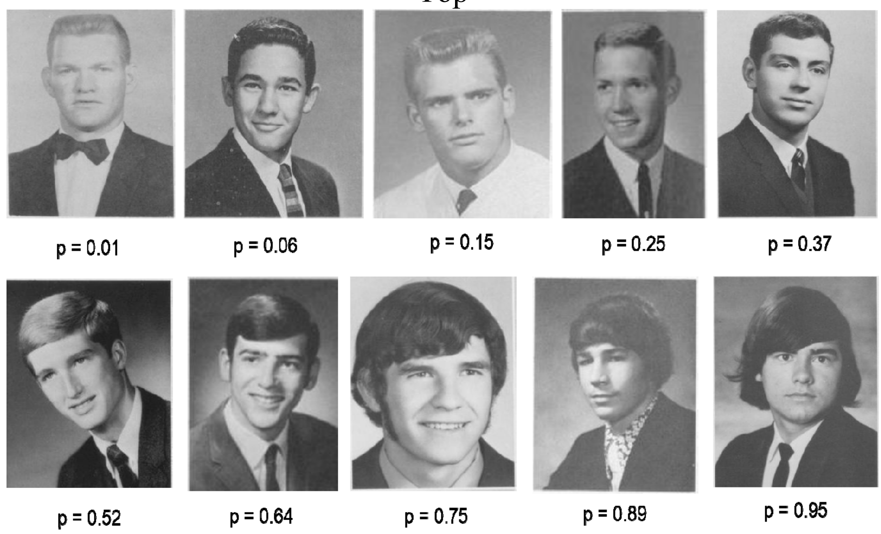

Figure A.10: Google Ngram for “Cultural Economics”

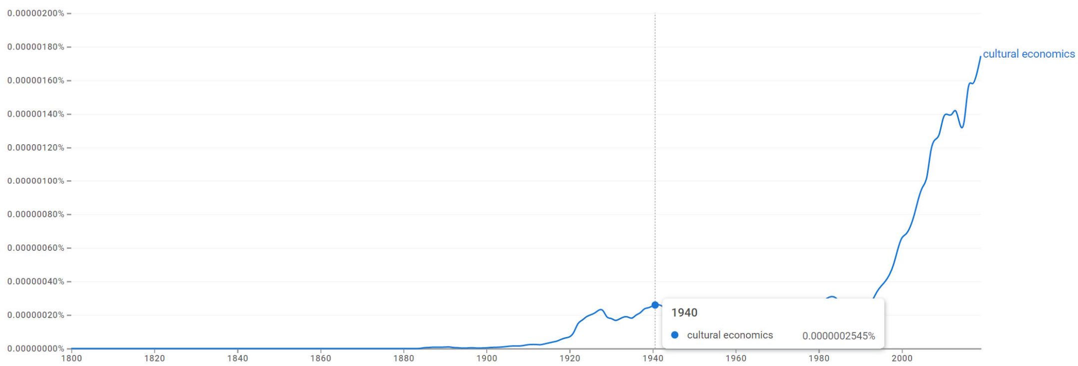

Appendix V: Tables

Table A.1: Summary Statistics US Sample

	Descriptive Statistics				
	Observations	**Mean**	**SD**	**Min**	**Max**
US Sample					
Predicted Probability of male	8,307,508	0.563	0.371	0.011	0.999
Predicted Probability of female	8,307,508	0.437	0.371	0.001	0.989
Predicted Probability of no tie	8,307,508	0.681	0.460	0.000	1.000
Predicted Probability of bow tie	8,307,508	0.060	0.231	0.000	1.000
Predicted Probability of tie	8,307,508	0.259	0.434	0.000	1.000
Predicted Probability of glasses	8,307,508	0.114	0.317	0.000	1.000
Predicted Probability of no glasses	8,307,508	0.886	0.317	0.000	1.000
Predicted Probability of facial hair	8,307,508	0.283	0.419	0.000	1.000
Predicted Probability of no facila hair	8,307,508	0.717	0.419	0.000	1.000
Predicted Probability of dress	8,307,508	0.260	0.413	0.000	1.000
Predicted Probability of shirt with collar	8,307,508	0.240	0.407	0.000	1.000
Predicted Probability of shirt without collar	8,307,508	0.201	0.376	0.000	1.000
Predicted Probability of suit	8,307,508	0.300	0.453	0.000	1.000
Predicted Probability of long hair	8,307,508	0.377	0.478	0.000	1.000
Predicted Probability of short hair	8,307,508	0.623	0.478	0.000	1.000
Predicted Probability of jewelry	8,307,508	0.220	0.409	0.000	1.000
Predicted Probability of no jewelry	8,307,508	0.780	0.409	0.000	1.000
Similarity - Within high-school cos simil at same year	8,307,508	0.785	0.125	0.020	1.000
Within high-school cos simil at t-1 years	7,955,443	0.727	0.140	0.009	1.000
Within high-school cos simil at t-2 years	7,616,503	0.737	0.141	0.000	1.000
Within high-school cos simil at t-3 years	7,028,798	0.746	0.143	0.000	1.000
Within high-school cos simil at t-4 years	6,939,912	0.741	0.145	0.000	1.000
Within high-school cos simil at t-5 years	6,872,336	0.737	0.146	0.000	1.000
Within high-school cos simil at t-6 years	6,794,470	0.732	0.148	0.000	1.000
Within high-school cos simil at t-7 years	6,724,771	0.728	0.150	0.000	1.000
Within high-school cos simil at t-8 years	6,646,108	0.724	0.151	0.000	1.000
Within high-school cos simil at t-9 years	6,575,551	0.720	0.152	0.000	1.000
Within high-school cos simil at t-10 years	6,470,332	0.717	0.153	0.000	1.000
Within high-school cos simil at t-11 years	6,384,757	0.714	0.154	0.000	1.000
Within high-school cos simil at t-12 years	6,293,761	0.711	0.155	0.000	1.000
Within high-school cos simil at t-13 years	6,207,637	0.708	0.155	0.000	1.000
Within high-school cos simil at t-14 years	6,128,178	0.706	0.156	0.000	1.000
Within high-school cos simil at t-15 years	6,050,348	0.704	0.157	0.000	1.000
Within high-school cos simil at t-16 years	5,955,388	0.702	0.157	0.000	1.000
Within high-school cos simil at t-17 years	5,856,351	0.700	0.157	0.000	1.000
Within high-school cos simil at t-18 years	5,756,390	0.698	0.158	0.000	1.000
Within high-school cos simil at t-19 years	5,635,429	0.697	0.158	0.000	1.000
Within high-school cos simil at t-20 years	5,544,238	0.695	0.158	0.000	1.000

Note: As we don’t have a balanced panel for every schools, there are some missing observations when calculating the cosine similarity between t and t-k.

Table A.2: Summary Statistics California Sample

	Descriptive Statistics				
	Observations	**Mean**	**SD**	**Min**	**Max**
California Data (25% of all available portraits in the state)					
Innovation - Forward Sim/Backward Sim -	918,685	1.000	0.046	0.545	1.545
Forward Similarity	918,685	0.824	0.107	0.112	0.956
Backward Similarity	918,685	0.824	0.106	0.106	0.963
Standardized values of Innovation	918,685	0.000	1.000	-9.900	11.865
Standardized values of Forward Similarity	918,685	-0.000	1.000	-6.672	1.237
Standardized values of Backward Similarity	918,685	-0.000	1.000	-6.801	1.313
Predicted Probability of no tie	918,685	0.625	0.479	0.000	1.000
Predicted Probability of bow tie	918,685	0.095	0.289	0.000	1.000
Predicted Probability of tie	918,685	0.281	0.445	0.000	1.000
Predicted Probability of glasses	918,685	0.061	0.239	0.000	1.000
Predicted Probability of no glasses	918,685	0.939	0.239	0.000	1.000
Predicted Probability of facial hair	918,685	0.308	0.429	0.000	1.000
Predicted Probability of no facila hair	918,685	0.692	0.429	0.000	1.000
Predicted Probability of dress	918,685	0.269	0.419	0.000	1.000
Predicted Probability of shirt with collar	918,685	0.208	0.385	0.000	1.000
Predicted Probability of shirt without collar	918,685	0.164	0.346	0.000	1.000
Predicted Probability of suit	918,685	0.359	0.473	0.000	1.000
Predicted Probability of long hair	918,685	0.429	0.488	0.000	1.000
Predicted Probability of short hair	918,685	0.571	0.488	0.000	1.000
Predicted Probability of jewelry	918,685	0.265	0.437	0.000	1.000
Predicted Probability of no jewelry	918,685	0.735	0.437	0.000	1.000

Table A.3: Summary Statistics Beatles Sample

	Descriptive Statistics				
	Observations	**Mean**	**SD**	**Min**	**Max**
Beatles Sample: all men in US sample between 1955-1975					
Predicted Probability of Beatles Hairstyle	1,444,383	0.350	0.247	0.000	1.000
Dummy - Percentile(Prob(Beatles)) > 75th	1,444,383	0.244	0.429	0.000	1.000
Km to closest CBS Station	1,444,383	58.474	49.267	0.643	289.305
Km to closest non CBS Station	1,444,383	48.577	49.293	0.508	358.689
Dummy - Distance to CBS Station < 60km	1,444,383	0.620	0.485	0.000	1.000
Dummy - Distance to non CBS Station < 60km	1,444,383	0.704	0.457	0.000	1.000

Table A.4: Demographic Statistics by county and CBS/No CBS

	No CBS	Yes CBS	Difference (No - Yes)
Total population, 1960	62807.705	172100.547	-109292.842**
Share of men, 1960	0.500	0.494	0.006**
Share non white people, 1960	0.092	0.089	0.004
Share of population (14 to 17) in schooling , 1960	0.869	0.877	-0.007
Share urban population, 1950	0.354	0.520	-0.166**
Unemployment Rate, 1950	0.042	0.041	0.001
Observations	319	318	

Descriptives are done at the county level. Table reports means and differences in demographic variables between CBS and No CBS group. We use Demographic Variables from Census 1950 and 1960 available at NHGIS. As the variable CBS/No CBS is constructed at high school level, there are some counties that have school with CBS signal and schools without CBS signal. Only counties without any school with CBS signal are considered as No CBS counties.