Coming apart? Cultural distances in the United States over time

Marianne Bertrand and Emir Kamenica*

July 2020

Abstract

We analyze temporal trends in cultural distances between groups in the US defined by income, education, gender, race, and political ideology. We measure cultural distance between two groups as the ability to infer an individual's group based on his or her (i) media diet, (ii) consumer behavior, (iii) time use, (iv) social attitudes, or (v) newborn's name. Gender difference in time use decreased between 1965 and 1995 and has remained constant since. Differences in social attitudes by political ideology, and somewhat by income, have increased over the last four decades. Whites and non-whites have diverged in consumer behavior. For all other demographic divisions and cultural dimensions, cultural distance has been broadly constant over time.

* University of Chicago Booth School of Business. We thank the University of Chicago Booth School of Business for financial support. Jihoon Sung and Feng Lin provided excellent research assistance. We thank Jann Spiess and Clara Marquardt for assistance with the implementation of the machine learning algorithms. Researchers own analyses calculated based in part on data from The Nielsen Company (US), LLC and marketing databases provided through the Nielsen Datasets at the Kilts Center for Marketing Data Center at The University of Chicago Booth School of Business. The conclusions drawn from the Nielsen data are those of the researchers and do not reflect the views of Nielsen. Nielsen is not responsible for, had no role in, and was not involved in analyzing and preparing the results reported herein.

What's great about this country is that America started the tradition where the richest consumers buy essentially the same things as the poorest. You can be watching TV and see Coca-Cola, and you know that the President drinks Coke, Liz Taylor drinks Coke, and just think, you can drink Coke, too. A Coke is a Coke and no amount of money can get you a better Coke than the one the bum on the corner is drinking. All the Cokes are the same and all the Cokes are good. Liz Taylor knows it, the President knows it, the bum knows it, and you know it.

- Andy Warhol

1 Introduction

While rural America watches *Duck Dynasty* and goes fishing and hunting, urban America watches *Modern Family* and does yoga in the park.[1] The economically better-off travel the world and seek out ethnic restaurants in their neighborhoods, while the less well-off don't own a passport and eat at McDonald's.[2] Conservatives give their boys masculine names like Kurt, while liberals opt for the more feminine-sounding options such as Liam.[3] While men play video games and watch pornography, women browse Pinterest and post pictures on Instagram.[4] These are just a few examples of the cultural distances across groups within America today. The presence of such cultural divides is not new – in the early 2000s, scholars emphasized racial differences in music tastes, language use, media diet, and consumer behavior[5] – but there is a perception that cultural distances are growing,[6] with a particular emphasis on increasing political polarization.[7]

These cultural distances may have important consequences. A large empirical literature in political economy documents that high levels of ethno-linguistic fragmentation hinder public good provision (Alesina et al. 1999; Alesina and La Ferrara 2005), decrease social capital (Alesina and La Ferrara 2000), and increase the probability of conflict (Montalvo and Reynal-Querol 2005). Moreover, Desmet et al. (2017) suggest that these outcomes especially worsen when cultural differences across ethnic groups are greater.

[1] https://www.nytimes.com/interactive/2016/12/26/upshot/duck-dynasty-vs-modern-family-television-maps.html

[2] https://www.theatlantic.com/national/archive/2011/03/americas-great-passport-divide/72399/

[3] See Oliver, Wood, and Bass (2016).

[4] http://www.pewinternet.org/2005/08/18/adult-content-online/; https://www.washingtonpost.com/news/the-switch/wp/2013/10/10/25-percent-of-men-watch-online-porn-and-other-facts-about-americans-online-video-habits/?utm_term=.450a3dfccb89; http://www.pewresearch.org/fact-tank/2015/08/28/men-catch-up-with-women-on-overall-social-media-use.

[5] See Waldfogel (2003), Wolfram and Thomas (2002), and Fryer and Levitt (2004).

[6] Fryer and Levitt (2004) document an increase in prevalence of distinctively black names over time. Focusing on differences across socio-economic groups within the white population, Murray (2012) writes: "It is not the existence of classes that is new, but the emergence of classes that diverge on core behavior and values – classes that barely recognize their underlying American kinship."

[7] See Kaufman (2002) on the increasing gender gap in party affiliation and Gentzkow (2016) on trends in polarization across party lines.

Sociologists such as Pierre Bourdieu (1984 [1979]) provide some theoretical foundations for the findings in the political economy literature. Bourdieu introduces the concept of cultural capital, which he associates with the set of tastes, mannerisms, or material belongings that one holds. Sharing cultural capital with others, Bourdieu argues, creates a sense of having a common identity. When cultural differences between groups increase, these groups find it more difficult to interact, communicate, and trust each other. Bourdieu was particularly concerned that cultural differences between the rich and the poor may damage social mobility. For example, students from poorer backgrounds might better integrate into college life if they can connect with better-off peers (Zimmerman 2017) but having little in common with those peers (e.g., having a different favorite TV show, different hobbies, different food preferences, etc.) may result in lower chances of forming friendships across income lines. Boneva and Rauh (2017) show that lower SES teens expect it will be more difficult to get along with other people in college, and that these expectations partly explain the socio-economic gap in the intentions to go to college. In other words, the lack of a shared culture may reduce the accumulation of both social and human capital. Bourdieu's logic also extends to groups not defined by income. African Americans and women may struggle to succeed in a predominantly white and male corporate America because of the greater difficulty of connecting with their majority-culture peers.[8]

Why might cultural divides be greater today than in the past? Technological progress could lead to cultural divergence: when there is only one channel to watch on television, or only one brand of ketchup to buy, all groups are mechanically constrained to share culture on these dimensions. Thus, increased choice sets might fuel cultural divergence. This, however, is not a foregone conclusion. First, it is possible that with only two TV shows, each show caters to one group or the other, but with thousands of shows, idiosyncratic preferences unrelated to group membership become the predominant driver of cultural choices. Second, universally-adopted new technologies might wipe out cultural differences; perhaps the rich and the poor used to spend their time differently from each other, but in the future everyone will just monitor their Facebook feed all day.

In this paper, we measure the extent of cultural distance across various groups in the US over time. In particular, we define groups based on income, education, gender, race, and political ideology.[9] We assemble multiple datasets in an attempt to capture as many aspects of people's cultural lives as possible, for as long as possible. Our data includes detailed information on media diet and consumer behavior (since 1992), attitudes (since 1976), time use (since 1965), and baby naming choices (since

[8] http://fortune.com/2016/08/11/african-american-executives-diversity-racism/

[9] In our Supplementary Materials, we also examine cultural distances by urbanicity (cf: Figure B.5) and age (cf: Figure B.6). Due to data constraints, we analyze cultural distances by urbanicity only in time use and social attitudes.

1960).[10] We define cultural distance in media diet between the rich and the poor in a given year by our ability to predict whether an individual is rich or poor based on her media diet that year.[11] We use an analogous definition for the other four dimensions of culture (consumer behavior, attitudes, time use, and baby names) and other group memberships. We use a machine learning approach to determine how predictable group membership is from a set of variables in a given year. In particular, we use an ensemble method that combines predictions from an elastic net, a regression tree, and a random forest (Mullainathan and Spiess 2017).

Figure 1 summarizes our findings. The results do not support the view that cultural divides are growing. With a few exceptions, the extent of cultural distance between groups has been broadly constant over time. One (unsurprising) exception is that men and women's time use became more similar from 1965 to 1995; perhaps more surprisingly, there has been no subsequent change in the gender differences in time use over the last 20 years.[12] Differences in social attitudes by political ideology, and somewhat by income, have increased since the 1970s. Whites and non-whites have diverged in consumer behavior. Nevertheless, our headline result is that for all other demographic divisions and cultural dimensions, cultural distance has been broadly constant over time.[13]

Two papers closest to ours are Alesina et al. (2017) and contemporaneous work by Desmet and Wacziarg (2019). Alesina et al. (2017) employ the European Value Survey and the General Social Survey (GSS) and report that, from the early 1980s to 2010, differences in social attitudes across countries in the EU and across nine large American states have somewhat increased. Desmet and Wacziarg (2019) define cultural distance between two groups as the share of total heterogeneity in responses to questions in the GSS that is not attributable to within-group heterogeneity. They examine eleven demographic divisions, including our five, and also report that cultural distances have been remarkably stable over time. Their results do somewhat contrast with ours as we find steady cultural divergences in social attitudes based on political ideology and income. The most important way in which our approaches differ is that Desmet and Wacziarg (2019) ignore correlation in answers across different questions in the GSS, but there are other differences as well.[14] In contrast to both Alesina

[10] As we discuss at greater length in the next section, we use Gfk Media Research Intelligence data for media diet and consumer behavior, General Social Survey for attitudes, American Heritage Time Use Study for time use, and California Department of Public Health birth records for baby names.

[11] This is the approach taken by Gentzkow et al. (2017) to measure differences between Democrats and Republicans in their Congressional speech.

[12] Both of the aforementioned patterns also hold if we only examine how men and women spend their time when they are not at work.

[13] These results are qualitatively the same if instead of employing the prediction method, we measure cultural distance simply as the Euclidean distance between average responses across groups (cf: Figure B.7).

[14] Suppose there are four equally sized groups, A, B, A', and B', who are asked a single binary question. Suppose the share of the individuals giving a specific response to this question is 0%, 5%, 10%, and 20% in the four groups, respectively. Our notion of cultural distance would indicate that groups A and B are closer together than groups A' and

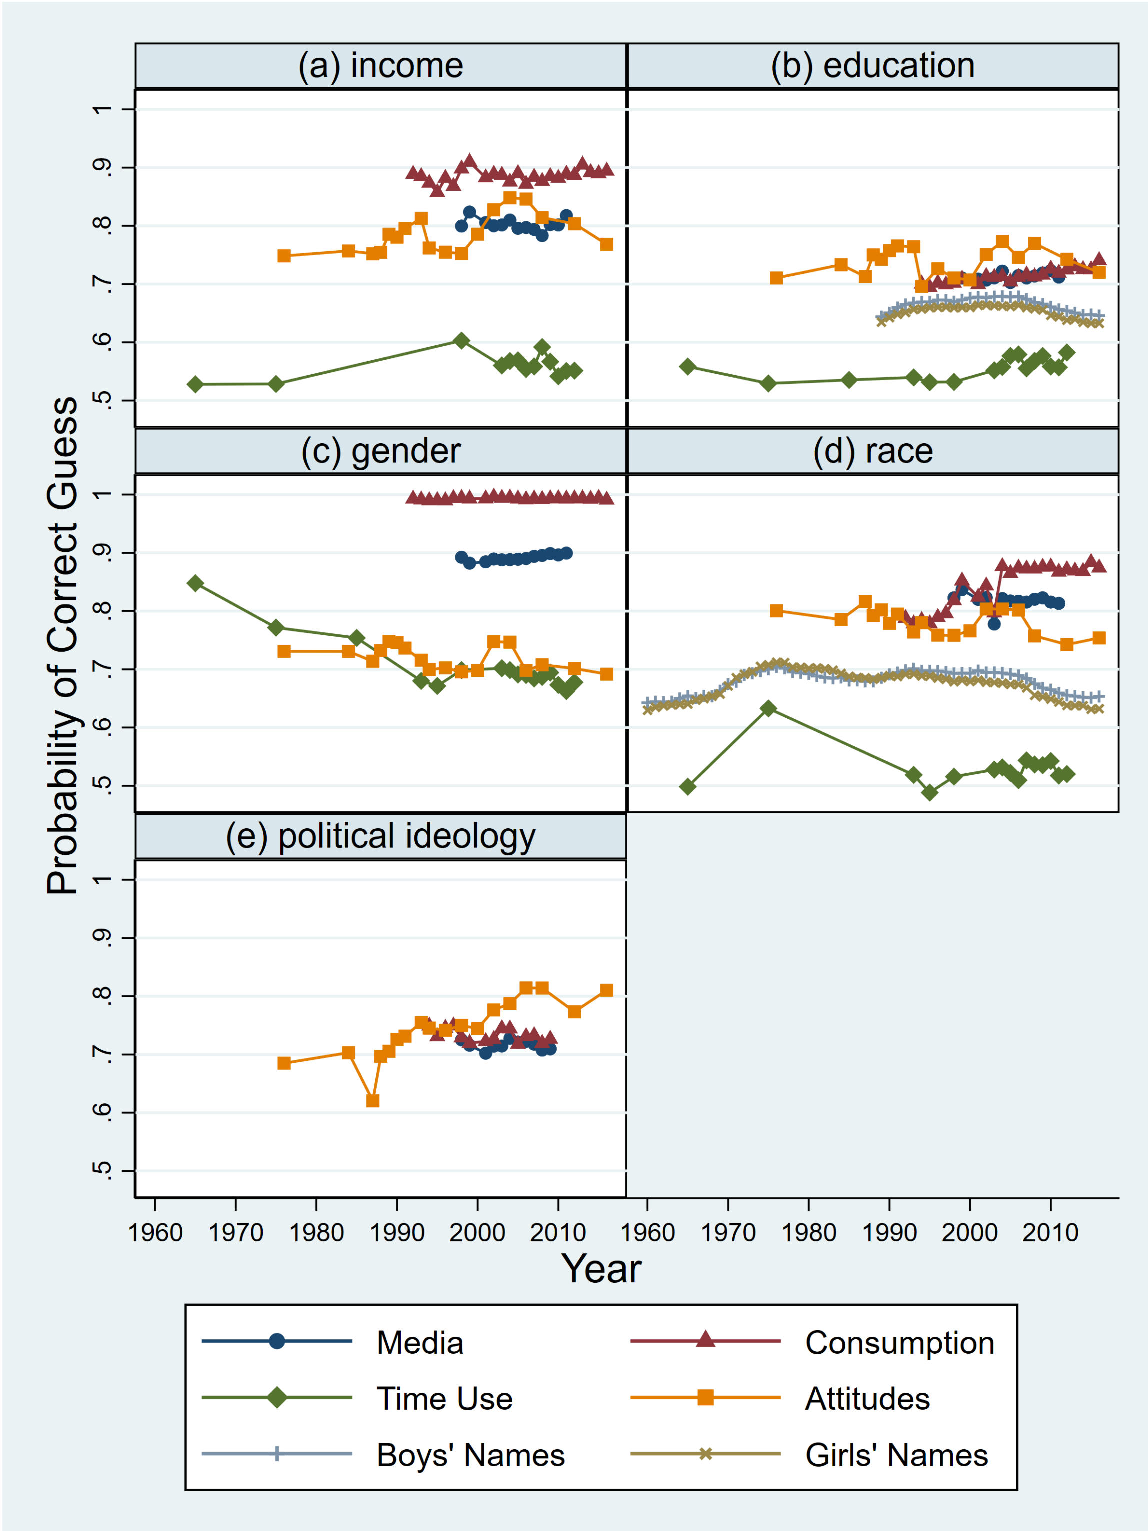

Figure 1: Cultural distances over time

Note: The likelihood, in each year, of correctly guessing (given a 50-50 prior) an individual's group membership (high or low income; high or low education; male or female; white or not; liberal or conservative) based on media diet, consumer behavior, time use, social attitudes, or a newborn child's name.

et al. (2017) and Desmet and Wacziarg (2019), we examine other dimensions of culture besides social attitudes, namely media diet, consumer behavior, time use, and baby names.

The rest of the paper is organized as follows. We briefly describe our datasets in Section 2. Section 3 lays out our empirical strategy and provides a discussion of our definition of cultural distance. The main results are reported in Section 4. Section 5 concludes.

2 Short data description

In this section, we provide a short description of the data that we utilize. The Data Appendix provides a more detailed description of our variables and sample construction.

Gfk Media Research Intelligence (MRI) data contains two questionnaires conducted each year between 1992 and 2016.[15] Demographic information, including household income, and some data pertaining to media exposure is obtained in a personal, face-to-face interview. The second questionnaire is left to be completed by the "principal shopper" of the household. This questionnaire asks whether the household has purchased, used, or owns a number of brands, products, and services. The second questionnaire also solicits data on which magazines the respondent reads and which TV shows and recently released movies the respondent has seen. For our media diet analysis, in any given year we use 879 to 1,129 binary (yes/no) answers about consumption of magazines, TV shows, and movies.[16] An example of a variable about media diet is "Have you seen the movie *Birdman* in the last 6 months?" For our consumer behavior analysis, we use 7,241 to 9,368 variables. Variables include questions such as "Has your household used Grey Poupon Dijon mustard in the last six months?", "Have you personally used a lipstick in the last six months?", "Do you own a dishwasher?", and "Have you personally used dry cleaning services in the last six months?"[17] For ease of exposition, from here on we will use the word 'product' to refer to 'products and services.' The sample size of the the MRI annual sample ranges from 15,352 to 22,033.[18]

B', whereas Desmet and Wacziarg (2019) would say that A and B exhibit a greater cultural distance than A' and B'. More generally, their definition allows for substantial cultural distance to be driven by arbitrarily rare behaviors, whereas our inference approach does not.

[15]Not all of the variables are available in every year. For example, we use data on magazines from 1992 to 2011, data on TV shows from 1992 to 2016, and data on movies from 1998 to 2016. That means that when we report trends in cultural distance based on overall media diet, we restrict our attention to years when all three of these subcomponents are available, namely 1998 to 2011.

[16]The MRI also asks questions about listening to radio and reading newspapers, but its coverage is too sparse to be useful for our purposes.

[17]As we discuss in greater length in next section, differences in how the rich and poor answer some of these questions (such as, "Do you own a dishwasher?") surely reflect the way that income affects budget sets rather than some notion of "cultural distance". We acknowledge the important distinction between income-constrained variables (such as the dishwasher one) and income-unconstrained ones (such as did you watch this movie or that one).

[18]An alternative dataset to MRI to predict group membership based on consumer behavior is the Kilts-Nielsen

To measure time use, we use the American Heritage Time Use Study (AHTUS). The AHTUS is a harmonized collection of diary data from 1965 to 2012. For the early years, the AHTUS covers roughly one year per decade, but since 2003 it includes annual surveys conducted by the Bureau of Labor Statistics.[19] Our harmonized data consists of 73 variables that indicate time spent on a specific activity (e.g., gardening), and we also include 8 aggregate activities (e.g., non-market work) from Aguiar and Hurst (2009). We restrict our sample to individuals who are employed full time. The sample size ranges greatly across years, from 666 to 10,107.

We use the General Social Survey (GSS) as our source of data on social attitudes. The GSS has been conducted annually since 1972. It collects stated attitudes on topics such as civil liberties, government policies, and morality. An example question is, "Are we spending too much, too little, or about the right amount on foreign aid?" We construct a set of variables with binary answers, e.g., in the example we construct three variables: (i) Are we spending too much on foreign aid?, (ii) Are we spending too little on foreign aid?, and (iii) Are we spending about the right about on foreign aid? The GSS also asks some questions about behavior, such as whether the respondent voted in the most recent Presidential election, which we also include in our analysis. We exclude questions about the respondent's assessment of his or her own financial situation.[20] There are many questions that the GSS asks only intermittently, so we drop some years in order to have a larger number of questions which are asked in every year we consider. This leaves us with 83 questions and 18 interspersed years between 1976 and 2016. The GSS often presents a given question to only two thirds of the survey participants. We impute the missing values for the remaining third based on the marginal distribution of responses in a given group.[21] The sample size of the GSS annual samples ranges from 1,074 to

Consumer Panel data (see https://research.chicagobooth.edu/nielsen). The Nielsen data tracks households' shopping behavior by asking participating households to scan the barcode of each purchased good after a shopping trip (using a scanning device provided by Nielsen). Nielsen differs from MRI in that it only covers products bought, not those used or owned. Moreover, MRI includes a broader set of products and services without a barcode. The main disadvantage of Nielsen over MRI for our purpose is a shorter time series: the Nielsen data only starts in 2004. Also, the Nielsen data has income brackets that are too broad to be able to identify top and bottom quartiles of the income distribution. Furthermore, the Nielsen data does not include information on political ideology. Figures B.8, B.9, B.10 compare predictability of respondents' education, gender and race based on consumer behavior in the MRI and Nielsen data. We restrict the analysis in these figures to single individuals, given Nielsen's focus on the household and MRI's focus on the "principal shopper." The results for education and race are similar across the two datasets in the overlapping years. As we discuss in Section 4.3.2, our empirical approach to measuring trends in cultural distance is not suitable for comparing consumer behavior of men and women in the MRI since we can perfectly predict group membership in this case. We do not reach this upper bound in the Nielsen data, where we observe convergence between men and women in consumer behavior. See also footnote 54.

[19] The AHTUS lacks information on household income in 1985, 1993, and 1995, so we drop those years from our analysis of cultural distance by income.

[20] When we examine cultural distance in social attitudes by ideology, we also exclude questions that directly pertain to ideology, namely political party affiliation and how the respondent voted in a presidential election.

[21] When we measure cultural distance by income, we impute the missing data based on the distribution of responses by the rich and the poor; when we measure cultural distance by education, we impute the missing data based on the distribution of responses by the more and less educated; etc.

3,669.

The California Department of Public Health (CDPH) birth data spans the years 1960 to 2016 and contains education and race of the mother in addition to the first name and biological gender of every baby born in California. In a given year, there are between 5,777 to 25,398 distinct boy's names and 9,739 to 35,341 distinct girl's names. The number of annual births in California ranges from 151,507 to 309,053 for boys and from 143,213 to 293,545 for girls.

All of our datasets study individuals in the United States who are 20 to 64 years old, except for the CDPH data where we put no restriction on mother's age.

3 Empirical approach

3.1 Compositional changes

Our empirical approach allows for the demographic composition of a group to change over time. For example, in the early 1970s, less than 10% of either the rich or the poor were Hispanic, but these days Hispanic individuals constitute 10% of the rich and 30% of the poor. Consequently, if Hispanic individuals are culturally distinct, this compositional change could lead to an increase in cultural distance between the rich and the poor. The same issue applies to other groups and other demographic characteristics; for instance, the share of people who are in our more educated group grows steadily over time and includes an ever-rising share of women.[22] Trends in cultural distance that are due to such compositional changes are something that we wish to capture rather than control for.[23]

3.2 Inference as a measure of cultural distance

In any given year, we say that two groups are further apart in their media diet (or consumer behavior, etc.) if we can predict more accurately which of the two groups a given individual belongs to based on his or her media diet (or consumer behavior, etc.). This approach follows Gentzkow et al. (2017), who measure partisanship of congressional speech by the ease with which one can infer a congressperson's

[22]Figure B.12 in the Supplementary Materials reports compositional changes in each of our groups.

[23]One exception may be the change in the age distribution of the rich and the poor. To the extent that trends in cultural distance are driven or hidden by the changes in the relative age between the rich and the poor, we may want to take those changes out, especially if we think that lifetime rather than contemporaneous income is a more meaningful way of defining who is rich and who is poor. In Figure B.13 in the Supplementary Materials, we define an individual as rich (poor) if he or she is in a household that is in the top (bottom) household income quartile among individuals in the same 5-year age bracket. We do not use information on household type (cf: discussion of household types in Section 4.1) since we do not have sufficient sample sizes to construct our groups based on both the age bracket and the household type. As seen in the figure, the results are mainly unaffected.

party from his or her speech.

Given some outcome space X, one could define the distance between two disjoint groups A and B based on any metric d on $\Delta(X)$ by letting the distance between the groups be equal to $d(\mu_A, \mu_B)$, where μ_A is the distribution of X in group A and μ_B is the distribution of X in group B.[24] Our predictability-based measure of distance implicitly sets d to be the total variation metric.[25] This measure has several features that are worth noting.

First, the measure takes no stance on which elements of X are close to another.[26] Suppose X consists of four elements: vodka, Sprite, 7 Up, and water. Suppose there are three equally-sized groups: in group A, 80% of people drink vodka and 20% drink water; in group B, 80% drink Sprite and 20% drink water; in group C, 80% drink 7 Up and 20% drink water. Our approach would say that the cultural distance between A and B is the same as the cultural distance between B and C, despite the fact that one might argue that B and C are closer since Sprite and 7 Up are more similar to each other than either is to vodka. Note, however, that this issue arises because of data limitations (no information on attributes of each $x \in X$) rather than because of our inference-based approach. Approaches employed by Alesina et al. (2017) and Desmet and Wacziarg (2019), as well as the Euclidean distance approach discussed below, all suffer from this issue. To address this issue completely, one would need access to auxiliary data that would answer questions such as: "Is gardening more similar to hunting or to playing with children?" or "Is film *Juno* more similar to *8 Mile* or to *The English Patient*?" That said, we can somewhat address these concerns by examining whether trends in cultural distance change when we aggregate across variables. Returning to the vodka/Sprite/7 Up/water example, suppose that we ascertain predictability of group membership based on product (soft-drink, alcohol, water) rather than brand consumption. In the hypothetical example above, we would see that groups B and C are closer to each other than they are to A. Yet, in our results below, no matter which demographic division we examine, we observe the same trends in cultural distance whether we study brand-level or product-level consumer behavior (cf: Figures 3, 6, 8, 10, and 13). This finding assuages concerns that our results overlook some systematic changes in the distances between individual brands.

Second, our measure of cultural distance has an upper bound that is achieved when one can

[24] In our setting, X would be the set of all possible vectors of answers to questions about media diet (or consumer behavior, etc.); $\Delta(X)$ denotes the set of probability distributions on X.

[25] If A and B are equally sized (as they are by construction in our approach), the ability to predict whether a person belongs to A or B is equal to $\frac{1}{2} + \frac{1}{2} d_{TV}(\mu_A, \mu_B)$, where d_{TV} denotes the total variation metric (cf: proof of Claim 3.30 in Mossel et al. 2014).

[26] Formally, this is related to the fact that total variation (unlike say the Prokhorov metric) does not require X to be a metric space itself.

perfectly predict group membership. Consequently, if some subset of variables is always sufficient to reach the upper bound, we would not be able to detect any changes in how similar the groups are on variables outside of that set. This turns out not to be an issue, however, since we are always far from the upper bound, except in the case of predicting gender using consumer behavior.[27] In fact, as the tables below show, the ability to predict group membership from any particular response item tends be substantially lower than the ability to do so using the entire set of responses.[28]

Third, a nice feature of our measure is that its units are easily interpretable. Contrast this with a measure that uses normalized Euclidean distance as the notion of cultural distance. Each of our datasets consists of a set of questions $q \in Q$. In the case of time use, each question is an activity and each answer is a number indicating the amount of time spent on that activity. In all of the other datasets, each question has a binary Yes/No answer. Let $\overline{q}_G$ denote the average answer in group G to question q;[29] we could then measure cultural distances between groups A and B by $\frac{\sqrt{\sum_{q \in Q}(\overline{q}_A - \overline{q}_B)^2}}{|Q|}$. In Figure B.7 in the Supplementary Materials, we replicate our results using this measure[30] and find qualitatively similar patterns, but with units of cultural distances that are harder to interpret. The fact that the inference-based metric and Euclidean distance paint the same overall picture suggests that our qualitative results are not highly sensitive to the choice of how to exactly measure cultural distance.

Finally, note that our inference-based approach does not allow us to aggregate across cultural dimensions that are not measured in the same dataset. For example, since we do not know the joint distribution of attitudes, time use, and income, we do not know how well one could predict income with both attitudes and time use. We do have data on media diet and consumer behavior in the same dataset, so in Figure B.14 in the Supplementary Materials we report cultural distance over time for these two aggregated dimensions. Again, we find no trends over time.[31]

[27] In the Supplementary Materials Figure B.11, we show cultural distance in consumer behavior by gender if we drop variables whose individual predictive power exceeds 75%. Once again we have nearly perfect ability to predict gender in every year and thus see no trend.

[28] For example, in 2009, the TV show that best discriminates liberals and conservatives is *The O'Reilly Factor*. Knowing whether an individual watches this show allows us to predict political ideology with 57% accuracy. By contrast, knowing the full set of TV shows that a person watches identifies political ideology with the accuracy of 71%.

[29] For time use, this is the average number of minutes spent on the activity; in all other cases, this is the share of respondents answering Yes.

[30] Note that this approach, like the memetic fractionalization index used by Desmet and Wacziarg (2019), ignores correlations between answers.

[31] If distance is measured based on the normalized Euclidean distance, it is possible to aggregate across datasets. The overall trend in cultural distance is then just the (weighted) average of the trends depicted in Figure B.7.

3.3 Machine learning

We use a machine-learning ensemble method to determine how predictable group membership is from the variables in each dataset in each year. The ensemble method consists of running separate prediction algorithms (elastic net, regression tree, and random forest) and then combining the predictions of these algorithms with weights chosen by OLS (Mullainathan and Spiess 2017). For each dataset, year, and group division (e.g., time use data by gender in 2010), we first split the dataset into a training sample (70% of the data) and a hold-out sample (30% of the data). We empirically tune each algorithm on the training sample by cross-validation. In particular, we partition the training data into five folds. For a given fold, we fit the algorithm on the other four folds for every value of the tuning parameter. We average the squared-error loss function for each tuning parameter over these four folds and choose the tuning parameter that minimizes the loss. Using the chosen tuning parameter, we obtain a prediction (e.g., probability that the respondent is a woman) for every observation in the given fold. Repeating this process for each fold gives us a prediction for every observation in the training sample for each of the three algorithms. We regress (using OLS) group membership on the three predictions (from the three algorithms) in the full training sample. We use the coefficients from this regression to combine the three algorithms into the ensemble prediction in the next step.

Before turning to our hold-out sample, we fit each algorithm on the full training sample for every value of the tuning parameter, and obtain the optimal tuning parameter that minimizes the mean squared-error loss. For each observation in the hold-out sample, we derive the prediction of each algorithm using the model estimated in the training sample under the optimal tuning parameter. We then compute the ensemble prediction for that observation using the aforementioned OLS coefficients. We then guess a respondent's group affiliation based on the ensemble prediction: if the probability that a respondent is in a group is above $\frac{1}{2}$, we guess that she is in that group; otherwise, we guess that she is in the other group. We define cultural distance (for each dataset, year, and demographic category) as the predictability of the group membership, i.e., the share of the guesses in the hold-out sample that are correct.[32]

When measuring differences in how groups name their babies, we do not use the ensemble method. One issue is that the hundreds of thousands of observations and tens of thousands of distinct names make the ensemble approach computationally very demanding. More importantly, the simple structure of the data whereby each observation takes on the value of zero for all but one of the tens of thousands of variables (each baby has only one name) means that there is no real benefit of employing machine

[32]All of our results about trends over time are qualitatively the same if we use any one of the algorithms (elastic net, regression tree, or random forest) rather than combining them into the ensemble prediction.

learning methods. After we split the data into training and hold-out samples, we simply guess that an individual in the hold-out sample belongs to the group that has a higher share of that given name in the training sample.[33] For example, in 2016, 0.72% of the more educated mothers and 0.92% of the less educated mothers in the training sample name their newborn girl Emily; since $92 > 72$, we guess that every Emily in the hold-out sample in 2016 was born to a less educated mother.

3.4 Data over time

We need to ensure that the "quality" of our datasets – in terms of number of observations and the availability of relevant variables – is constant over time. Otherwise, our ability to predict group membership might change over time for reasons unrelated to any changes in cultural distance. The solution to time-varying sample sizes is straightforward. For each dataset and demographic group, we equalize the number of observations in each year and demographic group as follows. Denoting by n the minimum sample size across years and groups (e.g., when computing the cultural distance in time use by education, the smallest year-group are the less educated in 1965), we randomly select n observations for every year-group. This yields a "balanced" dataset with the same number of observations in each year with half of the observations in each of the two groups. We then compute the predictability of group membership – as described in the previous subsection – in this balanced dataset. We repeat this procedure a number of times,[34] drawing a new random sample each time, and then we take the average predictability of group membership (averaged across the draws) as our measure of cultural distance. Note that this means that the sample sizes reported in Section 2 are larger than the balanced sample sizes that we use to make each prediction of group membership.

Another important consideration is related to the changes over time in the particular questions asked to survey participants. When it comes to the GSS and AHTUS data, we insist on having the same set of variables in each year. When it comes to the time use data, we think the set of activities that people can spend their time on has not changed that much over time, with the exception of spending time on a computer. Therefore, if the set of variables in the time use data expanded over time, this would likely be a reflection of improvement in data collection rather than a reflection of actual changes in the ways people are spending their time. Therefore, we use the crosswalk provided by the University of Oxford Center for Time Use research[35] to harmonize time use variables across

[33] If a name in the hold-out sample does not appear in the training sample, we assume the baby was born to whichever group has a greater tendency to give a newborn a unique name.

[34] In the GSS, AHTUS, and CDPH, we take 500 draws. In the MRI, we take 25 draws for the media diet and 5 draws for consumer behavior.

[35] See https://www.timeuse.org/ahtus/documentation.

years.[36] With regard to social attitudes, the GSS often asks a particular question only intermittently, and we do not believe that this is a reflection of the fact that this question was only relevant in the years the question was asked. Consequently, we limit the set of GSS variables and years we use in a way that ensures that each variable is available in each year.[37]

When it comes to the MRI data, we embrace the variations in the set of questions asked over time, both for media diet and consumer behavior. Our understanding is that the MRI seeks to include questions about all media items (magazines, TV shows, movies) and consumer products that are relevant at the time. For example, each year the MRI asks respondents about whether they had seen a number of newly released movies. While the number of movies that the MRI asks about is reasonably constant, ranging from 83 to 97 across years, the set of movies they ask about of course changes completely from year to year, reflecting the new releases. We assume that the changes in the variables about TV shows, magazines, products, and brands similarly reflect real changes in consumers' choice sets. While this assumption surely does not hold perfectly – for example, there is a big jump in the number of TV shows in the data in 2009 when the MRI added cable shows to the survey – it provides the most natural approach for measuring cultural distance when cultural elements are rapidly changing over time.

Just like for the MRI data, we embrace the year-to-year variation in the set of baby names in the CDPH data.

3.5 Confidence intervals

Throughout, we report our estimates of cultural distances without confidence intervals. One way to approach inference in our setting would be via subsampling (e.g., Politis et al. 1999), but our sample sizes are too small for the ensemble algorithm to perform well on partitioned data. That said, the fact that our measure of cultural distance tends to be pretty similar across years informally suggests that it is estimated reasonably precisely; otherwise, it would be highly unlikely for the estimates to fall so close to one another. We have also confirmed that if we add to the data a synthetically constructed

[36]The AHTUS asks about computer use only after 1985. We impute zero computer use for all respondents prior to 1985. The AHTUS does not ask about smartphone usage. All activities related to the use of computer and internet for leisure are aggregated under computer use.

[37]In contrast to the time use data, we are less confident in our decision to harmonize the set of GSS questions over time. It might very well be that the GSS changes questions it asks from one year to the next because the set of most important societal issues is changing, in which case there might be some argument for embracing the change in variables. Without a more specific model of how the GSS drafts their survey instrument each year, it is difficult to sign the potential bias induced by our harmonization choice. For example, if the GSS drops questions once everyone agrees on the answer and keeps only those questions where disagreement remains, our approach might underestimate cultural convergence over time. Alternatively, if the GSS systematically adds questions that have become more controversial, our approach might underestimate the increase in cultural distance over time.

variable whose correlation with group membership increases over time, we indeed observe a growing cultural distance using our method.

4 Results

We organize the results by group divisions: income, education, gender, race, and political ideology. For each group division, after a discussion of the overall patterns, we dive in greater detail into the five broad cultural components. For the media, we investigate the separate cultural influences due to TV watching, movie watching, and magazine readership. For consumer behavior, we investigate the separate roles of products vs. brands. For social attitudes, we consider the separate influence of thematic sub-categories, such as views related to the role of government in society or views related to civil liberties. Throughout, we try to enrich the results with a discussion of specific cultural traits that are most distinctive across groups at a given point in time. Rather than report every result for every group, we highlight the data we find most informative in the text and report the additional results in the Supplementary Materials.

4.1 Income

A vast literature in labor economics has documented the rise in income inequality in the US since the late 1970s. While a large share of this literature in recent years has focused on "top income inequality" (e.g., the share of total income going to the top 1 percent, or top 0.1 percent), it is also well understood that technological change and global competitive pressures have contributed to broader changes in income inequality across individuals and households (e.g., Autor et al. 2008, Meyer and Sullivan 2017). The causes of the rise in income inequality are now reasonably well understood, but the consequences are less clear. We are particularly interested in whether greater income inequality has led to a greater cultural gap between the rich and the poor. Technological change and a growing supply of goods and services also may have exacerbated or attenuated any changes in the cultural gap between the rich and the poor.[38] As discussed previously, increased cultural distance between rich and poor could be particularly damaging to social mobility. A high-income manager may promote the subordinate with whom she has the friendliest interactions around the water cooler, and that favorite subordinate will likely come from a high-income background if tastes, views, and experiences

[38] Jaravel (2017) documents that newly developed products in the US tend to target high-income households; this force could create a new set of goods around which a "culture of being rich" could coalesce. At the same time, other technological developments, such as certain forms of social media, could lead to cultural convergence between income groups by providing inexpensive goods that appeal to individuals of all income levels.

are sharply different between income classes.

We define an individual as rich (poor) if he or she is in a household that is in the the top (bottom) quartile of household income among households of the same type. We put households into four types: (i) a single adult with no dependents, (ii) two adults with no dependents, (iii) a single adult with dependent(s), and (iv) two adults with dependent(s).[39,40] We use the Current Population Survey to identify the distribution of household income for each household type in each year. We focus on the top and the bottom quartile (as opposed to, say, the top and the bottom half or the top and the bottom decile) to balance a desire to make the rich and the poor as different in their income as possible and the pragmatic need to keep our sample sizes sufficiently large.[41] Given our definition of rich and poor and our procedure for equalizing sample sizes described in Section 3.4, each prediction of income in a given year is based on 5,810 observations in the MRI, 394 observations in the GSS, and 200 observations in the AHTUS.

We also consider alternative definitions of rich and poor, comparing (i) top half vs. bottom half, (ii) top quartile vs. everyone else, and (iii) bottom quartile vs. everyone else. Under all of these alternative definitions, our qualitative results remain the same (cf: Figure B.16). Throughout the analysis, we use contemporaneous income rather than wealth or lifetime income. While the latter two measures might seem more closely related to what it means to be rich or poor, we do not have data on wealth or lifetime income.

Panel (a) of Figure 1 summarizes our results. There is no evidence of an increasing cultural gap between the rich and the poor based on media diet, consumer behavior, or time use. The patterns regarding media and consumer behavior, where our sample size is the largest, show that cultural distance is essentially the same in each year. Knowing what TV shows and movies someone watches

[39]We define a household as having dependents if there are children under 18 or if the household has more than two adults. This may induce some measurement error; we would code three roommates as two adults with a dependent and would code a single adult taking care of a parent or a sibling as two adults with no dependents.

[40]An alternative to this approach would be to use an equivalence scale to adjust for the size and the composition of the household. The downside of the alternative approach is that all standard scales (per-capita income, the Oxford scale, the OECD-modified scale, and the square root scale) systematically label households with (more) children as more likely to be poor. Consequently, the ability to predict household income then primarily stems from the ability to predict whether there is a child in the household: tell-tale signs of "being poor" become watching *SpongeBob SquarePants* or buying children's medications. Under our preferred approach, there is by construction no relationship between poverty and the presence of children, and the relationship between poverty and the number of children is weaker. If we do ignore household types and define rich (poor) as the top (bottom) quartile of household income divided by the square root of the household size, we observe the same temporal trends in cultural distances between the rich and the poor (cf: Figure B.15).

[41]Income variables in the GSS, the AHTUS, and the MRI are income brackets, so the top and bottom quartile cutoffs occur within an income bracket rather than at the boundary. We classify respondents into the top and bottom quartiles to minimize miscategorization (please refer to our Data Appendix for details). The share of miscategorization never exceeds 10 percent. More importantly, the extent of miscategorization does not explain any of the variance in measured cultural distance: if we regress measured cultural distance on a linear time trend and the dummy for the cultural dimension, adding the extent of mismeasurement increases R^2 from 0.4903 to 0.4905.

and what magazines a person reads allows us to correctly predict the person's income group about 80 percent of the time. Knowing what goods and services a person buys, including particular brands, allows us to correctly predict the person's income group between 85 to 89 percent of the time, with no apparent time trend. The gap in how rich and poor spend their time has also been relatively constant, without a clear trend. We do observe some divergence of attitudes between income groups between the mid-1990s and the mid-2000s.

4.1.1 Media diet and income

As we saw before, there has been no trend in the difference in overall media diet by income. Figure 2 shows how well one can predict income group over time based on the three separate components of media diet: TV programs, movies, and magazine readership.

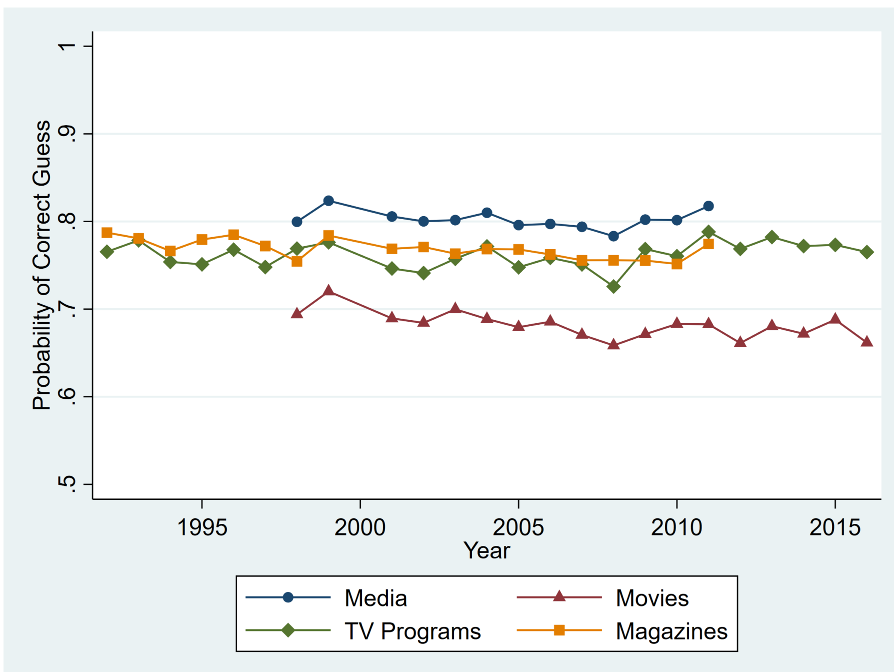

Figure 2: Cultural distance by income over time: media diet

Note: Data source is the MRI. Sample size each year is 5,810. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's income in the hold-out sample each year. The procedure to guess income in the hold-out sample was repeated 25 times, and the share of guesses reported is the average of these 25 iterations.

The results reveal that there has been no divergence between the rich and the poor in any of the sub-components of media diet. These groups watch different movies and TV shows and read different magazines, but the extent of these differences has been nearly constant across the last quarter-century. We also see that income differences in consumption of magazines and TV shows is somewhat greater than the difference in consumption of newly released movies.[42]

It is particularly interesting that predictability of income based on TV shows has been constant over time given that there have been substantial changes both in the number of TV shows available and the number of TV shows watched in each income group.[43] As we discussed in the introduction, there is no mechanical reason that necessitates a relationship between the number of options and the size of the cultural gap, but it is nonetheless striking that the cultural gap has been roughly constant even as the number of options has changed substantially.

Of course, the fact that cultural distance in media diet has been constant over time does not imply that the particular magazines, TV shows, and movies that drive this distance have been the same from year to year. This is most obvious in the case of movies where each year brings a crop of new releases. Panels (a), (b), and (c) of Table 1 report, respectively, the ten movies, TV shows, and magazines that are individually most informative of income. We do this for three separate years spanning the beginning, the middle, and the end of each dataset. Consider movies for example (Panel (b)). If we could ask a person a single question of the form "Did you see movie X" and then guess, based on the answer, whether the person is rich or poor, the best question to ask in 1998 would be "Did you see Jerry Maguire"? Guessing that the person is rich if they saw the movie and guess that they are poor otherwise (a higher share of the rich saw the movie) would lead to a correct guess 58.7% of the time. All of the ten most informative movies in 1998 are movies that are distinctly rich-people movies. By contrast, some of the most informative movies in 2007 are movies whose audiences are distinctly poor. As shown in Panel (a), even though cultural distance in TV consumption has been constant, the specific TV content that is most predictive of income has changed over time. Panel (c) reveals more stability in the list of magazines whose readership is most indicative of income group. While the rank ordering varies somewhat, three magazines are always in the top three: *Newsweek*, *Consumer Reports*, and *Time.*

4.1.2 Consumer behavior and income

Figure 3 reports the predictability of income over time separately based on products and brands individuals report buying or owning. We also again report the predictability of income based on the entire consumer behavior data as a benchmark.

[42]This comparison is more meaningful than comparing, say, the cultural distance in TV watching with cultural distance in time use, since those are measured in datasets with very different sample sizes. Of course, even with equal sample sizes, there are important differences in measurement across cultural dimensions. As shown in Figure B.18 in the Supplementary Materials, the average number of TV shows watched is about 30, while the average number of new movies watched is less than 10.

[43]Figure B.17 plots the number of TV shows in the data over time.

Table 1: TV shows, movies, and magazines most indicative of being high-income

Panel (a) TV shows

1992		2004		2016	
Watched *US Open (Golf)*	58.8%	Watched *2002 Winter Olympics*	63.9%	Watched NFL football games	60.6%
Didn't watch *Rescue 911*	58.6%	Watched NFL football games	58.8%	Watched MLB baseball games	57.6%
Watched NFL football games	57.4%	Watched MLB baseball games	58.1%	Watched *Property Brothers*	56.4%
Watched MLB baseball games	57.3%	Watched college football games	56.2%	Watched NCAA basketball games	56.3%
Watched *Wimbledon*	57.0%	Watched NCAA basketball games	55.9%	Watched *Love It or List It*	56.2%
Didn't watch *Cosby Show*	56.6%	Watched *US Open (Golf)*	55.1%	Watched *House Hunters*	56.2%
Didn't watch *Unsolved Mysteries*	56.5%	Didn't watch *Cops*	54.9%	Watched college football games	56.1%
Didn't watch *The Oprah Winfrey Show*	56.0%	Watched *Wimbledon*	54.7%	Watched *Academy Awards*	55.6%
Watched NCAA basketball games	55.9%	Watched *Academy Awards*	54.6%	Watched *Flip or Flop*	55.3%
Didn't watch *Cosby*	55.8%	Watched *The Masters*	54.6%	Watched *Anderson Cooper 360°*	55.0%

Panel (b) Movies

1998		2007		2016	
Watched *Jerry Maguire*	58.7%	Didn't watch *Big Momma's House 2*	54.2%	Watched *Gone Girl*	54.3%
Watched *First Wive's Club*	56.2%	Didn't watch *Final Destination 3*	54.0%	Didn't watch *Annabelle*	52.6%
Watched *The English Patient*	55.5%	Watched *Walk The Line*	53.6%	Watched *The Hunger Games*	52.5%
Watched *Michael*	54.2%	Didn't watch *Saw II*	53.2%	Didn't watch *TMNT*	52.5%
Watched *Air Force One*	53.9%	Watched *The Devil Wears Prada*	53.1%	Didn't watch *Ouija*	52.4%
Watched *Ransom*	53.1%	Watched *Pirates of The Caribbean 2*	52.8%	Didn't watch *No Good Deed*	52.3%
Watched *One Fine Day*	53.1%	Watched *The Da Vinci Code*	52.6%	Didn't watch *Let's Be Cops*	52.0%
Watched *My Best Friend's Wedding*	53.1%	Watched *Syriana*	52.4%	Watched *Interstellar*	52.0%
Watched *Evita*	53.0%	Watched *Brokeback Mountain*	52.3%	Watched *Kingsman*	51.7%
Watched *The Chamber*	52.7%	Watched *The Chronicles of Narnia 1*	52.3%	Watched *Unbroken*	51.6%

Panel (c) Magazines

1992		2002		2011	
Read *Newsweek*	61.8%	Read *Newsweek*	60.8%	Read *Consumer Reports*	58.4%
Read *Consumer Reports*	60.7%	Read *Time*	59.8%	Read *Newsweek*	57.8%
Read *Time*	60.5%	Read *Consumer Reports*	58.8%	Read *Time*	57.4%
Read *Parade*	59.5%	Read *U.S.News & World Report*	57.1%	Read *People*	56.8%
Read *Business Week*	59.0%	Read *Business Week*	57.1%	Read *Travel & Leisure*	55.9%
Read *US News & World Rpt.*	59.0%	Read *Money*	56.9%	Read *The Economist*	55.5%
Read *National Geographic*	58.1%	Read *Forbes*	56.8%	Read *The New Yorker*	55.3%
Read *Money*	58.1%	Read *Fortune*	56.4%	Didn't read *TV Guide*	55.0%
Read *People*	57.2%	Read *People*	56.2%	Read *Forbes*	55.0%
Read *Forbes*	56.7%	Read *Architectural Digest*	55.3%	Read *Real Simple*	54.9%

Note: Data source is the MRI. Sample size in all panels is 5,810. Reported in each column are the 10 cultural traits most indicative of being rich in that year. The numbers indicate the likelihood of guessing correctly whether an individual is rich or poor based on the answer to the question. For example, in 1992, knowing whether a person watched *US Open (Golf)* allows us to guess income correctly 58.8% of the time, whereas knowing whether a person watched *Rescue 911* allows us to guess income correctly 58.6% of the time. An affirmative answer to "Did you watch *US Open (Golf)*?" and a negative answer to "Did you watch *Rescue 911*?" indicate that the person is rich.

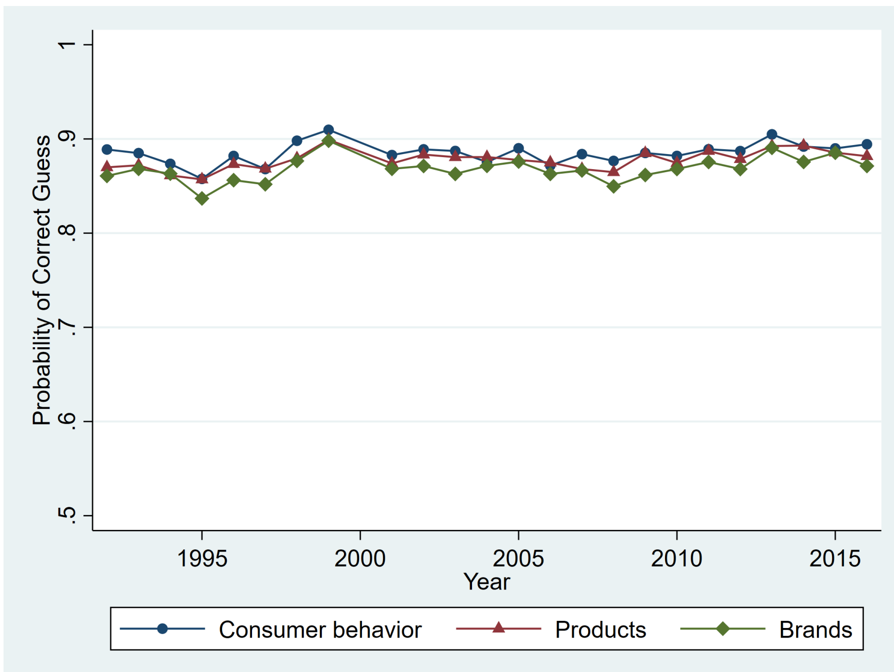

Figure 3: Cultural distance by income over time: consumer behavior

Note: Data source is the MRI. Sample size each year is 5,810. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's income in the hold-out sample each year. The procedure to guess income in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these 5 iterations.

Table 2: Products and brands most indicative of being high-income

Panel (a) Products

1992		2004		2016	
Own a built-in dishwasher	72.2%	Bought a new vehicle	74.0%	Traveled in the continental US	71.6%
Used dishwasher detergent	71.3%	Used dishwasher detergent	72.3%	Own bluetooth on vehicle	71.3%
Traveled domestically	69.2%	Own a dishwasher	71.6%	Own a valid passport	70.7%
Own a telephone answering machine	66.7%	Traveled in the continental US	71.4%	Own heated/cooled seat on vehicle	70.6%
Used dry cleaning services	66.0%	Own a stereo on vehicle	70.5%	Used dishwasher detergent	70.4%
Own a fireplace	65.9%	Own a personal computer	70.3%	Own a built-in dishwasher	69.9%
Own fine china dinnerware	65.8%	Own frequent flyer club membership	69.4%	Bought on internet	69.0%
Used tune-up service for vehicle	65.8%	Own air bags on passenger side	69.2%	Own frequent flyer club membership	68.8%
Own power locks on vehicle	65.8%	Traveled domestically with no travel agent	69.2%	Own sunroof/moonroof on vehicle	68.2%
Own a garage door opener	65.5%	Own a desktop computer	68.9%	Own side-impact air bags on vehicle	68.2%

Panel (b) Brands

1992		2004		2016	
Bought Kodak (film)	62.8%	Own AAA membership	62.1%	Own an iPhone	69.8%
Used Thomas' (English muffins)	61.9%	Own a Dell computer	60.6%	Own an iPad	67.4%
Used Federal Express	61.8%	Used Land O' Lakes (butter)	59.2%	Used Verizon Wireless (cellular)	61.5%
Used Grey Poupon Dijon (mustard)	61.5%	Used Bertolli (salad/cooking oil)	58.9%	Own AAA membership	60.9%
Own AAA membership	60.2%	Own a computer with Windows XP	58.9%	Own Hilton HHonors membership	60.6%
Used Johnson's (dental floss)	59.0%	Used Kikkoman (soy sauce)	58.8%	Used AMC	59.0%
Used Scotch Magic (transparent tape)	58.6%	Used AT&T (long distance call service)	58.5%	Not own an Android phone	58.9%
Used Cascade - Lemon (dish. detergent)	58.5%	Used Scotch Magic (transparent tape)	58.5%	Own United MileagePlus membership	58.9%
Used Philadelphia (cream cheese)	58.4%	Didn't use BIC (lighters)	58.5%	Used AT&T (cellular)	58.9%
Used Delta domestically	58.2%	Used Reynolds Wrap (aluminum foil)	58.3%	Used Kikkoman (soy sauce)	58.9%

Note: Data source is the MRI. Sample size in all panels is 5,810. Reported in each column are the 10 cultural traits most indicative of being rich in that year. The numbers indicate the likelihood of guessing correctly whether an individual is rich or poor based on the answer to the question. For example, in 1992, knowing whether a person owns a built-in dishwasher allows us to guess income correctly 72.2% of the time, whereas in 2004 knowing whether a person bought a BIC lighter allows us to guess income correctly 58.5% of the time. An affirmative answer to "Do you own a built-in dishwasher?" and a negative answer to "Did you buy a BIC lighter?" indicate that the person is rich.

We see that the flat trend line previously reported for consumer behavior extends to these two subsets of the consumer data: the probability of correctly guessing someone's income based on the products or brands consumed is essentially the same over the quarter-century of available data. Products and brands consumed have similar predictive power. Moreover, the aggregation of the product and brand features does not not change much the predictive power of the model compared to using products or brands separately.[44]

Panels (a) and (b) of Table 2 show the variables over time in the product and brand data respectively that are individually most indicative of income group. As in the previous subsection, the fact that cultural distance has not changed over time does not imply that the same features distinguish

[44] This is consistent with Neiman and Vavra (2018) in that their finding of increased household spending concentration over time in the Nielsen data is not driven by a widening gap between the goods purchased by different income groups.

income groups in each year. Consider panel (a) first. Household goods dominate the 1992 list: owning a dishwasher, a telephone answering machine, a fireplace, fine china dinnerware, and a garage door opener separate rich and poor. Household goods are again present in the 2004 list (dishwasher, personal computer) as are vehicle-related variables (new car, car stereo, airbags). While travel-related experiences and services are already present in the list of most indicative variables in 1992 and 2004 (domestic and international travel, frequent flyer programs), a travel-related feature reaches the top of the list in 2016. Beyond this, the ownership of car gadgets (bluetooth, heated seats, sunroof, airbags) remains an important income class differentiator in 2016.

Some of the most predictive brands, especially in the earlier years, are food items. Two of the brands most predictive of high income in 1992 are Thomas' English muffins and Grey Poupon Dijon mustard.[45] By 2004, Land O'Lakes butter, Bertolli oil, and Kikkoman soy sauce are among the most predictive brands. That said, by the end of the sample, ownership of Apple products (iPhone and iPad) tops the list. Knowing whether someone owns an iPhone in 2016 allows us to guess correctly whether the person is in the top or bottom income quartile 69.8 percent of the time. Across all years in our data, no individual brand is as predictive of being high-income as owning an Apple iPhone in 2016.

As the results above suggest, some of the differences in consumer behavior between the rich and the poor probably reflect differences in budgets rather than in anything that we should call culture. Presumably, the poor do not own dishwashers because they cannot afford them (or have no space for them) rather than because they have a cultural preference for washing dishes by hand. That said, many of the brands that distinguish the rich and the poor, such as mustard or soy sauce, may reflect cultural influence on food choices (cf. Atkin 2016). Also, the differences between the rich and the poor in other dimensions of culture, such as media diet and social attitudes, probably primarily reflect choices rather than opportunities.

4.1.3 Time use and income

As panel (a) of Figure 1 shows, the cultural distance in time use between the rich and the poor does not display a clear trend over time, but the issue is difficult to judge given the lack of data between 1975 and

[45] This fact is particularly interesting given the way Grey Poupon has been marketed. In the 1980s and early 1990s, the so-called "Pardon me" TV advertisements aired: a Rolls-Royce pulls up alongside another Rolls-Royce; a passenger in the back seat of one asks a passenger in the back seat of the other: "Pardon me, would you have any Grey Poupon?"; the other passenger responds, "But of course!" Since then, Grey Poupon has often been referenced in hip-hop lyrics as a symbol of status. For example, FM Static has a song with a verse, "And if I had money, then I'd only wear Sean John / Eat my cereal with Grey Poupon." An analysis by vox.com indicates that almost every year since 1992, at least one hip-hop song has been released referencing Grey Poupon. In 2011, 15 such songs were released.

1998. As we already mentioned, we focus our analysis of time use on respondents who are employed full-time, but even if we look at the full sample, we see no trends in cultural distance (cf: Figure B.19). In contrast to media diet, consumer behavior, and social attitudes, it is less straightforward to identify the individually most informative variables in time use. The problem is that the small sample sizes coupled with a rich potential set of responses (in contrast to the binary responses to questions about media, consumer behavior, and attitudes) drive a wedge between the empirical and the true distribution of responses in a given group.[46]

4.1.4 Social attitudes and income

Information on how people answer questions in the GSS was increasingly predictive of income from the mid-1990s to the mid-2000s but there is no clear trend over the entire time period from 1970s to today. To better understand these patterns, we also consider trends in cultural distance based only on subsets of questions in the GSS. In particular, we separate questions related to: law enforcement; marriage, sex and abortion; life and trust; politics and religion; civil liberties; confidence; and government spending. The Data Appendix reports the complete list of GSS questions included under each theme. Table 3 reports cultural distance based on each of these subsets of GSS questions for the years 1976, 1996, and 2016 (beginning, middle, and end of our sample). We report the full GSS results at the top for reference. The last two columns of Table 3 report the coefficient (and the t-statistic) from an estimate of a linear trend in cultural distance (based on that subset of questions) using all years in the GSS. The topics are listed in order of decreasing estimate of the trend.

Table 3 reveals that the rich and the poor have diverged the most in their attitudes toward law enforcement. At the other extreme, there is evidence of convergence on views about government spending. Table B.1 in the Supplementary Materials reports the ten social attitudes that are single-handedly most predictive of being rich in 1976, 1996, and 2016. The top of these lists is remarkably stable in each year: voting and trusting people are among the three most individually predictive variables in all three years. Figure 4 presents a more systematic analysis of stability of relative predictability over time. We rank order all of the variables based on how informative they are about income in 1976. We use this ranking to determine each variable's vertical position throughout the

[46]For example, in 1965, some poor respondents spent 5 minutes in a day on gardening, some spent 11, but none spent 12; among the rich respondents, some spent 5 minutes on gardening, none spent 11, but some spent 12. It would obviously be mistaken to conclude from this pattern of responses that anyone who spent 12 minutes a day on gardening must be rich while anyone who spent 11 minutes must be poor. This issue is the primary problem tackled by Gentzkow et al. (2017) in their analysis of polarization of political speech. We could follow their approach or other methods for dealing with the finite sample bias here, but for ease of exposition we try to avoid customizing our analyses for particular groups or dimensions of cultural distance.

Table 3: Cultural distance by income over time: social attitudes

	1976	1996	2016	Coefficient	T-statistic
All GSS	74.8%	75.5%	76.8%	0.17	2.58
Law Enforcement	60.0%	55.1%	61.9%	0.31	2.87
Life & Trust	66.3%	68.3%	67.6%	0.12	2.82
Marriage, Sex, Abortion	60.4%	57.9%	63.0%	0.13	1.38
Civil Liberties	65.6%	63.9%	60.7%	0.04	0.44
Politics & Religion	69.4%	65.1%	65.3%	0.00	0.05
Confidence	62.8%	62.1%	59.4%	-0.09	-1.38
Government Spending	65.7%	60.3%	58.1%	-0.14	-2.76

Note: Data source is the GSS. Sample size in each year is 394. Rows 2 to 8 present the results of the machine-learning ensemble method when performed only on the subset of GSS variables in that row. See Data Appendix for the list of GSS questions included in each row. See text and Data Appendix for details on sample construction and implementation of the ensemble. Columns 1 to 3 present the share of correct guesses of respondent's income in the hold-out sample in 1976, 1996, and 2016. The procedure to guess income in the hold-out sample was repeated 500 times. The remaining columns present results from a univariate regression of the share of correct guesses between 1976 and 2016 on a linear year trend, including the estimated coefficient on year (column 4) and the t-statistic associated with that estimated coefficient.

graph. We then color-code the relative informativeness of each variable in each year, with the most informative variables colored dark red, the least informative ones dark blue, and lighter shades of red and blue in between. If the relative informativeness of variables were perfectly stable, each horizontal line would be uniformly colored. Figure 4 reveals that there are substantial changes in the relative importance of specific questions over time, but a small set of highly predictive variables remain highly predictive each year.

4.2 Education

One concern about the classification of households into rich and poor based on income is that we only observe current income, while permanent income might be far more relevant. In this section, we analyze cultural divergence by education, which can be seen as a proxy for permanent income. Examining cultural divergence across education groups is also informative because of the role that education has played in the rise of income inequality (Katz and Murphy 1992). We classify people as less educated if they have at most completed high school, and more educated otherwise.[47,48]

[47] Given this definition and sample size equalization over time, each prediction of education level in a given year is based on 9,674 observations in the MRI, 650 observations in the GSS, 518 observations in the AHTUS, 181,156 observations for baby girls' names and 188,252 for baby boys' names.

[48] The definition of these two groups based on an absolute level of education means that we avoid the issue of potential miscategorization discussed in footnote 41. The fact that we mostly observe similar temporal trends as with income gives

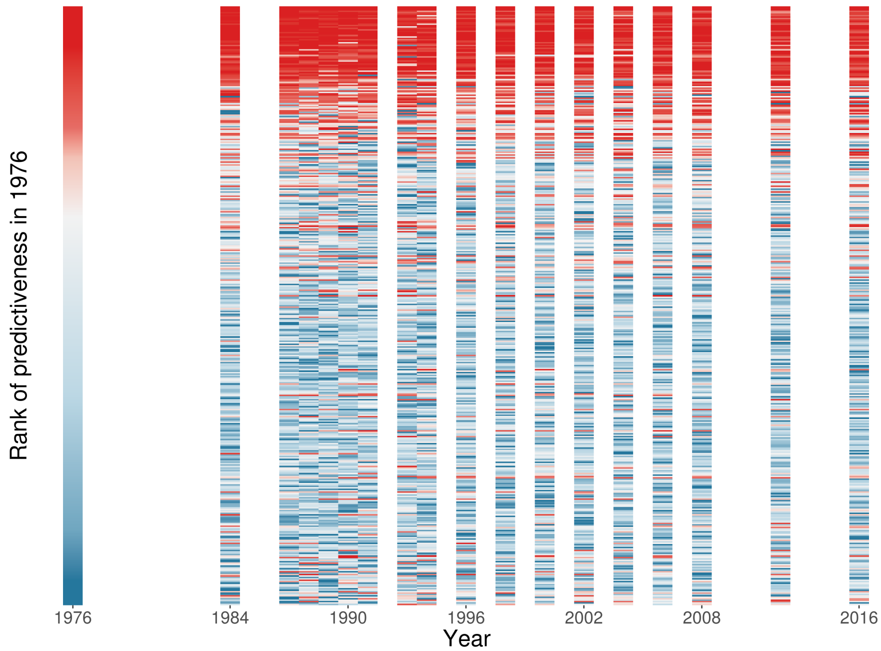

Figure 4: Stability over time of attitudes most indicative of income

Note: Data source is the GSS. Sample size is 394. Variables are ranked from bottom to top throughout the graph by increasing order of correctly guessing income in 1976 based on that variable only. Each variable's relative informativeness in subsequent years is color-coded, with the most informative variables in each year color-coded dark red and the least informative color-coded dark blue, and lighter shades of red and blue in between. See Data Appendix for implementation details.

Panel (b) of Figure 1 summarizes our results. In short, we find broadly similar patterns as with income. Because the results by education mostly match those by income, we do not go through all of them in detail here. More generally, any result reported for some group but not for another in the body of the paper is available in the Supplementary Materials.

4.2.1 Media diet and education

Figure 5 shows that the cultural distance in media diet by education has been stable both overall as well as based on any one of the three sub-categories of media (movies, TV shows, and magazines).

The figure also shows that we can predict education based on magazine readership about as well as we can do with the full set of media diet variables. Table B.2 reports the set of movies, TV shows, and magazines most indicative of higher education. Magazines that are highly indicative of being educated are stable over time, with *Newsweek* and *Time* topping the list throughout the sample.

us confidence that our results on income were not compromised by miscategorization.

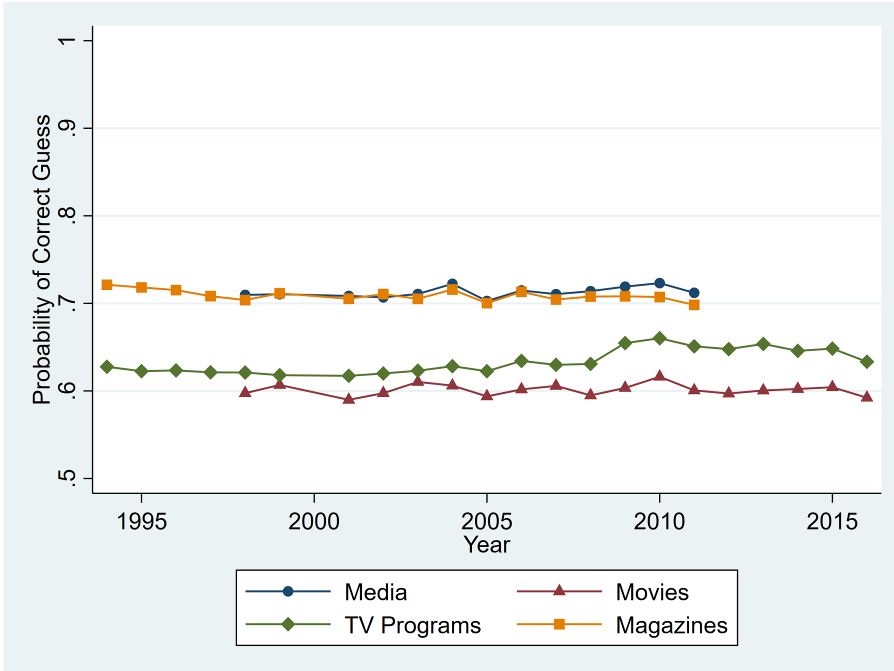

Figure 5: Cultural distance by education over time: media diet

Note: Data source is the MRI. Sample size each year is 9,674. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's education in the hold-out sample each year. The procedure to guess education in the hold-out sample was repeated 25 times, and the share of guesses reported is the average of these 25 iterations.

4.2.2 Consumer behavior and education

As indicated above, our ability to correctly guess one's education based on consumer behavior has remained mostly stable, with maybe a slight increase over time. This is also true when we attempt to predict education using products and brands separately, as shown in Figure 6.

Products and brands most indicative of having attended college are dominated by technological goods (cf: Table B.3). The specific goods reflect waves of technological innovation with the more educated separating themselves from the less educated by being earlier adopters of new technologies. For example, personal computers and internet purchases are more dominant in the early years, while smartphone usage is more relevant in recent years. For brands, using FedEx, owning Windows XP, and owning an iPhone are most informative about someone's education in 1994, 2005, and 2016, respectively. Throughout our sample, consumption of travel-related items also separate the more and the less educated.

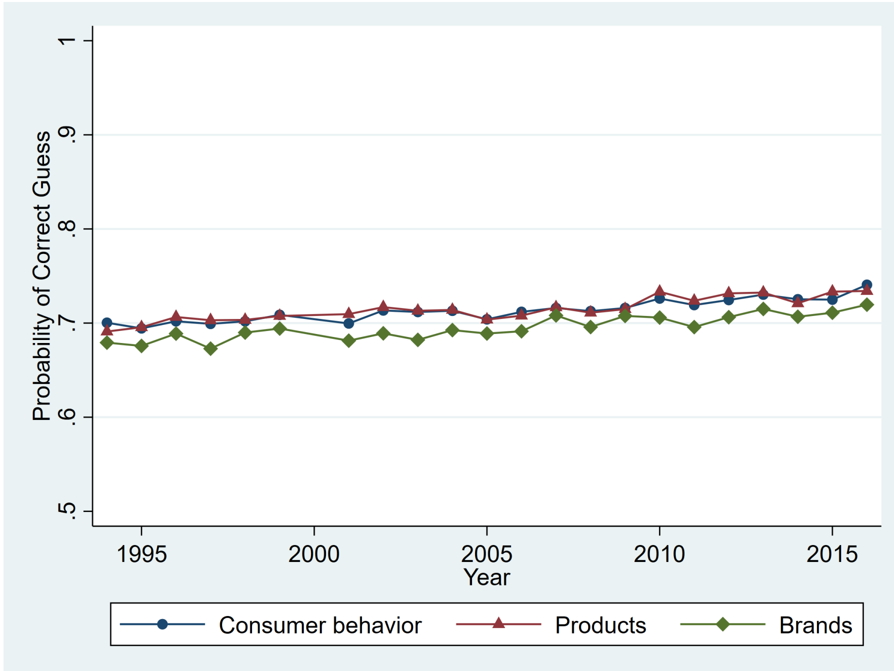

Figure 6: Cultural distance by education over time: consumer behavior

Note: Data source is the MRI. Sample size each year is 9,674. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's education in the hold-out sample each year. The procedure to guess education in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these 5 iterations.

4.2.3 Time use and education

Aguiar and Hurst (2007) find that less-educated individuals have experienced greater increases in leisure compared to their more-educated counterparts. Given this finding, it might seem surprising that we do not see an increase in cultural differences in time use by education level. The mean number of hours per week spent on leisure[49] was indeed roughly the same for the two education groups in 1975, but the less-educated were spending relatively more hours per week on leisure by 2003-2012. That said, the overall distribution of time spent on leisure for the two groups is very similar, both in 1975 and 2003-2012 (cf: Figure B.20). There is substantial heterogeneity in the amount of time spent on leisure within each group, and this heterogeneity is much greater than the mean difference across groups. Moreover, as noted by Aguiar and Hurst (2007), the within-group heterogeneity has also increased over time.[50] Consequently, time use is no more informative about education now than it was in the past.

[49]Leisure is defined as time spent watching TV, socializing, playing sports, reading, engaging in hobbies, sleeping, eating, and engaging in personal care.

[50]Aguiar and Hurst (2007) write: "We also document a significant dispersion of leisure within educational categories... while the growing leisure gap between educational groups is substantial, it is more than matched by the growing within-group dispersion."

Table 4: Culture distance by education over time: social attitudes

	1976	1996	2016	Coefficient	T-statistic
All GSS	71.1%	72.6%	72.0%	0.04	0.78
Law Enforcement	57.2%	54.6%	59.9%	0.11	1.70
Life & Trust	61.6%	60.8%	64.2%	0.06	1.68
Politics & Religion	64.2%	63.6%	62.7%	0.00	-0.07
Civil Liberties	67.7%	66.8%	64.1%	-0.06	-1.00
Marriage, Sex, Abortion	64.4%	59.0%	59.1%	-0.07	-1.15
Confidence	63.2%	63.0%	56.6%	-0.15	-3.30
Government Spending	62.7%	59.3%	56.9%	-0.19	-4.26

Note: Data source is the GSS. Sample size in each year is 650. Rows 2 to 8 present the results of the machine-learning ensemble method when performed only on the subset of GSS variables in that row. See Data Appendix for the list of GSS questions included in each row. See text and Data Appendix for details on sample construction and implementation of the ensemble. Columns 1 to 3 present the share of correct guesses of respondent's education in the hold-out sample in 1976, 1996, and 2016. The procedure to guess education in the hold-out sample was repeated 500 times. The remaining columns present results from a univariate regression of the share of correct guesses between 1976 and 2016 on a linear year trend, including the estimated coefficient on year (column 4) and the t-statistic associated with that estimated coefficient.

4.2.4 Social attitudes and education

We detect no clear time overall trend in the divide between the more and less educated based on social views. Table 4 shows that the more- and less-educated might have somewhat diverged over time in their answers to questions related to law enforcement and life and trust. On the other hand, it has become more difficult over time to tell the more- and less-educated apart based on confidence and views on government spending. Aggregating across all topics, the distance in social attitudes has been broadly constant over time.

4.2.5 Baby names and education

Another dimension of cultural distance we can investigate between education groups is baby naming choices. The use of baby names to study cultural differences between groups was first introduced in economics by Fryer and Levitt (2004), who study trends over time in the names whites and blacks give to their children. There is a much longer tradition in sociology to use patterns in first name usage to study culture and cultural change. As argued in Lieberson (2000), the names parents choose for their children reflect a combination of cultural and social influences, such as mass media and popular culture exposure or religious norms or parental beliefs about what the future of their children will be like. Naming patterns are particularly appealing to study culture in that affordability considerations are totally irrelevant. Lieberson and Bell (1992) was the first paper to document systematic differences in naming patterns by mother education. Among other things, Lieberson and Bell (1992) show difference in baby naming conventions by mother education that appear to reflect differences in the strength of gender norms: less educated mothers are more likely to give their daughters obviously feminine or "frilly" names and more likely to give their sons "strong" names.

Panel (b) of Figure 1 shows that our ability to correctly guess a mother's education based on her baby's name is essentially the same in 2016 as it was in 1989.[51] This is true both for boys' and girls' names. There is a mild increase in the predictability of education based on baby name in the early sample years that is fully compensated by a mild decrease since the mid-2000s. In other words, there is no evidence that more and less educated in California have been subject to diverging cultural and social forces that would result in diverging tastes in what to name their newborn.

4.3 Gender

Much has changed for women over the last half century. Educationally, women turned an educational deficit relative to men into an educational surplus (Goldin et al. 2006). Women's labor force participation rate increased, though it seems to have reached a plateau in the mid-1990s; women's labor market earnings converged towards those of men, but this convergence also appears to have slowed down in the most recent decade (Bertrand 2018). While these well-established trends may a priori suggest shrinking cultural divides between the genders, other forces may have pushed in the other direction. First, women's greater financial independence may have allowed them to also achieve greater cultural independence from their husbands, with the goods, the experiences, and the media they consume becoming more closely aligned with their own personal preferences. Similarly, the decline in marriage and the rise in divorce may have contributed to a cultural divergence between men and women as cultural choices became less likely to take place within the confines of the couple. Moreover, as noted by Edlund and Pande (2002), changes in family structure may have directly affected women's social attitudes, by strengthening their redistributive preferences, support for greater government spending, and overall support for more democratic political platforms.[52]

Panel (c) of Figure 1 summarizes our results by gender.[53] There is no evidence of an increasing cultural gap between men and women based on media diet or social attitudes. Our ability to predict gender based on consumer behavior is nearly perfect in every period, so this is one instance where

[51] While the CDPH starts in 1960, mother's education is only available starting in 1989. See Data Appendix for details.

[52] See also Montgomery and Stuart (1999) and Box-Steffensmeier, Boef, and Lin (2000).

[53] Given sample size equalization over time, each prediction of gender in a given year is based on 15,036 observations in the MRI, 984 observations in the GSS, and 666 observations in the AHTUS.

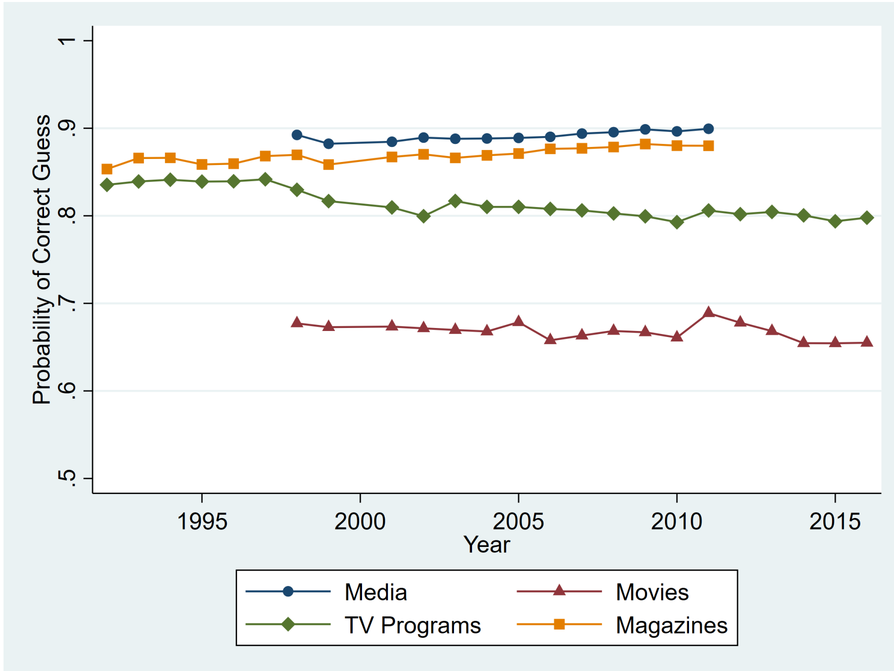

Figure 7: Cultural distance by gender over time: media diet

Note: Data source is the MRI. Sample size each year is 15,036. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's gender in the hold-out sample each year. The procedure to guess gender in the hold-out sample was repeated 25 times, and the share of guesses reported is the average of these 25 iterations.

our approach to measuring cultural distance is ill suited: the presence of a few highly gender-specific goods masks any potential changes in the gender gap in consumption of other goods. Finally, we see that men and women's time use became much more similar from 1965 to the mid-1990s, but there has been no further change in this dimension of cultural distance since the mid-1990s.

4.3.1 Media diet and gender

Figure 7 shows that the cultural distance in media diet by gender has been broadly stable in all the three sub-categories of media, but with hints of divergence in consumption of magazines and convergence in consumption of TV shows. The figure also reveals that magazine readership is most informative about gender, with limited gains in predictive power coming from adding data on TV shows and movies to the ensemble algorithm.

Table B.5 in the Supplementary Materials shows the movies that men and women sort on. Movies most indicative of males tend to fall into the action, thriller, and sci-fi categories while dramas and romantic comedies are most indicative of females. Moreover, Table B.5 reveals that in the early years, the most discriminating movies were those watched primarily by women. In the later years, gender-

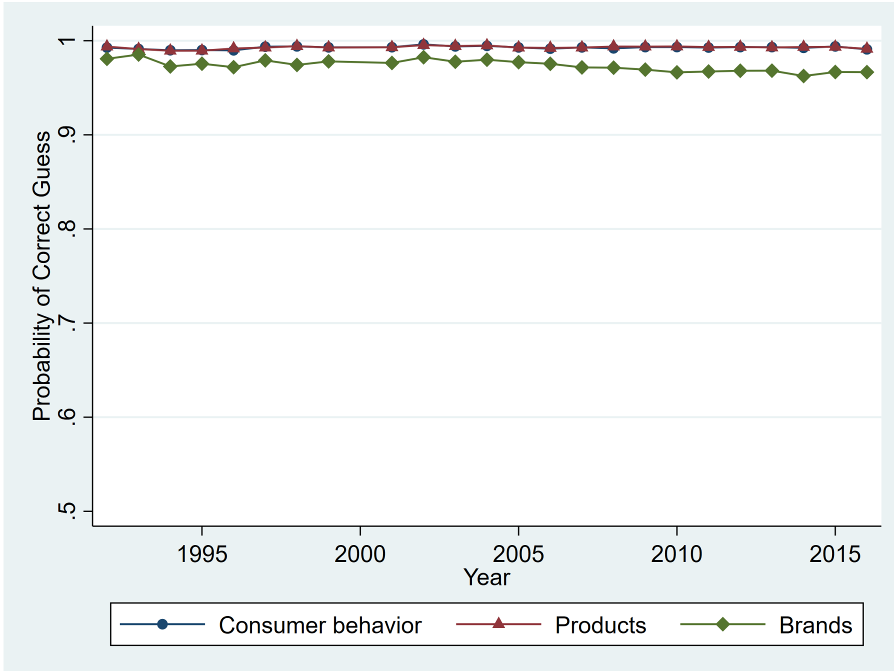

Figure 8: Cultural distance by gender over time: consumer behavior

Note: Data source is the MRI. Sample size each year is 15,036. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's gender in the hold-out sample each year. The procedure to guess gender in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these 5 iterations.

specific movies are mostly those watched primarily by men. Yet, despite these changes, the overall difference in movie-watching by gender has been constant over time.

Table B.5 also shows that, unsurprisingly, fashion and housekeeping magazines are the most distinctive of a female reader, while *Sports Illustrated* is indicative of a male reader. When it comes to television, gender differences are greatest when it comes to watching sports.

4.3.2 Consumer behavior and gender

As indicated above, our approach to measuring cultural distance is not suitable for comparing consumer behavior of men and women since there are some highly gender-specific products that collectively allow us to perfectly predict gender based on consumer behavior in every year.

Table B.6 in the Supplementary Materials lists the products and brands that are individually most indicative of gender. In particular, the consumption of personal care and makeup is so common and so segmented by gender that knowing whether one bought, used, or own these products provides sufficient information to infer gender nearly perfectly.[54]

[54] We could throw out some "overly gender-specific" products from our data and measure the predictability of the other variables but, as we discuss in footnote 46, for ease of exposition we prefer not to customize our approach for particular

4.3.3 Time use and gender

The changes we observe over time in the differences between how men and women spend their time are most striking. The time pattern we document in panel (c) of Figure 1 is reminiscent of the well-known time series on women's labor force participation in the US, with increases in women's labor force participation up until the mid-1990s and a subsequent plateau. However, recall that the sample of men and women we study in the time use data is restricted to the full-time employed, so the observed pattern cannot be due to changes in women's likelihood of being employed. In fact, we observe the same pattern if we predict gender based on shares of non-work time spent on various non-work-related activities (cf: Figure B.21); ways that men and women spend their time outside of work became more similar from 1965 to 1995, but this convergence has stopped since then. This pattern also echoes the fact that attitudes toward gender roles became more egalitarian over time, but this progress stalled in the mid-1990s (Fortin 2015).

4.3.4 Social attitudes and gender

As shown in panel (c) of Figure 1, we observe no divergence, and maybe some weak convergence, in the divide between men and women based on social views and norms.

Table 5 examines cultural distance by gender based on the seven thematic subsets of GSS questions. We observe only one topic for which the genders appear more divided today than in the past, namely marriage, sex, and abortion. There are two modules over which men and women have converged over time: views about life and trust and (but less so) views about government spending. We observe no systematic time trend in cultural distance for the other four themes covered in the GSS survey. The lack of divergence in views about government spending and politics and religion stands in contrast with the work which has argued that women have moved further over time to the political left of men (Edlund and Pande 2002).

4.4 Race

Our motivation for studying racial differences in culture over time is similar to our motivation for studying differences in culture by income groups over time. Just as growing cultural gaps between rich and poor may hinder social mobility, a growing cultural divide between whites and non-whites may

groups or dimensions of culture. In the Nielsen data, where we cannot perfectly predict gender based on consumer behavior, we observe convergence over time in shopping behavior between men and women (see Figure B.9). The main explanation for why we can perfectly predict gender in MRI but not Nielsen is because Nielsen only collects products and brands bought, not used or owned. For example, in 2004, 76% of women in MRI reported using lipstick or lip gloss; in contrast, only 21% of women in Nielsen had bought lipstick.

Table 5: Culture distance by gender over time: social attitudes

	1976	1996	2016	Coefficient	T-statistic
All GSS	73.1%	70.2%	69.2%	-0.10	-2.20
Marriage, Sex, Abortion	57.6%	61.8%	61.4%	0.12	3.04
Law Enforcement	58.3%	62.6%	58.0%	0.04	0.67
Civil Liberties	52.6%	52.8%	53.5%	0.03	0.52
Confidence	56.2%	56.3%	55.8%	0.01	0.17
Politics & Religion	57.5%	54.5%	53.8%	-0.03	-0.78
Government Spending	61.1%	58.5%	59.7%	-0.10	-1.92
Life & Trust	69.1%	64.2%	59.6%	-0.29	-7.42

Note: Data source is the GSS. Sample size in each year is 984. Rows 2 to 8 present the results of the machine-learning ensemble method when performed only on the subset of GSS variables in that row. See Data Appendix for the list of GSS questions included in each row. See text and Data Appendix for details on sample construction and implementation of the ensemble. Columns 1 to 3 present the share of correct guesses of respondent's gender in the hold-out sample in 1976, 1996, and 2016. The procedure to guess gender in the hold-out sample was repeated 500 times. The remaining columns present results from a univariate regression of the share of correct guesses between 1976 and 2016 on a linear year trend, including the estimated coefficient on year (column 4) and the t-statistic associated with that estimated coefficient.

cause continued economic struggles for minority groups in the US.

Given our limited sample sizes, we are unable to conduct our analysis across many racial categories.[55] Instead, we focus on comparison of whites and non-whites. Given this grouping and our procedure for equalizing sample sizes described in Section 3.4, each prediction of race in a given year is based on 4,150 observations in the MRI, 228 observations in the GSS, 298 observations in the AHTUS, 37,458 observations for baby girls' names and 39,402 observations for baby boys' names.

There is a rich literature on cultural differences by race. Fryer and Levitt (2004) discuss some of the prior work on the black-white cultural divide on dimensions such as musical tastes and linguistic patterns. Fryer and Levitt (2004) also highlight anecdotal evidence of racial differences in consumer behavior (e.g., the sharp racial differences in the popularity of different cigarette brands) and media diet (e.g., *Seinfeld*'s huge following among whites and limited appeal among blacks). However, a systematic documentation of changes over time in cultural differences by race, with the exception of Fryer and Levitt (2004)'s study of baby naming choices, is missing. We undertake this analysis in this section, focusing on differences between whites and all non-whites.

Panel (d) of Figure 1 summarizes our results. With the exception of an outlier data point in 1975, there is no apparent trend in our ability to predict one's race based on time use. There is also no

[55]The only exception is baby names where we have enough data to report findings across fine racial and ethnic categories.

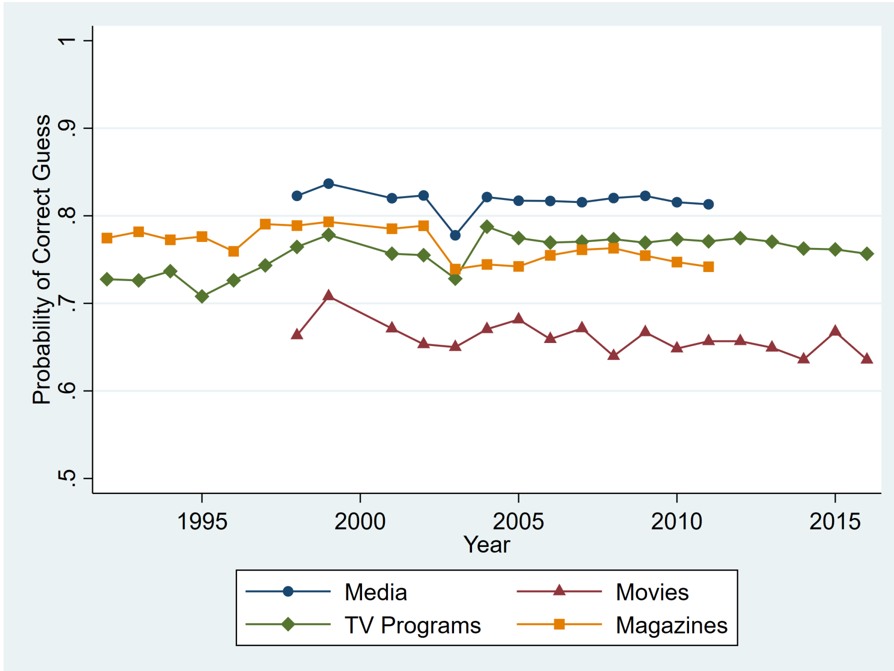

Figure 9: Cultural distance by race over time: media diet

Note: Data source is the MRI. Sample size each year is 4,150. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's race in the hold-out sample each year. The procedure to guess race in the hold-out sample was repeated 25 times, and the share of guesses reported is the average of these 25 iterations.

evidence of growing racial cultural divides based on patterns of media diet. Furthermore, there is no evidence of any growing racial divide, and instead some sign of a shrinking divide, based on social attitudes. While there was some divergence in baby name choices between the late 1960s and the mid-1970s, there has been convergence since the mid-2000s, with no overall time trend. The one dimension where we do observe increasing racial differences is with respect to consumption choices, with much of the increase in the racial gap occurring during the 1990s.

4.4.1 Media diet and race

Panel (a) of Figure 9 examines racial differences in consumption of the three media subcategories. While there are no steep trends, the differences in consumption of magazines and movies seem to have gotten somewhat smaller, while the differences in TV programs have somewhat increased.

Unlike with the prior results (income, education, gender), in the case of race, the combination of all data on media diet substantially increases the predictive power of our model compared to focusing on a subcategory (e.g., magazines).

Looking at specific media products that are individually most predictive of race (cf: Table B.8), a few patterns emerge. TV shows most indicative of being white vary quite a lot over time. Some popular sitcoms clearly have differential appeal across racial lines (e.g., *In Living Color* in 1992 was more popular among non-whites while *The Big Bang Theory* in 2016 was more popular among whites). There is also evidence that whites and non-whites sort into different sports on TV, with American football being more popular among whites and basketball more popular among non-whites. The same three magazines (*Ebony*, *Jet*, *Essence*) are most indicative of race throughout our sample period. The movies that are individually most predictive of race disproportionately cast black actors and/or are targeted at a black audience (such as *The Preacher's Wife* in 1998 or *Big Momma's House 2* in 2007).

4.4.2 Consumer behavior and race

As already indicated, we observe a growing divide between whites and non-whites in terms of consumption choices. The probability of correctly guessing race from consumer behavior grows from roughly 79 percent in 1992 to roughly 87 percent in 2016, with much of the increase taking place during the 1990s and early 2000s. Figure 10 reports on these patterns when we restrict the consumer data to only products or brands. These time series display the same patterns as the full consumer behavior data, with increases in the probability of correctly guessing race concentrated in the first half of the sample period.

The list of products and brands most indicative of being white is reported in Table B.9. Several of the items on the list (e.g. washing machine or air conditioning, or a Range Rover) may be a reflection of the systematic income differences between whites and non-whites. However, many other items are not expensive goods and hence more likely to truly pick up on cultural differences. For example, across years, owning a pet is often highly indicative of being white.[56] Also, whites are distinctive from non-whites in their use of baking products (e.g. bakeware, or baking soda and baking chocolate chips).

4.4.3 Time use and race

We do not see a clear trend in racial differences in time use, though the data show a blip in 1975. We suspect this data point is an accidental outlier, but without a formal take on standard errors (cf: Section 3.5), we cannot take a firm stance. If we restrict our attention to data since 2003, which gives us much larger sample sizes, we see a very stable difference in time use by race over that decade (cf: Figure B.22).

[56] Owning pets might simply reflect living in the suburbs; our data does not allow us to explore this. The MRI provides information on whether the respondent lives in a large core-based statistical area, but does not indicate whether the residence in an urban or a suburban area.

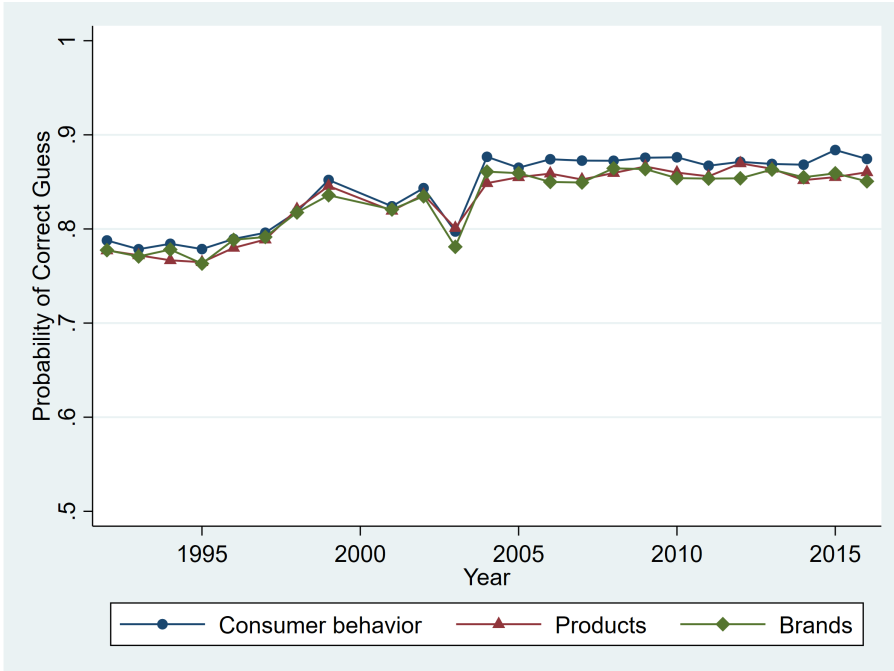

Figure 10: Cultural distance by race over time: consumer behavior

Note: Data source is the MRI. Sample size each year is 4,150. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's race in the hold-out sample each year. The procedure to guess race in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these 5 iterations.

4.4.4 Social attitudes and race

Cultural distance in social attitudes by race shows some sign of slight convergence, with the probability of an accurate guess decreasing from 80 percent to 75 percent over the 40 years of the data. This overall pattern, however, masks some sharp differences in trends across sub-categories of the GSS.

As shown in Table 6, whites and non-whites have grown apart on their views on law enforcement. In 1976, one could correctly predict race based on these views 62 percent of the time but by 2016 that number was up to 70 percent. There has also been divergence on views regarding civil liberties. On the other hand, whites and non-whites have sharply converged in their views on life and trust, politics and religion and government spending. For example, in 1976, one could correctly predict race based on views regarding politics and religion 73 percent of the time but this number was down to 62 percent by 2016.

4.4.5 Baby names and race

Several papers, starting with Lieberson and Mikelson (1995), have studied differences in baby naming conventions as a window into cultural differences by race. Using California birth data spanning the

Table 6: Culture distance by race over time: GSS

	1976	1996	2016	Coefficient	T-statistic
All GSS	80.1%	75.9%	75.4%	-0.11	-2.47
Law Enforcement	62.3%	66.1%	70.3%	0.23	4.28
Civil Liberties	48.9%	53.5%	58.7%	0.21	3.18
Marriage, Sex, Abortion	58.9%	50.2%	55.2%	-0.01	-0.16
Confidence	56.2%	53.8%	54.0%	-0.15	-2.47
Government Spending	74.3%	68.3%	56.9%	-0.36	-5.89
Politics & Religion	72.7%	67.9%	61.8%	-0.25	-6.31
Life & Trust	71.0%	56.7%	54.1%	-0.42	-7.21

Note: Data source is the GSS. Sample size in each year is 228. Rows 2 to 8 present the results of the machine-learning ensemble method when performed only on the subset of GSS variables in that row. See Data Appendix for the list of GSS questions included in each row. See text and Data Appendix for details on sample construction and implementation of the ensemble. Columns 1 to 3 present the share of correct guesses of respondent's race in the hold-out sample in 1976, 1996, and 2016. The procedure to guess race in the hold-out sample was repeated 500 times. The remaining columns present results from a univariate regression of the share of correct guesses between 1976 and 2016 on a linear year trend, including the estimated coefficient on year (column 4) and the t-statistic associated with that estimated coefficient.

years 1960 to 2001, Fryer and Levitt (2004) show that naming conventions among (non-Hispanic) whites and (non-Hispanic) blacks changed dramatically between the late-1960s and mid-1970s when baby names (especially baby girl names) became more racially distinctive, a change Fryer and Levitt attribute to the rise of the Black Power movement; the racial distinctiveness of names then somewhat stabilized in the last quarter of the 20th century. Another recent paper relying on the California birth data to study cultural distance between groups is Abramitzky et al. (2018), who use baby naming choices to proxy for the process of cultural assimilation among immigrants in the US; they find that such name-based assimilation is happening as fast in the last quarter century as it did in the earliest part of the 20th century.

In panel (d) of Figure 1, we plot our ability to correctly guess whether a mother is white based on the name of her baby. Consistent with Fryer and Levitt (2004), we see patterns of increasing racial differences in baby naming choices from the late-1960s to the mid-1970s, with the probability of a correct guess of race based on baby boys' (girls') names reaching a high of about 70 (71) percent in 1976. This probability is then roughly stable until the mid-2000s, after which it experiences a trend down. By 2016, the probability of correctly guessing mother's race based on boys' (girls') names was 65.3 (63.2) percent, about the same as in 1960.

High rates of immigration from Mexico and Central America into California over the last decades

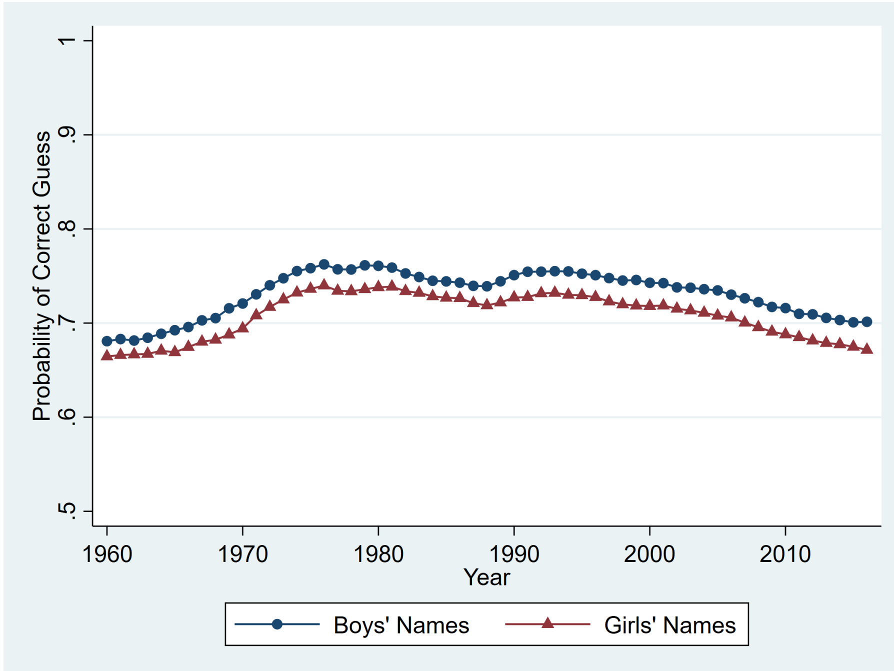

Figure 11: Cultural distance by race and ethnicity over time: baby names

Data source is the CDPH. Sample size each year is 91,320 for boys' names and 89,148 for girls' names. See text and data appendix for details on sample construction and implementation of the Bayesian method. Presented in the figure is share of correct guesses of mother's race and ethnicity (Non-Hispanic white or Other) in the hold-out sample each year. The procedure to guess race and ethnicity in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations.

naturally imply that the composition of the white group has changed a lot over time, with a growing representation of Hispanics in that group. While we have embraced these compositional changes throughout the paper, this is a case where one might reasonably ask how much they account for our findings. Therefore, in Figure 11, we compare non-Hispanic whites with the rest of the population. Qualitatively, the results are the same.

Figures B.23 and B.24 in the Supplementary Materials present results by finer race/ethnicity pairs for boys' names and girls' names respectively. The overall conclusion of no growing divides in baby naming conventions since the mid-1970s holds across these finer groups as well.

4.5 Political ideology

Of all the cultural divides under study in this paper, the divide that separates liberals from conservatives has received the most prior attention. A large literature documents the rising polarization of Democrats and Republicans in Congress,[57] and the discussions of political polarization is on the rise (Gentzkow 2016), but the literature to date has been far less conclusive on whether Republicans

[57] Recent work has also shown growing polarization on moral values between Democrat and Republican candidates in recent U.S. presidential elections (Enke 2018).

and Democrats in the US population overall have been growing apart. Most of the academic work has focused on differences in social attitudes captured in the GSS or the ANES, with some studies concluding that polarization is on the rise (e.g., Abramowitz and Saunders 2008; Draca and Schwarz 2018;) and others rejecting this conclusion (e.g., Fiorina and Abrams 2008; Glaeser and Ward 2006). Much of the disagreement between these studies is ultimately driven by differences in how polarization is measured. We contribute to this literature by examining what our definition of cultural distance implies for changes in the attitude-differences between liberals and conservatives. Furthermore, we extend the literature by also analyzing the divide between liberals and conservatives in other aspects of culture, namely media diet and consumer behavior. (Political affiliation is not available in our time use data.)

In the GSS, respondents are categorized as (a) extremely liberal, (b) liberal, (c) slightly liberal, (d) moderate, (e) slightly conservative, (f) conservative, or (g) extremely conservative; we define respondents as liberal if they identify themselves as (a), (b), or (c) and conservative if they identify themselves as (e), (f), or (g). In the MRI, respondents are categorized as (a) very liberal, (b) somewhat liberal, (c) middle of the road, (d) somewhat conservative, or (e) very conservative; we define respondents as liberal if they identify themselves as (a) or (b) and conservative if they identify themselves as (d) or (e). Given our sample-size-equalization procedure, this yields 4,864 observations per year in the MRI and 552 observations per year in the GSS. We discovered a sharp change in the MRI data in the number of missing observations on ideology after 2009, so based on our principle of keeping the quality of the data constant over time, we analyze media diet and consumer behavior by ideology only in the 1994-2009 period.

Panel (e) of Figure 1 summarizes our results. Our ability to predict someone's political ideology based on patterns of media diet or consumer behavior is essentially constant across years. On the other hand, we document a growing divide between liberals and conservatives based on their stated social attitudes.

4.5.1 Media diet and political ideology

Figure 12 shows that the stability of the cultural distance between liberals and conservatives based on the full media bundle broadly extends to the three separate media sub-components, with one exception: liberals and conservatives converged in terms of the TV shows they watch during the 1990s; since then, the difference in their TV habits has been mostly stable.

We also observe that TV shows are not only more predictive of ideology than movies or magazines, they are more predictive than all three subcomponents put together. While this might seem puzzling at

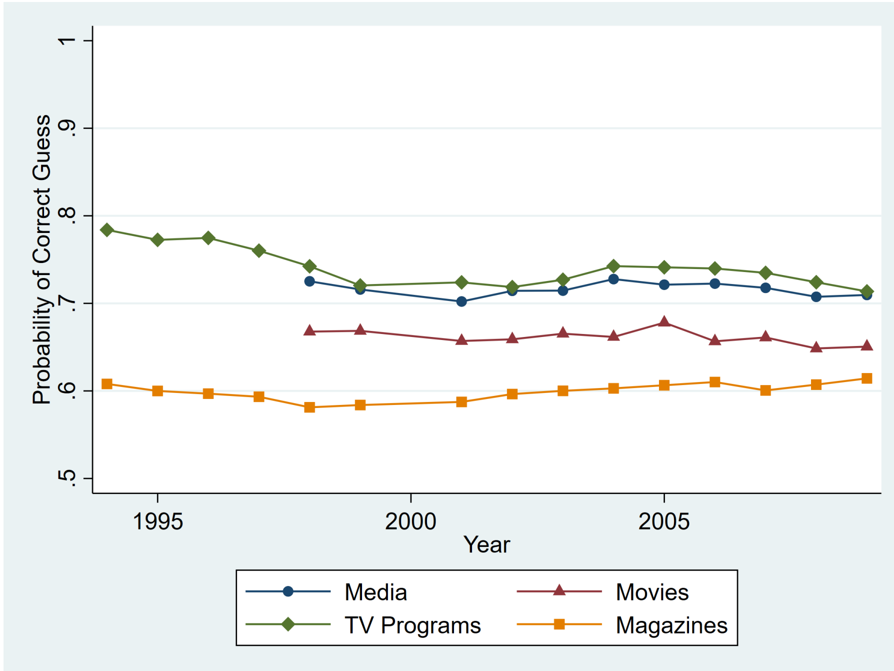

Figure 12: Cultural distance by political ideology over time: media diet

Note: Data source is the MRI. Sample size each year is 4,864. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's political ideology in the hold-out sample each year. The procedure to guess political ideology in the hold-out sample was repeated 25 times, and the share of guesses reported is the average of these 25 iterations.

first glance, it simply reflects the fact that our machine learning algorithm is not fully non-parametric, so inclusion of additional, less predictive variables can decrease predictive power of the estimated model.[58]

Table B.11 in the Supplementary Materials lists the TV programs, movies, and magazines that are single-handedly most indicative of political ideology. There is substantial variation over time in the list of most predictive TV shows. The contrast between the list of top shows in 2001 and 2009 is particularly interesting. TV programs most indicative of ideology in 2001 are college football games, *Academy Awards*, *Will & Grace*, and *Ally McBeal.* In contrast, the three TV programs most indicative of ideology by 2009 are all Fox news programs: *The O'Reilly Factor*, *Fox and Friends*, and *Hannity and Colmes*, with conservatives disproportionately watching all three. This contrast provides a good reminder of a key (potentially undesirable) feature of our measure of cultural distance, which is that the measure does not take any stance on either which elements of cultural behavior are important nor which elements are close to one another. One might argue that the distance between liberals and

[58]This is especially the case for random forests. If instead of ensemble (which includes random forest), we predict ideology using an elastic net or a regression tree, the predictive power is greater when we use all media variables rather than TV shows alone.

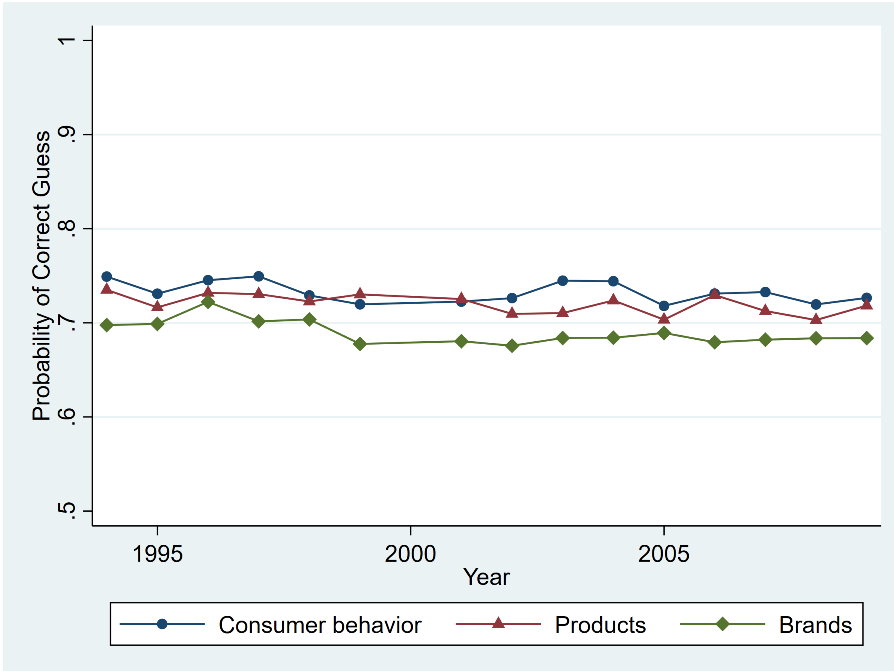

Figure 13: Cultural distance by political ideology over time: consumer behavior

Note: Data source is the MRI. Sample size each year is 4,864. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's political ideology in the hold-out sample each year. The procedure to guess political ideology in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these 5 iterations.

conservative is greater in 2009 than in 2001 in the sense that *The O'Reilly Factor* is more different from other TV shows than a college football game is. Moreover, even if one does not take a stance on the "distance" between *The O'Reilly Factor* and other programs, we might worry about differences in where people get their news much more than about differences in where people get their non-news entertainment.

4.5.2 Consumer behavior and political idiology

As indicated in Figure 1, our ability to correctly predict political ideology based on the basket of goods and brands consumed hovers around the low 70 percent range throughout the sample period. Figure 13 shows this pattern is broadly similar when we restrict the consumer behavior information to either products or brands.

The list of products individually most distinctive of ideology is interesting (cf: Table B.12). In all years but especially the first two listed in the table, liberals distinguish themselves from conservatives by drinking alcohol. Conservatives, on the other hand, are much more likely to engage in fishing. Another notable fact is that not using tampons is highly indicative of being conservative in the early

Table 7: Culture distance by political ideology over time: GSS

	1976	1996	2016	Coefficient	T-statistic
All GSS	68.5%	74.2%	81.0%	0.41	6.28
Religion	50.9%	59.3%	62.5%	0.32	5.53
Marriage, Sex, Abortion	63.0%	67.9%	75.2%	0.38	4.61
Confidence	59.5%	62.7%	68.3%	0.27	4.34
Government Spending	64.6%	65.6%	71.4%	0.24	3.96
Law Enforcement	64.7%	63.6%	67.3%	0.18	3.61
Life & Trust	52.7%	50.9%	58.3%	0.14	2.55
Civil Liberties	60.2%	57.7%	57.7%	0.03	0.36

Note: Data source is the GSS. Sample size in each year is 552. Rows 2 to 8 present the results of the machine-learning ensemble method when performed only on the subset of GSS variables in that row. See Data Appendix for the list of GSS questions included in each row. See text and Data Appendix for details on sample construction and implementation of the ensemble. Columns 1 to 3 present the share of correct guesses of respondent's political ideology in the hold-out sample in 1976, 1996, and 2016. The procedure to guess political ideology in the hold-out sample was repeated 500 times. The remaining columns present results from a univariate regression of the share of correct guesses between 1976 and 2016 on a linear year trend, including the estimated coefficient on year (column 4) and the t-statistic associated with that estimated coefficient.

years.[59] Ideology-specific brands are mostly food, with conservatives consuming Jell-O and CoolWhip and eating at Arby's while liberals go to Starbucks.[60]

4.5.3 Social attitudes and political ideology

Differences in social attitudes between liberals and conservatives is the dimension of cultural distance that has received the most prior attention. Based on our measure, we find that liberals and conservatives are more different today in their social attitudes than they have ever been in the last 40 years. Moreover, this divergence is not a recent phenomenon. Table 7 also shows that liberals and conservatives have diverged in their views on almost every one of the seven thematic subsets in the GSS.[61] In particular, while we detect no time trend in our ability to tell liberals and conservatives apart based on views towards civil liberties, we see cultural divergence in the remaining six dimensions. Divergence has been greatest in views on religion and view on marriage, sex and abortion.

[59] This is not about different shares of males and females for liberals and conservatives. Even within the sample of women, using tampons is highly indicative of being a liberal in the early years.

[60] Amusingly, in 2009 the most predictive TV show primarily watched by liberals is *The Daily Show with Jon Stewart*. In this show, making fun of *Arby's* (and its customers) is perhaps the most commonly repeated gag.

[61] Recall, as we mentioned in Footnote 20, that in analyzing ideological differences in social attitudes, we drop questions about the respondent's political affiliation and questions about how or whether the respondent voted; accordingly, the questions in the theme "Politics & Religion" are here replaced by just "Religion". This theme includes only questions about respondent's religion, attendance of religious services.

5 Conclusion

We study temporal trends in cultural distances as reflected in media diet, consumption choices, time use, social views, and baby naming choices between groups in the US defined by income, education, gender, race, and political ideology. We use a machine learning approach to measure cultural distance, an approach that is well suited to this particular application given the rich set of features and traits that define someone's culture. The main take-away of our analysis is that, except for a few noteworthy exceptions, cultural distances have remained broadly constant over time. This take-away runs against the popular narrative of the US becoming an increasingly divided society.

There are, however, a few important caveats to our main finding. First, our approach does not take a stance on what features of culture matter for healthy and productive interactions between groups in society. As we discussed earlier, social frictions may be more affected by whether we get our news from similar sources than by whether we watch different sitcoms. That said, perhaps people primarily connect by talking about sitcoms and sports rather than about the news. Nothing in our data provides guidance on which aspects of culture are most important for our ability to "get along."

A second limitation of our approach is that it can only analyze cultural distances one dimension at a time, since no single dataset encompasses data on media diet, consumption behavior, time use, and social attitudes. It is possible that there have been changes in the correlation between these components of culture over time and that an analysis that draws on an integrated dataset would come to different conclusions.

Finally, our assessment of the extent of cultural divides within the US has been focused on looking at pre-specified groups (rich vs. poor, more vs. less educated, man vs. woman, white vs. non-white, liberal vs. conservative). In future work, we plan to explore trends in polarization across social groups that are defined endogenously, based on their distinct cultural traits.

References

Abramitzky, Ran, Leah Platt Boustan and Katherine Eriksson. 2018. Cultural assimilation during two ages of mass migration. Working paper.

Abramowitz, Alan I. and Kyle L. Saunders. 2008. Is polarization a myth? *The Journal of Politics.* 70(2): 542-555.

Aguiar, Mark and Erik Hurst. 2007. Measuring trends in leisure: The allocation of time over five decades. *Quarterly Journal of Economics.* 122(3): 969-1006.

Aguiar, Mark and Erik Hurst. 2009. A summary of trends in American time allocation: 1965-2005. *Social Indicators Research.* 93(1): 57-64.

Alesina, Alberto, Reza Baqir, and William Easterly. 1999. Public goods and ethnic divisions. *Quarterly Journal of Economics.* 114(4): 1234-84.

Alesina, Alberto and Eliana La Ferrara. 2000. Participation in heterogeneous communities. *Quarterly Journal of Economics.* 115(3): 847-904.

Alesina, Alberto and Eliana La Ferrara. 2005. Ethnic diversity and economic performance. *Journal of Economic Literature.* 43(3): 762-800.

Alesina, Alberto, Guido Tabellini, and Francesco Trebbi. 2017. Is Europe an optimal political area? *Brookings Papers on Economic Activity.*

Autor, David H., Lawrence F. Katz and Melissa S. Kearney. 2008. Trends in U.S. wage inequality: Revising the revisionists. *The Review of Economics and Statistics.* 90(2): 300-323.

Atkin, David. 2016. The caloric costs of culture: Evidence from Indian migrants. *The American Economic Review.* 106(4): 1144-1181.

Bertrand, Marianne. 2018. Coase Lecture: the glass ceiling. *Economica.* 85(338): 205-231.

Boneva, Teodora and Christopher Rauh. 2017. Socio-economic gaps in university enrollment: The role of perceived pecuniary and non-pecuniary returns. Working paper.

Bourdieu, Pierre. 1984 [1979]. Distinction: A Social Critique of the Judgment of Taste. Translated by Richard Nice. Cambridge, MA: Harvard University Press.

Box-Steffensmeier, Janet, Suzanna De Boef, and Tse-Min Lin. (2004). The dynamics of the partisan gender gap. *American Political Science Review.* 98(3): 515-528.

Concordance of Activity Codes. American Heritage Time Use Survey Codebook. University of Oxford Center for Time Use Research.

Desmet, Klaus, Ignacio Ortuno-Ortin, and Romain Wacziarg. 2017. Culture, ethnicity, and diversity. *The American Economic Review.* 107(9): 2479-2513.

Desmet, Klaus and Romain Wacziarg. 2019. The cultural divide. Working paper.

Draca, Mirko and Carlo Schwarz. 2018. How polarized are citizens? Measuring ideology from the ground-up. Working paper.

Edlund, Lena and Rohini Pande. 2002. Why have women become left-wing? The political gender gap and the decline in marriage. *The Quarterly Journal of Economics.* 117(3): 917-961.

Enke, Benjamin. 2018. Moral values and voting: Trump and beyond. Working paper.

Fiorina, Morris P. and Samuel J. Abrams. 2008. Political polarization in the American public. *Annual Review of Political Science.* 11: 563-588.

Fortin, Nicole. 2015. Gender role attitudes and women's labor market participation: Opting-out, AIDS, and the persistent appeal of housewifery. *Annals of Economics and Statistics.* 117/118: 379-401.

Fryer, Roland G. Jr. and Steven D. Levitt. 2004. The causes and consequences of distinctively black names. *The Quarterly Journal of Economics.* 119(3): 767-805.

Gentzkow, Matthew. 2016. Polarization in 2016. Toulouse Network for Information Technology Whitepaper.

Gentzkow et al. 2017. Measuring polarization in high-dimensional data: Method and application to Congressional speech. Working paper.

Glaeser, Edward L. and Bryce A. Ward. 2006. Myths and realities of American political geography. *Journal of Economic Perspectives.* 20 (2): 119-144.

Goldin, Claudia, Lawrence F. Katz and Ilyana Kuziemko. 2006. The homecoming of American college women: The reversal of the college gender gap. *The Journal of Economic Perspectives.* 20(4): 133-156.

Jaravel, Xavier. 2017. The unequal gains from product innovations: Evidence from the US retail sector. Working paper.

Katz, Lawrence F. and Kevin M. Murphy. 1992. Changes in relative wages, 1963-1987: Supply and demand factors. *The Quarterly Journal of Economics.* 107(1): 35-78.

Kaufman, Karen. 2002. Culture wars, secular realignment, and the gender gap in party identification. *Political Behavior.* 24(3): 283-307.

Lieberson, Stanley. 2000. A Matter of Taste. New Haven, CT: Yale University Press.

Lieberson, Stanley and Kelly S. Mikelson. 1995. Distinctive African American names: an experimental, historical, and linguistic analysis of innovation. American Sociological Review, 60 (6): 928-946.

Lieberson, Stanley and Eleanor O. Bell. 1992. Children's first names: an empirical study of social taste. *American Journal of Sociology.* 98 (3): 511-554.

Meyer, Bruce D. and James X. Sullivan. 2017. Consumption and income inequality in the U.S. since the 1960s. NBER working paper 23655.

Montalvo, José, G. and Marta Reynal-Querol. 2005. Ethnic polarization, potential conflict, and Civil Wars. *American Economic Review*. 95(3): 796-816.

Montgomery, Robert and Charles Stuart. 1999. Sex and fiscal desire. Working paper.

Mossel, Elchanan, Allan Sly, and Omer Tamuz. 2014. Asymptotic learning on Bayesian social networks. *Probability Theory and Related Fields*. 158(1-2): 127-157.

Mullainathan, Sendhil and Jann Spiess. 2017. Machine learning: An applied econometric approach. *Journal of Economic Perspectives*. 31(2): 87-106.

Murray, Charles. 2012. *Coming Apart: The State of White America, 1960-2010*. New York: Crown Forum/Random House.

Neiman, Brent and Joseph Vavra. 2018. The rise in household spending concentration. Working paper.

Oliver, Eric, Thomas Wood, and Alexandra Bass. 2016. *Liberellas* versus *Konservatives*: Social status, ideology, and birth names in the United States. *Political Behavior*. 38:55–81.

Politis, Dimitris N., Joseph P. Romano and Michael Wolf. 1999. *Subsampling*. New York: Springer series in statistics.

Waldfogel, Joel. 2003. Preference externalities: An empirical study of who benefits whom in differentiated product markets. *The RAND Journal of Economics*. 34(3): 557–568.

Wolfram, Walt, and Erik Thomas. 2002. *The Development of African American English*. Oxford, UK: Blackwell Publishers.

Zimmerman, Seth. 2017. Making the One Percent: The Role of Elite Universities and Elite Peers. Working paper.

A Data Appendix

A.1 Sample Construction

General Social Survey

We use the General Social Survey (GSS) to measure cultural distance for social attitudes. We use 18 interspersed years from 1976 to 2016 (1976, 1984, 1987, 1988, 1989, 1990, 1991, 1993, 1994, 1996, 1998, 2000, 2002, 2004, 2006, 2008, 2012, and 2016). While the GSS is available from 1972 to 2016 (annually from 1972 to 1991, 1993, and bi-annually from 1994 to 2016), we restrict the analysis to the 18 interspersed years above as the preferred trade-off between maximizing the number of years (and time coverage) and maximizing the number of common questions asked in each year.

We use 83 questions from the GSS. We define a variable as a dummy variable for each response to a question. For example, the question "Are you happy?" has five possible responses: 1) very happy, 2) pretty happy, 3) not too happy, 4) don't know, and 5) no answer. We define a variable "Are you happy - very happy" as a dummy variable that equals 1 for response 1) to the question and 0 otherwise. We do the same for the other responses. We organize the full list of variables in seven themes:

Civil liberties: Allow atheists to teach; allow communists to teach; allow militarists to teach; allow racists to teach; allow homosexuals to teach; allow atheists' books in library; allow communists' books in library; allow militarists' books in library; allow racists' books in library; allow homosexuals' books in library; allow atheists to speak; allow communists to speak; allow militarists to speak; allow racists to speak; allow homosexuals to speak.

Confidence: confidence in military; confidence in business; confidence in organized religion; confidence in education; confidence in executive branch; confidence in financial institutions; confidence in US Supreme Court; confidence in organized labor; confidence in Congress; confidence in medicine; confidence in the press; confidence in scientific community; confidence in TV.

Government spending: foreign aid; military & defense; solving problems of large cities; halting crime rate; dealing with drug addiction; education; environment; welfare; health care; affirmative action; space exploration programs; income tax too high/adequate/too low.[62]

Law enforcement and gun control: courts dealing with criminals; should marijuana be legal; approve of police striking citizens if: citizen said vulgar things; citizen attacked policemen with fists; cit-

[62]For the eleven first questions in the government spending module, the GSS has a "split ballot" design since 1984, where one-third of the respondents were asked the original version of the question and another one-third of the respondents were asked a slightly differently worded version of the question. For these questions, we merge the two questions and treat them as the same despite the slight change in wording. For example, for government spending on education, the original question was worded as: "We are faced with many problems in this country, none of which can be solved easily or inexpensively. I'm going to name some of these problems, and for each one I'd like you to name some of these problems, and for each one I'd like you to tell me whether you think we're spending too much money on it, too little money, or about the right amount. Are we spending too much, too little, or about the right amount on improving the nation's education system?" The altered version use the word "education" instead of "the nation's education system."

izen attempted to escape custody; citizen questioned as murder suspect; ever approve of police striking citizen; favor/oppose death penalty for murder; favor/oppose gun permits; have gun/pistol/rifle/shotgun at home.

Life, life outlook, and trust: should aged live with their children; afraid to walk at night in neighborhood; opinion of how people get ahead; general happiness; condition of health; people helpful or looking out for selves; any opposite race in neighborhood; if rich, continue or stop working; can people be trusted.

Marriage, sex, and abortion: approve of legal abortion if: strong change of serious defect; woman's health seriously endangered; married – wants no more children; low income – cannot afford more children; pregnant as result of rape; not married; happiness of marriage; homosexual sex relations; feelings about porn laws; extramarital sex; seen X-rated movie in the last year.

Politics and religion: political party affiliation; liberal vs. conservative; voted for D, R, I or other presidential candidate; voted in the election; how often attend religious services; religion & denomination; how fundamentalist; belief in life after death.[63,64,65,66,67]

We use all respondents of ages 20 to 64.

We use the following demographic variables in the GSS for our analysis: income, education,

[63] When predicting political ideology, we drop variables related to the following four questions: Liberal vs. conservative; Political party affiliation; Voted for D, R, I or other presidential candidate; Voted in the election.

[64] For the question voted for D, R, I or other presidential candidate, we use the following questions in the GSS: PRES72, PRES80, PRES84, PRES88, PRES92, PRES96, PRES00, PRES04, PRES08, PRES12. Each of these questions asked which presidential candidate the respondent voted for in the election in year 19XX or 20XX. These questions were asked only for the four years after the election. For example, VOTE88 exists in the GSS for years 1989-1992 only.

[65] For the question voted in the election, we use the following questions in the GSS: VOTE72, VOTE80, VOTE84, VOTE88, VOTE92, VOTE96, VOTE00, VOTE04, VOTE08, VOTE12. Like the PRESXX questions, each of these variables asked whether they voted in the election in year 19XX or 20XX, and were asked only for the four years after the election.

[66] We derived our presidential vote variable (with this following values: voted for D candidate, voted for R candidate, voted for I or other candidate, didn't vote, don't know, and no answer) from the question voted for D, R, I or other presidential candidate and the question voted in the election in the following way: 1. Respondents who responded "didn't vote" in either the vote question or the presidential vote question are assigned "didn't vote;" 2. Respondents who responded "don't know" in either the vote question or the presidential vote question are assigned "don't know;" 3. Respondents who responded "refused" or "no answer" in either the vote question or the presidential vote question are assigned "no answer;" 4. Respondents who responded "not eligible" to the vote question are assigned missing code and we impute their responses later.

[67] For the religion and denomination questions, we merged the religion question and the Christian denomination question such that we have a response for each Christian denomination and for each non-Christian religion.

gender, race,[68,69], political ideology,[70] urbanicity,[71] and age. All demographic variables are available in all years that we analyze.

The income variable available in the GSS is family income, and it is reported in income brackets. The income brackets change across years.[72]

To implement the ensemble and Bayesian algorithms, we equalize sample size across years by demographic. For each binary demographic, we find the smallest sample size among all demographic group-year combinations. The algorithms first read in the entire cleaned dataset and then randomly draw a balanced sample with the two demographic groups having the same number of observations equal to the aforementioned smallest sample size. The sample size that we use for each demographic is listed in A.1. Column "Demographics" shows the demographic for the rows, column "Groups" shows all available demographic groups for the given demographic, column "Smallest Group – Year" shows the demographic group-year combination with the smallest sample size, column "Smallest Group Size" shows the corresponding smallest sample size, and column "Balanced Sample Size" shows the size of the balanced sample that the algorithms use.

In the GSS, for most questions, the data is missing for approximately one-third of the sample. This is because the "sociopolitical attitude and behavior questions are administered using a "split-ballot" design - in which items are assigned to two of three ballots, each of which is answered by a random two-thirds of most GSS samples" (Smith et al. 2014).[73] We impute the missing data as follows. In each demographic-year, among respondents with non-missing values for each question, we compute the distribution of answers (for example, 40% answer "Republican," 30% answer "Independent," and 40% answer "Democrat" to the party affiliation question). Then, for each demographic-year, we use the distribution of answers among respondents with non-missing values to randomly impute the response for respondents with missing responses in the same proportions.[74]After imputing for missing values,

[68]In the GSS, we use the question RACE for our race specification. The responses to this question are "white", "black", or "other." This question is available for all years of the GSS.

[69]In the GSS, there is a question HISPANIC, which identifies whether or not the respondent is Hispanic and has values for detailed country of origin in the Hispanic world (for example, Mexican, Puerto Rican, Cuban, etc.). This variable is available since year 2000. We do not use this variable for our race specification.

[70]For political ideology, the GSS question that we use is POLVIEW, which has the following responses: extremely liberal; liberal; slightly liberal; moderate; slightly conservative; conservative; and extremely conservative. We define political ideology as equal to one if the responses are extremely liberal, liberal, or slightly liberal. We define political ideology as equal to zero if the responses are slightly conservative, conservative, and extremely conservative. We drop observations with the response moderate.

[71]For urbanicity, the GSS question that we use is SRCBELT, which has the following responses: 12 largest SMSA's; 13-100 SMSA's; suburb of 12 largest SMSA's; suburb of 13-100 largest SMSA's; other urban; and other rural. We define urbanicity as equal to one for all responses other than "other rural", zero otherwise.

[72]There are 12 brackets for 1976, 17 brackets for the period 1982 to 1985, 20 brackets for the period 1986 to 1990, 21 brackets for the period 1991 to 1996, 23 brackets for the period 1998 to 2004, and 25 brackets for the period 2006 to 2012, and 26 brackets for 2016.

[73]Smith, Tom W, Peter Marsden, Michael Hout, and Jibum Kim. 2014. *General Social Surveys: Cumulative Codebook.*

[74]We note that the above method of imputation uses only the marginal distribution (the distribution of each variable

Table A.1: Sample Size for GSS

Demographic	**Groups**	**Smallest Group – Year**	**Smallest Group Size**	**Balanced Sample Size**
Income	Top Quartile Bottom Quartile	Top Quartile 1990	197	394
Education	Some College or More High School or Less	Some College or More 1976	325	650
Gender	Male Female	Male 1990	492	984
Race	White Non-White	Non-White 1976	114	228
Political Ideology	Liberal Conservative	Liberal 2004	276	552
Urbanicity	Urban Rural	Rural 1988	115	230
Age	40 Years or Older Less Than 40 Years Old	40 Years or Older 1990	479	958

we reshaped the data into dummy variables for each question-response.

American Heritage Time Use Survey (AHTUS)

We use the American Heritage Time Use Survey (AHTUS) to measure cultural distance for time use. We use all available years: 1965, 1975, 1985, 1993, 1995, 1998 and annually from 2003 to 2012.

We equalize the set of activities across years using an activities crosswalk that is based on the official documentation published by the University of Oxford Center for Time Use Research. After equalizing the set of activities across years, we use all of the 73 available activities, as well as the 8 aggregates of activities from Aguiar and Hurst (2009).[75] We define a variable as minutes spent on the activity per day. The full list of variables is: general or other personal care; sleep; naps and rest; wash, dress, personal care; personal medical care; meals at work; other meals and snack; main paid work (not at home); paid work at home; second job, other paid work; work breaks; other time at workplace; time looking for work; regular schooling, education; homework; short course or training; occasional lectures and other education or training; food preparation, cooking; set table, wash/put away dishes; cleaning;

X by demographic group) and not the joint distribution (the joint distribution of variable X, Y, and Z by demographic group).

[75]The 8 aggregates of activities are: market work; home maintenance; obtain goods and services; other home production; non-market work; child care; leisure; and other.

laundry, ironing, clothing repair; home repairs, maintain vehicle; other domestic work; purchase routine goods; purchase consumer durables; purchase personal services; purchase medical services; purchase repair, laundry services'; financial/government services; purchase other services; general care of older children; medical care of children; play with children; supervise/help with homework; read to/with, talk with children; other child care; adult care; general voluntary acts; political and civic activity; worship and religious acts; general out-of-home leisure; attend sporting event; go to cinema; theater, concert, opera; museums, exhibitions; café, bar, restaurant; parties or receptions; sports and exercise; walking; physical activity/sports with child; hunting, fishing, boating, hiking; gardening; pet care, walk dogs; receive or visit friends; other in-home social, games; artistic activity; crafts; hobbies; relax, think, do nothing; read books, periodicals, newspapers; listen to music; listen to radio; watch television, video; writing by hand; conversation, phone, texting; and use computer.[76]

Travel: travel to or from work; travel related to education; travel related to consumption; travel related to child care; travel related to volunteering and worship; other travel.

We use full-time employed respondents of ages 20 to 64.

We use the following demographic variables in the AHTUS for our analysis: income, education, gender, race,[77,78] urbanicity, and age. Not all variables are available in all years and we do not use all demographic variables in all years. While the income variable is available in all years, we exclude 1985, 1993, and 1995 from our analysis using income, because the available income data are too coarse (only approximate income quartiles are available in those years). The race variable is not available in 1985.

The income variable available in the AHTUS is family income, and it is available in income brackets. The income brackets change across years.[79]

To implement the ensemble and Bayesian algorithms, we equalize sample size across years by demographic. For each binary demographic, we find the smallest sample size among all demographic group-year combinations. The algorithms first read in the entire cleaned dataset and then randomly draw a balanced sample with the two demographic groups having the same number of observations

[76]The variable "use computer" first appears in the data in 1985. We assign 0 minutes for "use computer" for all observations prior to 1985.

[77]In AHTUS, we use the variable ETHNIC2 for our race specification. The values of this variable are "white", "black", "some other race", "missing or dirty", or "not applicable." We drop observations that have the values "missing or dirty" or "not applicable." We define the binary race variable as equal to 1 if the value is "white" and 0 if the value is "black" or "some other race." This variable is available for all years of AHTUS.

[78]In AHTUS, there is a variable called HISP which identifies respondent's Hispanic origin. The variable has values "Yes" or "No" for respondent's Hispanic origin. This variable is available since year 1995. We do not use this variable for our race specification.

[79]There are 10 brackets for 1965, 18 brackets for 1975, 7 brackets for 1998, and 16 brackets for the period 2003 to 2012.

equal to the aforementioned smallest sample size. The sample size that we use for each demographic is listed in A.2. Column "Demographics" shows the demographic for the rows, column "Groups" shows all available demographic groups for the given demographic, column "Smallest Group – Year" shows the demographic group-year combination with the smallest sample size, column "Smallest Group Size" shows the corresponding smallest sample size, and column "Balanced Sample Size" shows the size of the balanced sample that the algorithms use.

Table A.2: Sample Size for AHTUS

Demographic	**Groups**	**Smallest Group – Year**	**Smallest Group Size**	**Balanced Sample Size**
Income	Top Quartile Bottom Quartile	Bottom Quartile 1998	100	200
Education	College or More High School or Less	High School or Less 1995	259	518
Gender	Male Female	Female 1995	333	666
Race	White Non-White	Non-White 1995	149	298
Urbanicity	Urban Rural	Rural 1985	353	706
Age	40 Years or Older Less Than 40 Years Old	40 Years or Older 1995	306	612

Gfk Media Research Intelligence Survey of the American Consumer (MRI)

We use the Gfk Media Research Intelligence Survey of the American Consumer (MRI) to measure cultural distance for media diet and consumer behavior. We use all the years that we have access to, which is annually from 1992 to 1999 and annually from 2001 to 2016. The types of variables that we use are:

Movies: "Did you watch movie X in the last 6 months?"

Magazines: "Did you read magazine X in the last 6 months?"[80]

TV programs: "Did you watch TV program X yesterday / in the last 7 days / 30 days / 12 months?"

[80]We did not use magazines which do not require subscription (such as magazines of airlines and retail stores) because exposure to these types of magazines may not capture people's preferences for reading these magazines.

Products: “Do you own product X / Did you use product X / Did you buy product X in the last 30 days / 6 months / 12 months?” [81]

Brands: “Do you own product from brand X / Did you use product from brand X / Did you buy product from brand X / Did you shop at store X in the last 6 months / 12 months?”[82]

As each question in the MRI has a yes (1) or no (0) answer, we define a variable as a dummy variable equal to 1 for a positive response, 0 otherwise.

MRI includes other variables that we did not use in the analysis. These include: attitudes (political affiliation, health[83], fashion[84], general[85], attitudes towards advertisements[86], personal attitudes[87], passionate about topic X[88]), time use (political activity, miles driven on a car, number of nights spent on overnight camping trips, hours listened to the radio, hours watched TV, interests, hours per week spent on doing X, time spent using the internet, hours spent playing video-game system X/video-game type X, music type X listened to in the last 6 months, hobby X, volunteered for charitable organization, member of an organization or club, leisure activity X), other consumer behavior (non-principle shopper’s purchase), other media consumption (newspapers[89],visited social networking site X in the last 30 days, visited website X in the last 30 days).

The number of variables for each module is 83 to 97 variables each year for movies, 177 to 237 variables each year for magazines, 517 to 839 variables each year for TV programs, 1,577 to 2,484 variables each year for products, and 5,664 to 6,930 variables each year for brands. We pool movies, magazines, and TV programs together as the media module; there are 879 to 1,129 variables each year for the media module. We also pool products and brands together as the consumer module; there are 7,241 to 9,368 variables each year for the consumer module.

Not all variables are available for all years. While products, brands, and TV programs are available for all years, movies are available for 1998, 1999, and annually from 2001 to 2016. Also, we only use magazines annually from 1992 to 1999 and annually from 2001 to 2011.[90] Hence, for the media module, we only use the overlapping years for movies, magazines, and TV programs, which are 1998, 1999, and

[81] We use all products except for financial and insurance products. Same for brands. We also treat travel destinations as products.

[82] We only use the question “Did you shop at store X?” if the store mainly sells products of its own brand.

[83] An example is “I go to the doctor regularly for check-ups.”

[84] An example is “Comfort is one of the most important factors when selecting fashion products to purchase.”

[85] An example is “Buying American products is important to me.”

[86] An example is “Advertising helps me keep up-to-date about products and services that I need or would like to have.”

[87] An example is “Having material possessions is important.”

[88] Example topics include health care, cooking, and grocery.

[89] Newspapers are not used because of the small number of newspapers included in the dataset; regional newspapers are not included in the US-level data that we have access to.

[90] While magazine data exist in the MRI Media Survey post-2011, the time period was reduced to the last 7 days for the weekly magazines and the last 14 days for the bi-weekly magazines starting in 2012. This makes the “Did you read magazine X” variables in 2012-2016 not comparable to those prior to 2012.

annually from 2001 to 2011.

We use all respondents from ages 20 to 64.

We use the following demographic variables in the MRI for our analysis: income, education, gender, race,[91,92] political ideology, and age.[93] Furthermore, not all demographics are available for all years. The income, race, and gender variables are available for all years, education and political ideology are available annually from 1994 to 1999 and annually from 2001 to 2016. While we use all available years for education, for political ideology we only use data from 1994 to 1999, and from 2001 to 2009. This is because the share of respondents who do not respond to the political ideology question in the period 2010 to 2013 is substantially higher than in the period 1994 to 2009, while the share in the period 2014 to 2016 is substantially lower than in the period 1994 to 2009. This suggests that the quality of the political ideology question in the period 2010 to 2016 is not the same as in the period 1994 to 2009.

The income variable available in the MRI is household income, and it is available in income brackets. The income brackets change across years.[94]

To implement the ensemble and Bayesian algorithms, we equalize sample size across years by demographic. For each binary demographic, we find the smallest sample size among all demographic group-year combinations. The algorithms first read in the entire cleaned dataset and then randomly draw a balanced sample with the two demographic groups having the same number of observations equal to the aforementioned smallest sample size. The sample size that we use for each demographic is listed in A.3. Column "Demographics" shows the demographic for the rows, column "Groups" shows all available demographic groups for the given demographic, column "Smallest Group – Year" shows the demographic group-year combination with the samllest sample size, column "Smallest Group Size" shows the corresponding smallest sample size, and column "Balanced Sample Size" shows the size of the balanced sample that the algorithms use.

[91] In MRI, the race variable has the following values for the listed years: 1992-1997 - "White," "African American," or "Other;" 1998-2002 - "White," "African American," "Asian," or "Other;" 2003-2016 - "White," "African American," "American Indian or Alaska Native," "Asian," or "Other."

[92] In MRI, there is a variable that identifies whether the respondent is of Hispanic origin. This variable is available since year 2007. We do not use this variable for our race specification.

[93] Age is only available in five-year age groups (20 to 24,..., 60 to 64).

[94] There are 14 brackets for 1992 and 1993, 15 brackets for the period 1994 to 2001, 16 brackets for the period from 2002 to 2008, and 17 brackets for the period 2009 to 2016.

Table A.3: Sample Size for MRI

Demographic	Groups	Smallest Group – Year	Smallest Group Size	Balanced Sample Size
Income	Top Quartile Bottom Quartile	Bottom Quartile 1995	2,905	5,810
Education	College or More High School or Less	High School or Less 2015	4,837	9,674
Gender	Male Female	Female 1996	7,518	15,036
Race	White Non-White	Non-White 1992	2,075	4,150
Political Ideology	Liberal Conservative	Liberal 1996	2,432	4,864
Age	40 Years or Older Less Than 40 Years Old	Less Than 40 Years Old 2013	7,243	14,486

California Department of Public Health Birth Record (CDPH)

We use the California Department of Public Health Birth Record (CDPH) to measure cultural distance for newborn's name. We use all the years that includes the demographics we are interested in, which is annually from 1960 to 2016.

The number of names for each year is 5,777 to 25,398 each year for boys and 9,739 to 35,341 each year for girls.

We use the following demographic variables in the CDPH for our analysis: mother's education, mother's race, and mother's Hispanic origin. Not all demographics are available for all years. Mother's education is available from 1989 to 2016. Mother's race (white, black, Asian, and other) is available from 1970 to 2016, and we use child's race as a proxy for mother's race from 1960-1969.[95] Mother's Hispanic origin is availabe from 1960 to 2016, which we define with the following procedure based on Fryer and Levitt (2004):[96]

[95]There are more than one race variables from 2000 to 2016, but we only use the primary one. The race variable has the following values for the listed years: 1960-1967: "White (Includes Mexican, Puerto Rican, and All Other Whites)," "Black," "American Indian (Includes Alaskan)," "Chinese," "Japanese," "Aleut," "Eskimo," "Filipino," "Hawaiian (Includes Part Hawaiian)" ("Part Hawaiian" is a separate code in 1960-1961); 1968-1977: "White," "Black," "American Indian," "Chinese," "Japanese," "Filipino" (added in 1974), "All Others;" 1978-1981: "White," "Black," "American Indian," Asian," "Other;" 1982-2016: "White," "Black," "American Indian," "Asian-Unspecified," "Asian-Specified," "Asian-Chinese," "Asian-Japanese," "Asian-Korean," "Asian-Vietnamese," "Asian-Cambodian," "Asian-Thai," "Asian-Laotian" (added in 1989), "Asian-Hmong" (added in 2000), "Other Specified," "Asian-Indian (Excluding American Indian, Aleut, Eskimo)," "Filipino," "Hawaiian," "Guamanian," "Samoan," "Eskimo," "Aleut," "Pacific Islander (Excluding Hawaiian, Guamanian, Samoan)" (added in 1985). We later define the race codes with "Asian" and Pilipino as Asian.

[96]When our definition of race involves mother's Hispanic origin, a mother would be considered as being Hispanic

1. We calculate the share of Hispanic mothers and fathers with a maiden/last name among all mothers and fathers who have that name and non-missing Hispanic code.[97] If at least 50% of mothers and fathers with a maiden/last name are Hispanic, we define the maiden/last name as a Hispanic last name.

2. If a newborn has a Hispanic last name, we define his or her mother as being Hispanic. If a newborn's last name is not matched with any maiden/last name that is ever associated with a mother or father with a non-missing Hispanic code, we drop the observation.[98]

We sample all records of newborns with non-missing first names and non-missing biological gender.

To implement the Bayesian algorithms, we equalize sample size across years by demographic. For each binary or multinary demographic, we find the smallest sample size among all demographic group-year combinations. The algorithm first reads in the entire cleaned dataset and then randomly draws a balanced sample with the demographic groups having the same number of observations equal to the aforementioned smallest sample size. The sample size that we use for each demographic is listed in A.4. Column "Demographics" shows the demographic for the rows, column "Groups" shows all available demographic groups for the given demographic, column "Smallest Group – Year" shows the demographic group-year combination with the samllest sample size, column "Smallest Group Size" shows the corresponding smallest sample size, and column "Balanced Sample Size" shows the size of the balanced sample that the algorithm uses.

A.2 Ensemble Algorithm

We use a machine learning approach to determine how predictable group membership is from a set of variables in a given year. In particular, we use an ensemble method that consists in running multiple separate algorithms and then averaging the prediction of these algorithms with weights chosen by cross-validation (Mullainathan and Spiess, 2017). We use three machine learning algorithms: elastic net regression (tuned by lambda and alpha), regression tree (tuned by the minimal node size of each tree), and random forest (tuned by the minimal node size of each tree and the proportion of variables used in each tree). We "ensemble" across algorithms with weights determined by OLS. The ensemble

regardless of her race code.

[97] Mothers' maiden names (recorded with 15 characters) are available from 1978 to 2016, and fathers' last names (recorded with 15 characters) are available from 1989 to 2016. Hispanic origins of mothers and fathers are available from 1982 to 2016. The Hispanic variable has the following values for the listed years: 1982-2016: "Not Spanish/Hispanic," "Mexican / Mexican-American / Chicano," "Puerto Rican," "Cuban," "Central/South American" (added in 1985), "Other Spanish/Hispanic (Born Outside The U.S.)," "Other Spanish/Hispanic (Born In The U.S.)" ("Other Spanish/Hispanic" is split into the last two options in 1985).

[98] The share of observations dropped (by gender) varies from 1.7% to 7.9%.

Table A.4: Sample Size for CDPH Birth Data

Panel (a) Newborn's Gender: Male

Mother's Demographic	**Groups**	**Smallest Group – Year**	**Smallest Group Size**	**Balanced Sample Size**
Education	College or More High School or Less	High School or Less 2016	94,126	188,252
Race	White Other	Other 1960	19,701	39,402
Race and Ethnicity	Non-Hispanic White Other	Other 1960	45,660	91,320
Race and Ethnicity (Pairwise Comparison)	Non-Hispanic White Black Hispanic Asian	Asian 1971	2,434	4,868

Panel (b) Newborn's Gender: Female

Mother's Demographic	**Groups**	**Smallest Group – Year**	**Smallest Group Size**	**Balanced Sample Size**
Education	College or More High School or Less	High School or Less 2016	90,578	181,156
Race	White Other	Other 1973	18,729	37,458
Race and Ethnicity	Non-Hispanic White Other	Other 1960	44,574	89,148
Race and Ethnicity (Pairwise Comparison)	Non-Hispanic White Black Hispanic Asian	Asian 1971	2,220	4,440

algorithm yields a prediction (posterior probability) that the respondent is in the given group (top income quartile, some college or more, etc.) for each respondent. We define "guess" as 1 if the prediction is greater than or equal to 0.5, 0 otherwise. We report the share of correct guesses in the hold-out sample (30%). The procedure is as follows.

1. Draw a balanced sample from the full sample, and then partition the balanced sample into a training sample (70%) and a hold-out sample (30%).

2. Tuning step (general)

    (a) Divide the training sample randomly into 5 folds. We use the same 5 folds for all three algorithms.

(b) For each fold, fit the algorithm for every tuning parameter value on all 4 other folds. Choose the optimal parameter that minimizes the mean squared-error loss over these 4 folds. Use the optimal parameter to obtain a prediction for every observation in the given fold.

(c) From 2(b), obtain one prediction for each observation in the full training sample.

(d) Repeat steps 2(b)-2(c) for each algorithm (elastic net regression, regression tree, random forest).

3. Tuning parameters (specific to each algorithm)

   (a) Elastic net regression

      i. In 2(c), elastic net regression is fit for a grid of values of lambda and alpha for the following objective function: $\min_{\beta_0,\beta} \frac{1}{N} \sum_{i=1}^{N} w_i l(y_i, \beta_0 + \beta^T x_i) + \lambda[(1-\alpha)||\beta||_2^2 + \alpha||\beta||_1]$

         A. Lambda ranges from e^{-8} to e^{10}, in increments of 0.5 for the exponent (i.e. -8, -7.5, ..., 9.5, 10). Lambda controls the penalty on the coefficients. As lambda grows larger, the penalty grows stronger, and coefficients are forced closer to zero.

         B. Alpha grid is 0, 0.5, and 1. $\alpha = 1$ case is LASSO, $\alpha = 0$ case is the ridge regression, and $\alpha = 0.5$ is the intermediate case. Alpha specifies the type of penalty applies to the coefficients. When $\alpha = 1$ (LASSO), coefficients are penalized based on the sum of their absolute values (L1 penalty). When $\alpha = 0$ (ridge regression), coefficients are penalized based on the sum of their squared values (L2 penalty). When alpha is between 0 and 1, the coefficients are penalized based on both L1 and L2 penalties, and the weights are determined by alpha.

   (b) Regression tree

      i. In 2(c), regression tree is fit for a grid of values of minimum node size ("minbucket"), where node size is the number of observations belonging to a terminal node. The grid for node size is (1, 5, 10, 20, 30, 40, 50, 70, 100, 150, 500). The depth of the tree is determined by the node size: the smaller the node size, the deeper the tree.

   (c) Random forest

      i. In 2(c), random forest is fit for a grid of values of 1) minimum node size of each tree ("node sizes") and 2) the proportion of variables used in each tree ("pmtry"). The number of trees is set to 100. The grid for node sizes is (5, 10, 20, 50, 100, 200, 400, 1000) and the grid for pmtrys is (0.1, 0.2, 0.3, 0.4).

4. Ensemble step
    (a) From 2, we have obtained one prediction for each algorithm for every observation in the full training sample.
    (b) Fit weights by running a linear regression (OLS) of the outcome on the predicted values for each algorithm in the full training sample, and store the resulting linear model.
    (c) Fit each algorithm on the full training sample and obtain optimal parameters that minimize the mean squared-error loss over the full training sample.
    (d) To predict in the hold-out sample, use the optimal parameters from 4(c) to obtain predictions for each algorithm on the hold-out sample, and then ensemble the predictions with the linear model obtained in 4(b).
5. Ensemble algorithm implementation
    (a) For each dataset-year, implement the ensemble algorithm where:
        i. LHS = Income / Education / Gender / Race / Political Ideology / Urbanicity / Age (dummy variables)
        ii. RHS = Dataset
    (b) Iterate the ensemble algorithm for X number of random subset of the dataset (X=500 for attitudes and time use, X=25 for media, movies, TV programs, magazines, X=5 for consumer behavior, products, and brands).
    (c) For each iteration, compute the hold-out sample share of correct guesses.
        i. The ensemble algorithm outputs the predictability that a respondent is in the income / demographic group for each year.
        ii. We guess whether the respondent is in that income / demographic group if the predictability is greater than or equal to / less than 0.5.
        iii. Then, for each respondent, we have the true income / demographic of the respondent and our guess using the RHS variables. We compute the hold-out sample share of correct guesses.
        iv. The ensemble algorithm uses 70% of the data to generate a prediction model (training sample), and designates the remaining 30% as the hold-out sample. We only use the hold-out sample to compute the share of correct guesses.
    (d) For each dataset-year, average the hold-out sample share of correct guesses across the iterations.

A.3 Bayesian Algorithm for Newborn's Name

We use a Bayesian approach to determine how well we can predict a mother's membership in a demographic group based on her child's name in a given year. We report the average share of correct guesses in the hold-out sample across 500 iterations. The procedure is as follows.

1. Randomly draw a balanced sample from the full sample, and then randomly partition the balanced sample into a training sample (70%) and a hold-out sample (30%).

2. In the training sample, calculate the shares of newborns with a certain name (e.g. Alice) conditional on the mothers' membership in a demographic group. Also calculate the shares of newborns with unique names (a name that appears only once in the training sample) conditional on the mothers' membership in a demographic group.

3. In the hold-out sample, guess a mother to be in the demographic group that is associated with a higher share of newborns with her child's name based on the calculation in Step 2. If her child's name does not appear in the training sample at all, guess the mother to be in the demographic group that is associated with a higher share of newborns with unique names based on the calculation in Step 2.

4. Calculate the hold-out sample share of correct guesses.

5. Repeat steps (1) to (4) 500 times. Obtain the average hold-out sample share of correct guesses across iterations.

A.4 Bayesian Algorithm for the Most Indicative Traits

We use a Bayesian approach to determine how well we can guess group membership based on a single variable in a given year. We use the results from the Bayesian approach to produce a) the table of top 10 cultural traits that are most indicative of membership in a demographic group and b) the heat map of cultural traits that are indicative of membership in a demographic group (for attitudes only).[99] The procedure is as follows:

1. Randomly draw a balanced sample from the full sample, and then randomly partition the balanced sample into a training sample (80%) and a hold-out sample (20%).

[99] When producing the tables of top 10 TV programs that are most indicative of membership in a demographic group, we create one aggregate variable for each of the following sports programs: NBA, NCAA basketball games, MLB baseball games, NFL football games, college football games, US Open (golf), and US Open (tennis). For each of these sports programs, we first sort out all variables associated with them. We then assign 1 to the aggregate variable if a respondent has a positive response to any of these variables, and assign 0 to the aggregate variable if a respondent have negative responses to all of these variables.

2. In the training sample, calculate the share of positive responses for a given variable (e.g. watched *Fox and Friends*) conditional on the respondents' membership in a demographic group.

3. In the hold-out sample, guess a respondent to be in the demographic group that is associated with a higher share of positive responses for a given variable based on the calculation in Step 2.

4. Calculate the hold-out sample share of correct predictions.

5. Repeat steps (1) to (4) 100 times. Obtain the average hold-out sample share of correct guesses across iterations.

6. In the full sample, calculate the share of positive responses for a variable (e.g. watched *Fox and Friends*) conditional on the respondents' membership in a demographic group.

The procedure for producing the table of the top ten cultural traits that are most indicative of group membership is then as follows. First, we rank each response in decreasing order of the average hold-out sample share of correct predictions obtained in Step 4. Second, we report the average hold-out sample share of correct guesses for the ten responses with the highest share of correct guesses. Third, we guessS a respondent to be in the demographic group that is associated with a higher share of positive responses for a given variable based on the calculation in Step 5 (e.g., watching *Fox and Friends* is predictive of being conservative).

The procedure for producing the heat map of cultural traits that are indicative of group membership (for attitudes only) is as follows. First, we rank each variable in increasing order of the average hold-out sample share of correct guesses obtained by the Bayesian procedure for the first year (1976 for attitudes). Variables are vertically ranked throughout the heat map figure based on that 1976 order. Second, in each subsequent year, we assign to each variable its rank in increasing order of the average hold-out sample share of correct guesses for that year. We then assign color-code to each variable's relative rank in each year, with the most informative variables being color-coded dark red and the least informative color-coded dark blue, and lighter shades of red and blue in between.

A.5 Defining income quartile cutoffs by household groups using the Current Population Survey (CPS)

We use family income for the GSS and AHTUS and household income for the MRI. Note that the income variables in all three of our main datasets are in income brackets, not continuous dollar amounts. As the CPS top / bottom income quartile cutoffs by household groups most often occur within an income bracket, using income brackets does not exactly capture the top / bottom income quartiles in the CPS. We describe below the method we use to minimize this mismeasurement.

First, we define household groups as follows. We define the households with one adult and no children as household group 1, households with two adults and no children as household group 2, households with two adults and children as household group 3, and households with one adult and children. Households with more than two adults were classified into household group 3; adults other than the two primary adults are regarded as dependents.

The procedure for defining the income quartile dummy variable is as follows. For every year-household group, we obtain from the CPS the top and bottom quartile income cutoffs as well the full income distribution. For each of the three datasets (GSS, AHTUS, MRI), we then consider all possible assignments of observations to top and bottom quartiles based on the income brackets available in that dataset-year. For each possible assignment, we count the number of observations that actually are in top / bottom quartile according to the CPS but not assigned as such, as well as the number of observations that actually are not top / bottom quartile according to the CPS but assigned as such. We call the sum of these two numbers the number of mis-measured observations. For each dataset-year-household group, we then generate the top and bottom quartile variables by choosing the assignment that minimizes the number of mis-measured observations.

The share of mis-measured observations, when averaged across household groups (with weights corresponding to the number of observations in each household group), are summarized below.

1. Top quartile:
   (a) GSS: average - 2.7%, minimum - 1.3%, maximum - 5.4%
   (b) AHTUS: average - 5.0%, minimum - 0.6%, maximum - 8.5%
   (c) MRI: average - 4.0%, minimum - 1.6%, maximum - 7.0%
2. Bottom quartile
   (a) GSS: average - 1.6%, minimum - 0.2%, maximum - 4.6%
   (b) AHTUS: average - 2.8%, minimum - 1.9%, maximum - 7.3%
   (c) MRI: average - 1.2%, minimum - 0.6%, maximum - 2.8%

While the share of mis-measured observations is less than 5% for most dataset-quartiles, the share is larger than 5% (and thus not negligible) for: MRI for years 2007-2013 for the top quartile; AHTUS for years 1965, 1998, and 2006-2012 for the top quartile; AHTUS for year 1998 for the bottom quartile;

and GSS for year 1984 for the top quartile. To investigate the effect of mismeasurement on our ability to predict, we regress the average hold-out sample share of correct guesses on an intercept, average share of mismeasurement for the top and bottom quartiles, year, and dataset dummies. First, we find that the coefficient on the average share of mismeasurement is not statistically significant (coefficient = -0.14, t-statistic = -0.27). Second, we find that the R-squared increases only minimally when we include the average share of mismeasurement; in fact, the adjusted R-squared decreased. From these two observations, we conclude that while the level of mismeasurement is not negligible, its effect on our ability to predict does not appear to be substantive.

B Supplementary Materials

B.1 Main Additional Results

B.1.1 Income

Table B.1: Attitudes and norms most indicative of being high-income

1976		1996		2016	
Trust people	67.7%	Trust people	65.4%	Voted for pres. candidate	63.7%
Voted for pres. candidate	67.2%	Voted for pres. candidate	64.5%	Trust people	62.9%
Allow homosexuals to speak	66.1%	People are helpful	62.5%	Allow abortion for married women	61.9%
Spending on space expl. isn't too much	65.8%	My health condition is very good	60.0%	Ever approve of police striking citizens	61.0%
Allow homosexuals' book in library	65.4%	Confident in the scientific community	59.7%	Allow abortion for single women	60.3%
Allow homosexuals to teach	64.6%	Federal income tax is too high	59.6%	Allow abortion for low income women	59.8%
Allow communists to speak	64.4%	Allow abortion for single women	59.1%	I am happy	59.7%
Allow anti-religionists to speak	63.6%	Allow anti-religionists to teach	58.9%	Homosexual sex isn't wrong at all	59.6%
Allow communists' book in library	63.4%	Ever approve of police striking citizens	58.9%	Not afraid to walk at night in neigh.	59.4%
People are helpful	63.0%	Allow communists to speak	58.8%	My health condition is more or less than fair	59.1%

Note: Data source is the GSS. Sample size is 394. Reported in each column are the 10 cultural traits most indicative of being rich in that year. The numbers indicate the likelihood of guessing correctly whether an individual is rich based on the answer to the question. For example, in 1976, knowing whether a person trusts people allows us to guess income correctly 67.7% of the time, whereas knowing whether a person thinks spending on space exploration is too much allows us to guess income correctly 65.8% of the time. An affirmative answer to "Do you trust people?" and a negative answer to "Is spending on space exploration too much?" indicate that the person is rich.

B.1.2 Education

Table B.2: TV shows, movies, and magazines most indicative of being more educated

Panel (a) TV shows

1994		2005		2016	
Didn't watch *Rescue 911*	55.3%	Watched *2004 Summer Olympics*	53.5%	Watched NFL football games	53.7%
Watched NCAA basketball games	54.6%	Watched *Academy Awards*	53.4%	Watched *House Hunters*	53.2%
Watched *Wimbledon*	54.4%	Didn't watch *Cops*	53.3%	Watched *Love It or List It*	53.2%
Didn't watch *Unsolved Mysteries*	54.3%	Watched NCAA basketball games	53.1%	Watched *Academy Awards*	53.1%
Watched NFL football games	53.6%	Watched *Wimbledon*	52.6%	Watched *Property Brothers*	53.1%
Didn't watch *Country Music Awards*	53.5%	Watched college football games	52.4%	Watched NCAA basketball games	52.8%
Didn't watch *In the Heat of the Night*	53.4%	Didn't watch *NASCAR Daytona 500*	52.4%	Watched *Flip or Flop*	52.7%
Didn't watch *The Oprah Winfrey Show*	53.4%	Didn't watch *Cops*	52.3%	Watched MLB baseball games	52.5%
Watched college football games	53.2%	Watched *The Masters*	52.3%	Watched college football games	52.3%
Watched MLB baseball games	53.2%	Watched NFL football games	52.3%	Didn't watch *News*	52.2%

Panel (b) Movies

1998		2007		2016	
Watched *Jerry Maguire*	55.0%	Watched *The Chronicles of Narnia 1*	53.1%	Watched *Gone Girl*	53.1%
Watched *The English Patient*	53.5%	Watched *Walk The Line*	52.9%	Watched *The Hunger Games*	52.6%
Watched *First Wive's Club*	52.9%	Watched *Pirates of The Caribbean 2*	52.3%	Watched *Interstellar*	52.4%
Watched *Air Force One*	52.4%	Watched *The Devil Wears Prada*	52.3%	Didn't watch *Annabelle*	51.6%
Watched *Star Wars - Special Edition*	52.2%	Watched *Brokeback Mountain*	52.3%	Watched *Guardians of the Galaxy*	51.5%
Watched *The Empire Strikes Back-Special edn.*	52.2%	Watched *The Da Vinci Code*	52.2%	Watched *The Theory of Everything*	51.3%
Watched *Star Trek First Contact*	52.1%	Watched *The Constant Gardener*	52.2%	Watched *Big Hero 6*	51.2%
Watched *Ransom*	52.1%	Watched *Harry Potter 4*	52.1%	Watched *Into the Woods*	51.2%
Watched *Evita*	52.0%	Watched *Memoirs Of A Geisha*	52.0%	Watched *Unbroken*	51.1%
Watched *One Fine Day*	52.0%	Didn't watch *Big Momma's House 2*	52.0%	Watched *Birdman*	51.1%

Panel (c) Magazines

1994		2002		2011	
Read *Newsweek*	60.3%	Read *Newsweek*	60.4%	Read *Time*	57.4%
Read *Time*	59.2%	Read *Time*	59.1%	Read *Newsweek*	57.2%
Read *U.S. News & World Report*	58.4%	Read *U.S.News & World Report*	56.7%	Read *Consumer Reports*	56.2%
Read *Consumer Reports*	58.0%	Read *Consumer Reports*	56.6%	Read *People*	55.8%
Read *National Geographic*	57.0%	Read *People*	55.9%	Read *National Geographic*	54.9%
Read *Business Week*	56.7%	Read *National Geographic*	55.7%	Read *The New Yorker*	54.5%
Read *Money*	55.9%	Read *Business Week*	54.3%	Read *Us Weekly*	54.1%
Read *People*	55.8%	Read *Forbes*	54.3%	Read *Real Simple*	54.1%
Didn't read *National Enquirer*	55.4%	Read *Money*	54.0%	Read *Forbes*	54.1%
Read *Smithsonian*	55.2%	Read *Fortune*	54.0%	Read *O, The Oprah Magazine*	53.8%

Note: Data source is the MRI. Sample size in all panels is 9,674. Reported in each column are the 10 cultural traits most indicative of being educated in that year. The numbers indicate the likelihood of guessing correctly whether an individual is educated based on the answer to the question. For example, in 1994, knowing whether a person watched NCAA backetball games allows us to guess education correctly 54.6% of the time, whereas knowing whether a person watched *Rescue 911* allows us to guess education correctly 55.3% of the time. An affirmative answer to "Did you watch NCAA backetball games?" and a negative answer to "Did you watch *Rescue 911*?" indicate that the person is educated.

Table B.3: Products and brands most indicative of being more educated

Panel (a) Products

1994		2005		2016	
Traveled in the continental US	59.8%	Own a personal computuer	63.4%	Used email on cellphone	65.8%
Own an imported car	59.1%	Own computer software	63.2%	Used a search engine on cellphone	64.1%
Own a personal computer	59.1%	Own computer peripherals	62.5%	Used an app on cellphone	63.4%
Traveled domestically by air	58.8%	Bought on internet	61.7%	Own a tablet or e-reader	63.2%
Used dishwasher detergent	58.8%	Own a desktop computer	61.2%	Bought on internet	63.0%
Own computer peripherals	58.5%	Traveled in the continental US	61.2%	Traveled in the continental US	62.6%
Own an answering machine	58.3%	Own word processing software	60.6%	Own computer software	62.5%
Own computer software	58.3%	Own a valid passport	60.1%	Used a website for maps on cellphone	62.4%
Own word processing software	58.0%	Own a cd rom drive	59.8%	Used internet on cellphone	62.2%
Own a desktop computer	58.0%	Own a ink-jet printer	59.5%	Own a valid passport	62.1%

Panel (b) Brands

1994		2005		2016	
Used Federal Express	56.5%	Own a computer with Windows XP	58.0%	Own an iPhone	62.9%
Bought Kodak (film)	55.6%	Own a Dell computer	56.5%	Own an iPad	61.0%
Own AAA membership	55.6%	Own AAA membership	55.9%	Own AAA membership	55.9%
Used Johnson & Johnson (dental floss)	55.1%	Bought at Starbucks	55.1%	Used Verizon Wireless (cellular)	55.9%
Own AT&T calling cards	54.9%	Used Kikkoman (soy sauce)	55.0%	Used AMC	55.7%
Used Kikkoman (soy sauce)	54.6%	Own a Sony television	54.3%	Bought at Starbucks (fast food)	55.5%
Used Grey Poupon Dijon (mustard)	54.4%	Used Bertolli (salad/cooking oil)	54.3%	Used AT&T (cellular)	55.4%
Didn't use Little Debbie (snack cakes)	54.3%	Own a Sony compact disc player	54.2%	Own an HP printer/fax machine	55.2%
Didn't use BIC (lighters)	53.8%	Didn't use BIC (lighters)	54.1%	Bought at Chipotle (fast food)	55.2%
Drank Diet Coke	53.7%	Used Grey Poupon Dijon (mustard)	53.9%	Used Expedia.com for advise about travel arrangement	55.2%

Note: Data source is the MRI. Sample size in all panels is 9,674. Reported in each column are the 10 cultural traits most indicative of being educated in that year. The numbers indicate the likelihood of guessing correctly whether an individual is educated based on the answer to the question. For example, in 1994, knowing whether a person traveled in the continental US allows us to guess education correctly 59.8% of the time, whereas in 2005, knowing whether a person bought a BIC lighter allows us to guess education correctly 54.1% of the time. An affirmative answer to "Do you own an imported car?" and a negative answer to "Did you buy a BIC lighter?" indicate that the person is educated.

Table B.4: Attitudes and norms most indicative of being more educated

1976		1996		2016	
Allow anti-religionists to teach	66.2%	Voted for pres. candidate	62.8%	Voted for pres. candidate	63.1%
Allow communists to speak	65.8%	Allow communists to speak	61.9%	Trust people	62.7%
Allow militarists to speak	64.3%	Allow communists' book in library	61.0%	Allow communists to teach	61.6%
Allow communists' book in library	64.1%	Allow militarists to speak	61.0%	Allow communists to speak	61.1%
Allow communists to teach	63.9%	Allow militarists to speak	60.5%	Allow communists' book in library	60.6%
Homosexual sex isn't always wrong	63.8%	Allow communists to teach	60.0%	Homosexual sex isn't wrong at all	59.5%
Allow anti-religionists to speak	63.0%	Trust people	59.8%	People are helpful	59.4%
Allow anti-religious' book in library	63.0%	Allow anti-religionists to teach	59.2%	Ever approve of police striking citizens	59.0%
Allow homosexuals' book in library	63.0%	Allow abortion for single women	59.0%	Allow abortion for low income women	59.0%
Allow homosexuals to speak	62.6%	Allow anti-religious' book in library	59.0%	Allow militarists to speak	58.5%

Note: Data source is the GSS. Sample size is 650. Reported in each column are the 10 cultural traits most indicative of being educated in that year. The numbers indicate the likelihood of guessing correctly whether an individual is educated based on the answer to the question. For example, in 1976, knowing whether a person thinks anti-religionists should be allowed to speak allows us to guess education correctly 66.2% of the time, whereas knowing whether a person thinks homosexual sex is not always wrong allows us to guess education correctly 63.8% of the time. An affirmative answer to "Should anti-religionists be allowed to speak?" and a negative answer to "Is homosexual sex always wrong?" indicate that the person is educated.

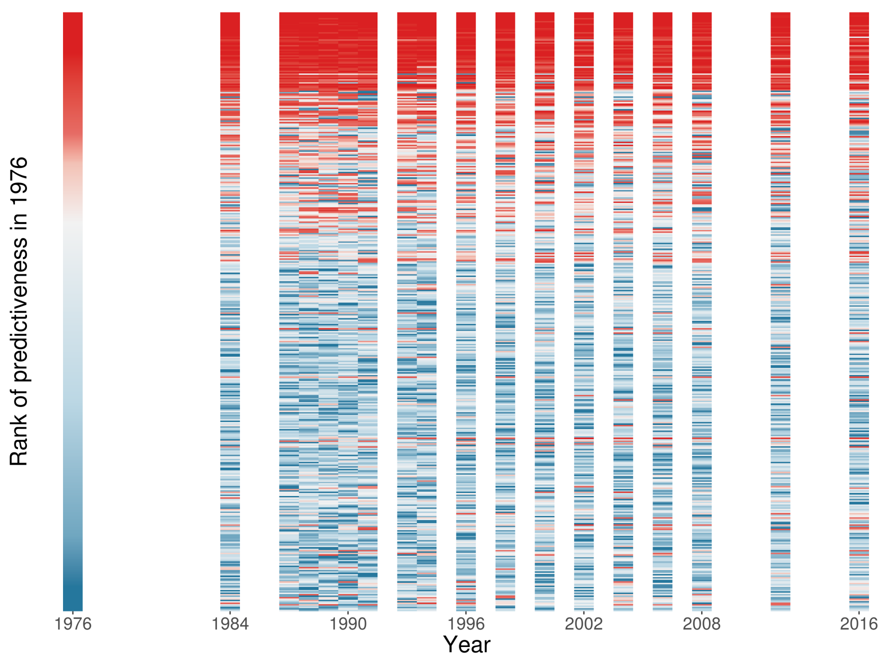

Figure B.1: Stability over time of attitudes most indicative of education

Note: Data source is the GSS. Sample size is 650. Variables are ranked from bottom to top throughout the graph by increasing order of correctly guessing education in 1976 based on that variable only. Each variable's relative informativeness in subsequent years is color-coded, with the most informative variables in each year color-coded dark red and the least informative color-coded dark blue, and lighter shades of red and blue in between. See Data Appendix for implementation details.

B.1.3 Gender

Table B.5: TV shows, movies, and magazines most indicative of being male

Panel (a) TV shows

1992		2004		2016	
Watched NFL football games	64.2%	Watched NFL football games	63.0%	Watched NFL football games	61.7%
Watched MLB baseball games	61.2%	Watched college football games	61.7%	Watched college football games	60.3%
Watched college football games	60.6%	Watched MLB baseball games	60.3%	Watched NBA games	58.5%
Watched NCAA basketball games	59.9%	Watched NBA games	58.7%	Watched MLB baseball games	58.4%
Watched NBA games	59.7%	Watched NCAA basketball games	58.5%	Watched NCAA basketball games	58.1%
Didn't watch *The Oprah Winfrey Show*	55.4%	Watched *ABC Little League World Series*	54.7%	Watched *Football Night in America*	55.4%
Didn't watch *Barbara Walters*	55.3%	Watched *NASCAR Daytona 500*	54.5%	Watched *SportsCenter*	54.9%
Didn't watch *Golden Girls*	54.8%	Watched *ABC Big 12 Championship*	54.5%	Didn't watch *Love It or List It*	54.6%
Watched *World League of American Football*	54.8%	Didn't watch *ABC Barbara Walters Specials*	54.5%	Didn't watch *House Hunters*	54.4%
Watched *US Open (Golf)*	54.5%	Watched *US Open (Golf)*	54.4%	Watched *American Pickers*	54.1%

Panel (b) Movies

1998		2007		2016	
Didn't watch *First Wive's Club*	55.3%	Didn't watch *In Her Shoes*	52.6%	Watched *Interstellar*	52.6%
Didn't watch *The Mirror Has Two Faces*	53.5%	Watched *Transporter 2*	52.2%	Watched *John Wick*	52.5%
Didn't watch *Dalmatians*	53.5%	Didn't watch *The Family Stone*	51.9%	Watched *The Hobbit 3*	52.2%
Didn't watch *The Preacher's Wife*	53.1%	Didn't watch *The Devil Wears Prada*	51.9%	Watched *Fury*	52.2%
Didn't watch *One Fine Day*	52.8%	Didn't watch *Cheaper By The Dozen 2*	51.9%	Watched *Guardians of the Galaxy*	52.1%
Didn't watch *My Best Friend's Wedding*	52.3%	Watched *Underworld: Evolution*	51.8%	Didn't watch *Gone Girl*	51.8%
Didn't watch *Jerry Maguire*	52.3%	Didn't watch *Pride And Prejudice*	51.8%	Watched *The Equalizer*	51.6%
Watched *Star Trek First Contact*	52.2%	Didn't watch *Rumor Has It*	51.8%	Didn't watch *Into the Woods*	51.6%
Didn't watch *Fly Away Home*	51.9%	Watched *King Kong*	51.8%	Didn't watch *Annie*	51.5%
Didn't watch *Michael*	51.4%	Didn't watch *Memoirs Of A Geisha*	51.7%	Watched *Mad Max*	51.5%

Panel (c) Magazines

1992		2002		2011	
Didn't read *Woman's Day*	67.8%	Didn't read *Woman's Day*	65.3%	Didn't read *Better Homes & Gardens*	65.2%
Didn't read *Family Circle*	67.5%	Didn't read *Good Housekeeping*	64.4%	Didn't read *Woman's Day*	64.8%
Didn't read *Good Housekeeping*	66.3%	Didn't read *Better Homes & Gardens*	64.3%	Didn't read *Good Housekeeping*	63.9%
Didn't read *Ladies' Home Journal*	64.4%	Didn't read *Family Circle*	63.7%	Didn't read *People*	62.4%
Didn't read *Better Homes & Gdns*	64.0%	Read *Sports Illustrated*	62.4%	Read *Sports Illustrated*	62.2%
Didn't read *McCall's*	63.5%	Didn't read *Ladies' Home Journal*	61.5%	Didn't read *Family Circle*	62.1%
Read *Sports Illustrated*	62.8%	Didn't read *Martha Stewart Living*	60.5%	Didn't read *O, The Oprah Magazine*	61.9%
Didn't read *Redbook*	61.6%	Didn't read *Glamour*	60.4%	Didn't read *Glamour*	61.5%
Didn't read *Glamour*	59.5%	Didn't read *Cosmopolitan*	60.3%	Didn't read *Cosmopolitan*	60.7%
Didn't read *Cosmopolitan*	59.4%	Didn't read *People*	59.2%	Didn't read *In Style*	60.4%

Note: Data source is the MRI. Sample size in all panels is 15,036. Reported in each column are the 10 cultural traits most indicative of being male in that year. The numbers indicate the likelihood of guessing correctly whether an individual is male based on the answer to the question. For example, in 1992, knowing whether a person watched NFL football games allows us to guess gender correctly 64.2% of the time, whereas knowing whether a person watched *The Oprah Winfrey Show* allows us to guess gender correctly 55.4% of the time. An affirmative answer to "Did you watch NFL football games?" and a negative answer to "Did you watch *The Oprah Winfrey Show*?" indicate that the person is male.

Table B.6: Products and brands most indicative of being male

Panel (a) Products

1992		2004		2016	
Didn't use perfume/cologne for women	90.8%	Didn't use lipstick & lip gloss	87.9%	Didn't use hair care products for women	88.4%
Didn't use lipstick & lip gloss	90.0%	Didn't use perfume and cologne for women	87.4%	Didn't use perfume and cologne for women	84.8%
Didn't use hair care products for women	87.7%	Didn't use hair care products for women	87.1%	Didn't buy women's clothing	83.5%
Didn't use a blusher	86.3%	Didn't use facial moisturizers	84.2%	Didn't use lipstick & lip gloss	83.4%
Used aftershave lotion/cologne for men	84.5%	Didn't use a blow dryer	83.2%	Didn't use mascara	83.2%
Didn't use mascara	83.6%	Didn't buy women's clothing	82.8%	Didn't use a blow dryer	82.2%
Didn't buy stockings/pantyhose	82.5%	Didn't use mascara	82.0%	Didn't buy women's lingerie/undergarments	82.0%
Didn't use foundation make-up	82.4%	Didn't use foundation make-up	80.4%	Didn't use foundation make-up	80.7%
Didn't use face creams and lotions	82.4%	Didn't use a blusher	78.1%	Didn't use eye liner	79.2%
Didn't use a blow dryer	82.1%	Used aftershave lotion & cologne for men	77.8%	Didn't use eye shadow	77.4%

Panel (b) Brands

1992		2004		2016	
Didn't use Cutex (nail polish remover)	68.3%	Didn't use Cutex (nail polish remover)	62.6%	Didn't buy Victoria's Secret (lingerie)	60.7%
Didn't buy L'eggs (stockings)	63.2%	Didn't use Lady Bic (disposable razors)	58.6%	Didn't use Bath & Body Works (perfume)	59.2%
Didn't use Massengill Douche (hygiene douches)	59.0%	Didn't use Bath & Body Works (h/b cream)	58.2%	Didn't use Cutex (nail polish remover)	58.3%
Didn't use Tampax (tampon)	58.8%	Didn't buy at Bath & Body Works	57.4%	Didn't buy Old Navy (women's clothing)	57.7%
Used Mennen Speed Stick (deodorants)	58.0%	Didn't use Bath & Body Works (bath additives)	57.2%	Didn't use Bath & Body Works (h/b cream)	57.6%
Didn't use Oil of Olay (face creams)	57.2%	Used Norelco (electric shavers)	56.7%	Didn't use OPI (nail care products)	57.5%
Didn't use Avon (lipstick & lip gloss)	57.1%	Didn't use Tampax Cardboard Applicator (tampons)	56.3%	Didn't buy at Bath & Body Works	57.2%
Own a Range Rover	57.0%	Didn't use Bath & Body Works (body wash)	56.1%	Didn't buy Hanes (lingerie)	57.1%
Didn't buy No Nonsense (stockings)	56.9%	Didn't use Bath & Body Works (perfume)	56.1%	Didn't use Secret Invisible Solid (deodorants)	56.9%
Used Old Spice (aftershave lotion & cologne)	56.5%	Used Gillette Mach 3 (razor blades)	56.0%	Didn't use Dove Solid (deodorants)	56.8%

Note: Data source is the MRI. Sample size in all panels is 15,036. Reported in each column are the 10 cultural traits most indicative of being male in that year. The numbers indicate the likelihood of guessing correctly whether an individual is male based on the answer to the question. For example, in 1992, knowing whether a person bought aftershave lotion/cologne for men allows us to guess gender correctly 84.5% of the time, whereas knowing whether a person bought perfume/cologne for women allows us to guess gender correctly 90.8% of the time. An affirmative answer to “Did you buy aftershave lotion/cologne for men?” and a negative answer to “Did you buy perfume/cologne for women?” indicate that the person is male.

Table B.7: Attitudes and norms most indicative of being male

1976		1996		2016	
Not afraid to walk at night in neigh.	68.6%	Not afraid to walk at night in neigh.	63.9%	Watched an X-rated movie in the last year	61.3%
Spending on space expl. isn't too much	60.7%	Porn shouldn't be illegal to all	61.1%	Not afraid to walk at night in neigh.	60.2%
Watched an X-rated movie in the last year	58.6%	Approve of police striking citizens who escape custody	58.1%	Spending on space expl. is too little	57.8%
Oppose gun permits	58.1%	Ever approve of police striking citizens	57.5%	Porn shouldn't be illegal to all	57.7%
Porn shouldn't be illegal to all	58.1%	Oppose gun permits	57.5%	Ever approve of police striking citizens	57.2%
Spending on military is too little	57.2%	Own shotgun in home	57.1%	Not confident in banks/fin. institutions	56.3%
Favor death penalty for murder	56.8%	Watched an X-rated movie in the last year	56.7%	Extramarital sex isn't always wrong	55.8%
Not moderate	56.2%	Spending on space expl. isn't too much	56.6%	Trust people	55.7%
Not confident in organized labor	56.0%	Favor death penalty for murder	56.6%	Spending on health care isn't too little	55.6%
Marijuana should be made legal	55.8%	Own gun in home	56.0%	Federal income tax isn't too high	55.6%

Note: Data source is the GSS. Sample size is 984. Reported in each column are the 10 cultural traits most indicative of being male in that year. The numbers indicate the likelihood of guessing correctly whether an individual is male based on the answer to the question. For example, in 1976, knowing whether a person watched an X-rated movie in the last year allows us to guess gender correctly 58.6% of the time, whereas knowing whether a person is afraid to walk at night in the neighborhood allows us to guess gender correctly 68.6% of the time. An affirmative answer to "Did you watch an X-rated movie in the last year?" and a negative answer to "Are you afraid to walk at night in the neighborhood?" indicate that the person is male.

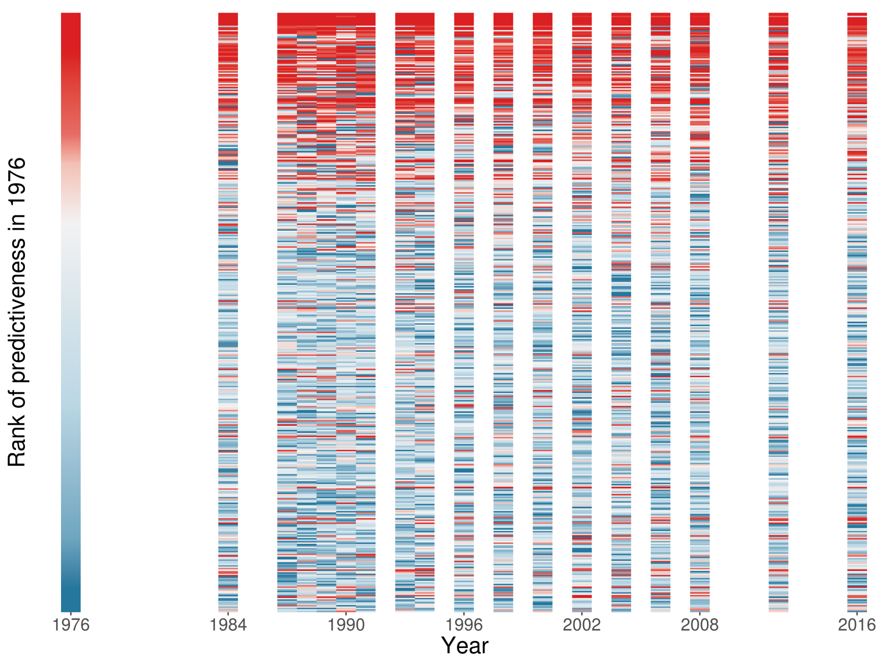

Figure B.2: Stability over time of attitudes most indicative of gender

Note: Data source is the GSS. Sample size is 984. Variables are ranked from bottom to top throughout the graph by increasing order of correctly guessing gender in 1976 based on that variable only. Each variable's relative informativeness in subsequent years is color-coded, with the most informative variables in each year color-coded dark red and the least informative color-coded dark blue, and lighter shades of red and blue in between. See Data Appendix for implementation details.

B.1.4 Race

Table B.8: TV shows, movies, and magazines most indicative of being white

Panel (a) TV shows

1992		2004		2016	
Didn't watch *In Living Color*	58.5%	Watched *2002 Winter Olympics*	61.1%	Didn't watch NBA games	57.0%
Didn't watch *Cosby Show*	58.0%	Didn't watch NBA games	56.2%	Watched *American Pickers*	54.9%
Didn't watch *Arsenio Hall*	57.9%	Didn't watch *The Parkers*	55.4%	Watched NFL football games	54.4%
Didn't watch *A Different World*	57.4%	Watched *NASCAR Daytona 500*	55.1%	Didn't watch *Empire*	54.4%
Watched *National Geographic Specials*	55.6%	Watched NFL football games	55.0%	Watched *Macy's Thanksgiving Day Parade*	54.4%
Didn't watch *Cosby*	55.4%	Watched *Dick Clark's New Years Rockin' Eve*	54.8%	Watched MLB baseball games	54.2%
Watched *Tournament of Roses Parade*	55.2%	Watched *Macy's Thanksgiving Day Parade*	54.7%	Watched *Rudolph the Red-Nosed Reindeer*	54.1%
Didn't watch *In Heat of the Night*	55.1%	Didn't watch *Soul Train Music Awards*	54.4%	Watched *The Big Bang Theory*	54.0%
Didn't watch *True Colors*	55.0%	Watched MLB baseball games	54.2%	Watched *SNL Specials*	53.7%
Watched *Country Music Awards*	54.8%	Watched *NASCAR Talladega 500*	54.1%	Watched *NHL Stanley Cup Finals*	53.7%

Panel (b) Movies

1998		2007		2016	
Didn't watch *The Preacher's Wife*	55.6%	Watched *Walk The Line*	55.7%	Didn't watch *No Good Deed*	54.2%
Watched *Jerry Maguire*	54.6%	Didn't watch *Big Momma's House 2*	55.6%	Didn't watch *The Equalizer*	53.5%
Watched *Michael*	54.5%	Didn't watch *Final Destination 3*	53.6%	Didn't watch *Furious 7*	52.8%
Watched *First Wive's Club*	53.9%	Didn't watch *Get Rich Or Die Tryin'*	53.4%	Didn't watch *Selma*	52.4%
Watched *The English Patient*	53.0%	Didn't watch *Tyler Perry's Madea's Reunion*	53.3%	Didn't watch *Annabelle*	52.3%
Didn't watch *Space Jam*	52.9%	Didn't watch *Saw II*	53.0%	Watched *The Hunger Games*	52.0%
Didn't watch *How to Be a Player*	52.7%	Watched *The Chronicles of Narnia 1*	52.6%	Didn't watch *Annie*	51.9%
Watched *One Fine Day*	52.5%	Didn't watch *Transporter 2*	52.6%	Didn't watch *Let's Be Cops*	51.9%
Watched *Fly Away Home*	52.4%	Watched *Pirates of The Caribbean 2*	52.5%	Didn't watch *Beyond The Lights*	51.9%
Watched *Dalmatians*	52.3%	Didn't watch *King Kong*	52.4%	Didn't watch *Top Five*	51.8%

Panel (c) Magazines

1992		2002		2011	
Didn't read *Ebony*	69.3%	Didn't read *Ebony*	72.0%	Didn't read *Ebony*	63.6%
Didn't read *Jet*	68.0%	Didn't read *Jet*	71.7%	Didn't read *Essence*	61.6%
Didn't read *Essence*	63.1%	Didn't read *Essence*	68.1%	Didn't read *Jet*	61.4%
Didn't read *Black Enterprise*	56.7%	Didn't read *Black Enterprise*	61.5%	Didn't read *Black Enterprise*	57.4%
Read *National Geographic*	55.5%	Didn't read *Vibe*	60.5%	Didn't read *TV Guide*	55.5%
Read *Modern Maturity*	55.4%	Didn't read *The Source*	57.2%	Didn't read *Vogue*	54.4%
Read *Consumer Reports*	54.9%	Didn't read *Gentlemen's Quarterly*	54.0%	Didn't read *Life & Style Weekly*	54.4%
Read *Country Living*	54.2%	Didn't read *TV Guide*	53.8%	Didn't read *ESPN The Magazine*	54.2%
Read *Reader's Digest*	53.8%	Didn't read *National Enquirer*	53.7%	Didn't read *People en Español*	54.1%
Read *Field & Stream*	53.6%	Didn't read *Vogue*	53.6%	Didn't read *Seventeen*	54.0%

Note: Data source is the MRI. Sample size in all panels is 4,150. Reported in each column are the 10 cultural traits most indicative of being white in that year. The numbers indicate the likelihood of guessing correctly whether an individual is white based on the answer to the question. For example, in 1992, knowing whether a person watched *National Geographic Specials* allows us to guess race correctly 55.6% of the time, whereas knowing whether a person watched *In Living Color* allows us to guess race correctly 58.5% of the time. An affirmative answer to "Did you watch *National Geographic Specials*?" and a negative answer to "Did you watch *In Living Color*?" indicate that the person is white.

Table B.9: Products and brands most indicative of being white

Panel (a) Products

1992		2004		2016	
Own a pet	62.9%	Own high-ticket sport/recreation equipment	65.5%	Own a battery flashlight	64.1%
Own a washing machine	62.6%	Own a pet	65.0%	Own a pet	63.2%
Own a microwave oven	62.1%	Own a battery flashlight	64.4%	Own a smoke/fire detector	62.9%
Own high-ticket sport/recreation equipment	62.1%	Used dishwasher detergent	64.1%	Own sport/recreation equipment	62.8%
Own a refrigerator	61.8%	Own a hot water heater	64.0%	Own a hot water heater	62.3%
Own a smoke/fire detector	61.8%	Own an automatic coffee maker	63.8%	Own low-ticket lawn/porch furniture	62.0%
Used suntan & sunscreen products	61.8%	Own low-ticket sport/recreation equipment	63.7%	Used dishwasher detergent	62.0%
Own a climate control appliance	61.4%	Own a smoke/fire detector	63.6%	Own a gas grill	61.9%
Own a hot water heater	61.3%	Own cruise control on vehicle	63.5%	Own glass ovenware/bakeware	61.9%
Own a shovel	61.2%	Own a washing machine	63.5%	Own an air conditioner	61.6%

Panel (b) Brands

1992		2004		2016	
Bought Kodak (film)	59.3%	Used Scotch Magic (transparent tape)	60.3%	Used Verizon Wireless (cellular)	60.2%
Used Scotch Magic (transparent tape)	59.1%	Used Nestlé (baking chips)	59.2%	Used Nestlé (baking chips)	57.5%
Bought BIC (pens)	58.0%	Used Arm & Hammer (baking soda)	57.6%	Used Thomas' (English muffins)	56.9%
Used Arm & Hammer (baking soda)	57.7%	Used Cut-Rite (waxed paper)	57.0%	Didn't use Dove (soaps)	56.6%
Used AT&T (long distance call service)	57.6%	Used Pam Regular (cooking products)	56.8%	Used Scotch Magic (transparent tape)	56.6%
Used Philadelphia (cream cheese)	57.5%	Used Heinz (ketchup)	56.4%	Used Shout (laundry pre-treatments)	56.2%
Used Nestlé (baking chips)	57.3%	Used French's (mustard)	56.2%	Didn't use Fabuloso (household cleaners)	56.0%
Used Elmer's (glue)	57.0%	Used Vlasic (pickles)	56.2%	Didn't use T-Mobile (cellular)	55.9%
Used Cut-Rite (waxed paper)	56.8%	Used Elmer's (glue)	56.0%	Used Sweet Baby Ray's Barbecue Sauce	55.7%
Own a Range Rover	56.4%	Own a Ford	56.0%	Didn't use Ajax Lemon (dishwashing liquid)	55.7%

Note: Data source is the MRI. Sample size in all panels is 4,150. Reported in each column are the 10 cultural traits most indicative of being white in that year. The numbers indicate the likelihood of guessing correctly whether an individual is white based on the answer to the question. For example, in 1992, knowing whether a person owns a pet allows us to guess race correctly 62.9% of the time, whereas in 2016, knowing whether a person bought Dove (soaps) allows us to guess race correctly 56.6% of the time. An affirmative answer to "Do you own a pet?" and a negative answer to "Did you buy Dove (soaps)?" indicate that the person is white.

Table B.10: Attitudes and norms most indicative of being white

1976		1996		2016	
Spending on blacks isn't too little	75.1%	Spending on blacks isn't too little	68.9%	Ever approve of police striking citizens	66.7%
Not a fundamentalist	70.2%	Ever approve of police striking citizens	64.7%	Approve of police striking citizens who escape custody	63.7%
Trust people	66.8%	Spending on welfare isn't too little	62.8%	Approve of police striking citizens who attack with fists	61.8%
Voted for Republican pres. candidate	66.2%	Spending on space expl. isn't too much	61.7%	Spending on blacks isn't too little	60.8%
None opposite race in neighborhood	65.2%	Own gun in home	61.6%	Own shotgun in home	60.8%
Ever approve of police striking citizens	63.5%	Voted for Republican pres. candidate	61.6%	Own rifle in home	60.6%
People are helpful	63.3%	Own rifle in home	61.5%	Own gun in home	60.6%
Approve of police striking citizens who escape custody	62.0%	Approve of police striking citizens who escape custody	61.5%	Allow communists' book in library	60.5%
Favor death penalty for murder	61.5%	Favor death penalty for murder	60.9%	Didn't voted for Democrat pres. candidate	60.1%
Confident in the scientific community	60.8%	Own shotgun in home	60.4%	Homosexual sex isn't wrong at all	59.9%

Note: Data source is the GSS. Sample size is 228. Reported in each column are the 10 cultural traits most indicative of being white in that year. The numbers indicate the likelihood of guessing correctly whether an individual is white based on the answer to the question. For example, in 1976, knowing whether a person trusts people allows us to guess race correctly 66.8% of the time, whereas knowing whether a person thinks spending on blacks is too little allows us to guess race correctly 75.1% of the time. An affirmative answer to "Do you trust people?" and a negative answer to "Is spending on blacks too little?" indicate that the person is white.

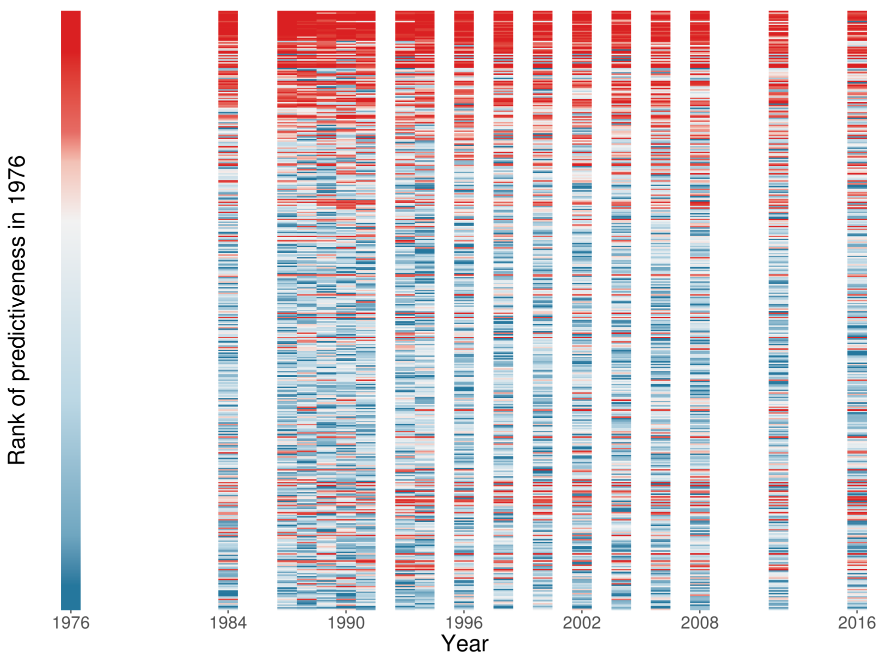

Figure B.3: Stability over time of attitudes most indicative of race

Note: Data source is the GSS. Sample size is 228. Variables are ranked from bottom to top throughout the graph by increasing order of correctly guessing race in 1976 based on that variable only. Each variable's relative informativeness in subsequent years is color-coded, with the most informative variables in each year color-coded dark red and the least informative color-coded dark blue, and lighter shades of red and blue in between. See Data Appendix for implementation details.

B.1.5 Political Ideology

Table B.11: TV shows, movies, and magazines most indicative of being liberal

Panel (a) TV shows

1994		2001		2009	
Watched *Academy Awards*	55.1%	Didn't watch college football games	53.8%	Didn't watch *The O'Reilly Factor*	57.2%
Didn't watch *Bob Hope Specials*	54.8%	Watched *Academy Awards*	53.4%	Didn't watch *Hannity & Colmes*	56.4%
Didn't watch *Rush Limbaugh*	54.3%	Watched *Will & Grace*	53.3%	Didn't watch *Fox and Friends*	55.8%
Watched *SNL Anniv. Specials*	54.0%	Watched *Ally McBeal*	52.9%	Watched *The Daily Show with Jon Stewart*	55.4%
Watched *Grammy Awards*	54.0%	Watched *Grammy Awards*	52.9%	Didn't watch *Fox Report with Shepard Smith*	54.8%
Watched *Roseanne*	54.0%	Watched *Friends*	52.8%	Didn't watch college football games	54.7%
Watched *MTV Music Video Awards*	53.8%	Watched *Golden Globe Awards*	52.5%	Didn't watch *Fox News Sunday*	54.3%
Watched *Northern Exposure*	53.7%	Didn't watch *Country Music Awards*	52.4%	Watched *Academy Awards*	54.2%
Didn't watch college football games	53.6%	Watched *The Simpsons*	52.4%	Didn't watch *NASCAR Daytona 500*	54.1%
Watched *Murphy Brown*	53.4%	Watched *Saturday*	52.4%	Watched *The Colbert Report*	54.1%

Panel (b) Movies

1998		2004		2009	
Watched *Jerry Maguire*	55.0%	Watched *Chicago*	55.1%	Watched *Juno*	56.2%
Watched *The English Patient*	53.8%	Watched *The Hours*	54.6%	Watched *No Country for Old Men*	54.2%
Watched *The People vs. Larry Flynt*	53.7%	Watched *About Schmidt*	53.9%	Watched *Michael Clayton*	53.7%
Watched *Ransom*	53.0%	Watched *Adaptation*	53.6%	Watched *Sweeney Todd*	52.9%
Watched *First Wive's Club*	52.9%	Watched *8 Mile*	53.5%	Watched *Atonement*	52.7%
Watched *Evita*	52.4%	Watched *Catch Me If You Can*	53.2%	Watched *American Gangster*	52.5%
Watched *Michael*	52.3%	Watched *Lord of the Rings 2*	53.1%	Didn't watch *National Treasure 2*	52.3%
Watched *Mars Attacks*	52.3%	Watched *Harry Potter 2*	53.0%	Watched *The Golden Compass*	52.1%
Watched *Sleepers*	52.2%	Watched *The Pianist*	52.9%	Watched *Sex And The City*	52.1%
Watched *William Shakespeare's Romeo & Juliet*	52.1%	Watched *Gangs of New York*	52.7%	Watched *Charlie Wilson's War*	52.1%

Panel (c) Magazines

1994		2001		2009	
Read *Rolling Stone*	54.6%	Read *Cosmopolitan*	53.7%	Read *The New Yorker*	54.8%
Read *Cosmopolitan*	54.4%	Read *Entertainment Weekly*	53.6%	Read *Rolling Stone*	54.8%
Didn't read *Reader's Digest*	53.9%	Read *Rolling Stone*	53.6%	Read *Vanity Fair*	54.6%
Read *Newsweek*	53.7%	Read *Vogue*	53.4%	Read *Vogue*	54.5%
Read *Time*	53.4%	Read *Vanity Fair*	52.9%	Read *Time*	54.1%
Read *Vanity Fair*	53.3%	Read *People*	52.7%	Read *People*	53.7%
Read *The New Yorker*	53.3%	Read *The New Yorker*	52.7%	Read *Us Weekly*	53.6%
Read *Elle*	53.2%	Didn't read *Reader's Digest*	52.7%	"Read *O, The Oprah Magazine*"	53.6%
Read *Vogue*	53.2%	Read *Newsweek*	52.4%	Read *Newsweek*	53.5%
Read *People*	53.0%	Read *Elle*	52.4%	Read *Entertainment Weekly*	53.2%

Note: Data source is the MRI. Sample size in all panels is 4,864. Reported in each column are the 10 cultural traits most indicative of being liberal in that year. The numbers indicate the likelihood of guessing correctly whether an individual is liberal based on the answer to the question. For example, in 1994, knowing whether a person watched *Academy Awards* allows us to guess political ideology correctly 55.1% of the time, whereas knowing whether a person watched *Bob Hope Specials* allows us to guess political ideology correctly 54.8% of the time. An affirmative answer to "Did you watch *Academy Awards*?" and a negative answer to "Did you watch *Bob Hope Specials*?" indicate that the person is liberal.

Table B.12: Products and brands most indicative of being liberal

Panel (a) Products

1994		2001		2009	
Drank bottled water & seltzer	56.3%	Drank imported beer	56.7%	Not own a fishing rod	56.8%
Drank beer	56.1%	Drank alcoholic beverages	56.6%	Bought a novel	56.7%
Didn't use gelatin and gelatin desserts	56.1%	Drank distilled liquor	56.4%	Not own fishing lures or hooks	56.3%
Used tampons for women	55.9%	Drank other alcoholic beverages	56.1%	Not own a fishing reel	56.2%
Drank alcoholic beverages	55.9%	Drank beer	55.9%	Not own a domestic vehicle	56.1%
Drank imported beer	55.9%	Bought alternative music (tapes & discs)	55.9%	Didn't use disposable plates	55.7%
Drank white goods (alcohol)	55.6%	Bought a novel	55.6%	Bought a book	55.7%
Bought audio tapes & discs	55.3%	Drank mixed drinks	55.2%	Not own other fishing equipment	55.5%
Drank other alcoholic beverages	55.2%	Drank white goods (alcohol)	55.1%	Drank imported beer/ale	55.4%
Not own a truck/van/suv	55.2%	Drank wine	55.0%	Didn't use refrigerated/frozen bread and dough products	55.3%

Panel (b) Brands

1994		2001		2009	
Didn't use Jell-O Regular	54.9%	Didn't buy at Cracker Barrel (family rest.)	53.3%	Bought at Starbucks (fast food)	54.6%
Didn't use Morton (salt)	53.6%	Didn't use Cool Whip (whipped topping)	53.3%	Bought at Ikea	54.3%
Didn't use Arm & Hammer (baking soda)	53.2%	Didn't use Hunts (canned tomatoes)	53.1%	Didn't use Cool Whip (whipped topping)	54.1%
Didn't use Crisco Regular (shortening)	53.1%	Didn't use Crisco Regular (shortening)	53.0%	Didn't buy at Arby's (fast food)	53.8%
Didn't use French's (mustard)	53.1%	Didn't buy at Arby's (fast food)	52.9%	Didn't use Bush's Best Baked Beans (canned)	53.8%
Didn't buy at Arby's (fast food)	53.1%	Didn't use Star Kist (canned tuna)	52.9%	Not own a Chevrolet	53.5%
Bought Trojan (condoms)	53.1%	Didn't use Green Giant (canned or jarred vegetables)	52.8%	Used Burt's Bees (lip care)	53.3%
Didn't buy at Dairy Queen	52.9%	Used Ben & Jerry's (ice cream)	52.8%	Didn't use Nestlé (baking chips)	53.2%
Didn't use Elmer's (glue)	52.8%	Didn't use Little Debbie (snack cakes)	52.8%	Didn't use Jimmy Dean (sausage)	53.2%
Bought at The Gap	52.8%	Didn't use Gold Medal (flour)	52.8%	Used Ben & Jerry's (ice cream)	53.1%

Note: Data source is the MRI. Sample size in all panels is 4,864. Reported in each column are the 10 cultural traits most indicative of being liberal in that year. The numbers indicate the likelihood of guessing correctly whether an individual is liberal based on the answer to the question. For example, in 1994, knowing whether a person bought bottled water and seltzer allows us to guess political ideology correctly 56.3% of the time, whereas knowing whether a person bought gelatin and gelatin desserts allows us to guess political ideology correctly 56.1% of the time. An affirmative answer to "Did you buy bottled water and seltzer?" and a negative answer to "Did you buy gelatin and gelatin desserts?" indicate that the person is liberal.

Table B.13: Attitudes and norms most indicative of being liberal

1976		1996		2016	
Marijuana should be made legal	65.5%	Homosexual sex isn't always wrong	66.6%	Allow abortion for single women	71.2%
Extramarital sex isn't always wrong	63.1%	Allow abortion for low income women	63.6%	Allow abortion for married women	70.4%
Oppose death penalty for murder	62.9%	Allow abortion for single women	63.0%	Allow abortion for low income women	68.9%
Spending on blacks is too little	62.3%	Spending on the environment is too little	61.6%	Homosexual sex isn't wrong at all	67.2%
Spending on big cities is too little	62.1%	Spending on welfare isn't too much	61.4%	Spending on military is too much	66.0%
Homosexual sex isn't always wrong	61.6%	Spending on military is too much	61.3%	Spending on the environment is too little	65.1%
Allow anti-religionists to teach	61.4%	Allow abortion for married women	61.0%	Spending on blacks is too little	64.8%
Allow communists to teach	61.1%	Marijuana should be made legal	60.4%	Oppose death penalty for murder	63.7%
Porn shouldn't be illegal to all	61.1%	Spending on health care is too little	60.2%	Extramarital sex isn't always wrong	63.7%
Spending on military is too much	60.5%	Spending on blacks is too little	59.9%	Allow abortion for rape victims	62.7%

Note: Data source is the GSS. Sample size is 552. Reported in each column are the 10 cultural traits most indicative of being liberal in that year. The numbers indicate the likelihood of guessing correctly whether an individual is liberal based on the answer to the question. For example, in 1976, knowing whether a person thinks marijuana should be made legal allows us to guess political ideology correctly 65.5% of the time, whereas knowing whether a person thinks extramarital sex is always wrong allows us to guess political ideology correctly 63.1% of the time. An affirmative answer to "Should marijuana be made legal?" and a negative answer to "Is extramarital sex always wrong?" indicate that the person is liberal.

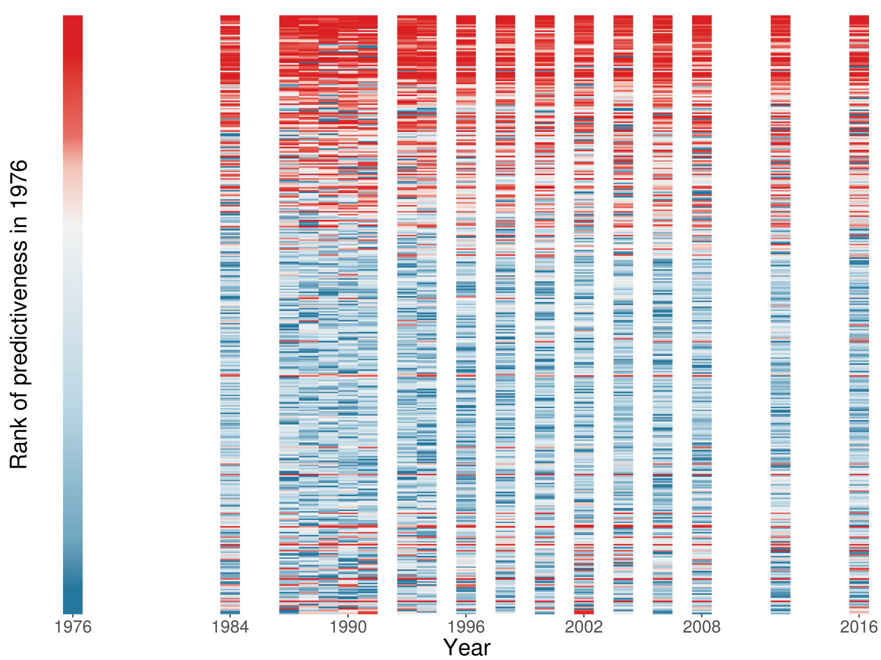

Figure B.4: Stability over time of attitudes most indicative of political ideology

Note: Data source is the GSS. Sample size is 552. Variables are ranked from bottom to top throughout the graph by increasing order of correctly guessing political ideology in 1976 based on that variable only. Each variable's relative informativeness in subsequent years is color-coded, with the most informative variables in each year color-coded dark red and the least informative color-coded dark blue, and lighter shades of red and blue in between. See Data Appendix for implementation details.

B.2 Robustness

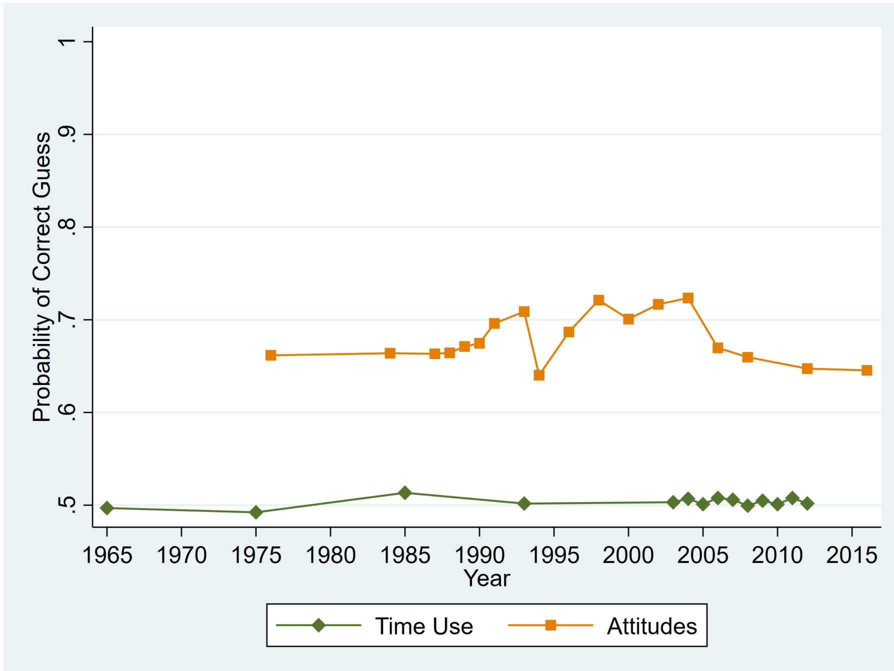

Figure B.5: Cultural distance by urbanicity

Note: Data sources are the GSS and the AHTUS. Sample sizes each year are 706 for time use and 230 for attitudes. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's urbanicity in the hold-out sample each year. The procedure to guess urbanicity in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations.

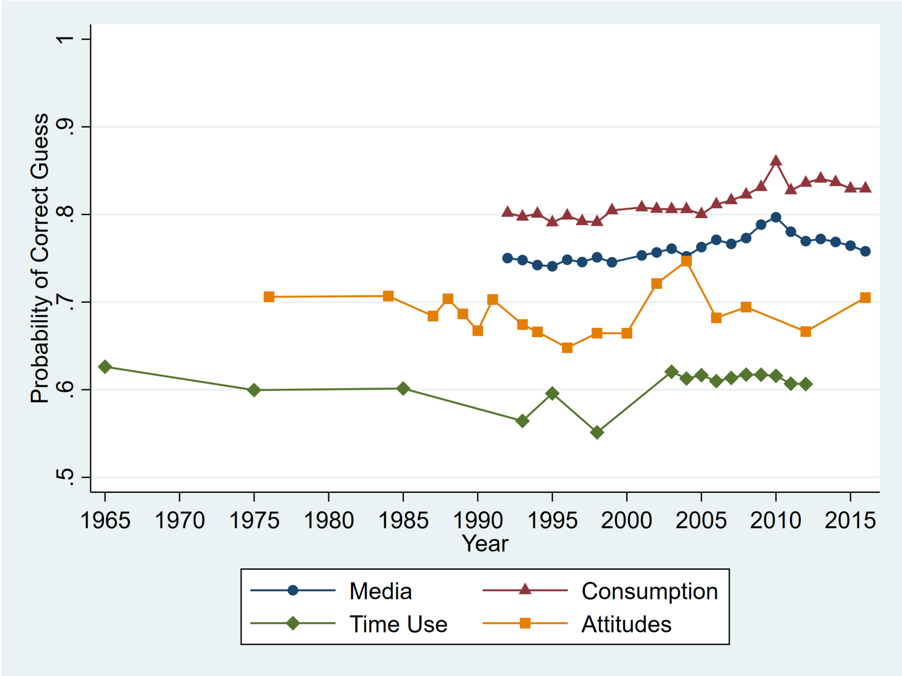

Figure B.6: Cultural distance by age

Note: Data sources are the GSS, the AHTUS, and the MRI. Sample sizes each year are 14,486 for media and consumption, 612 for time use, and 958 for attitudes. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's age in the hold-out sample each year. The procedure to guess age in the hold-out sample was repeated 5 times for consumption, 25 times for media, and 500 times for time use and attitudes, and the share of guesses reported is the average of these iterations.

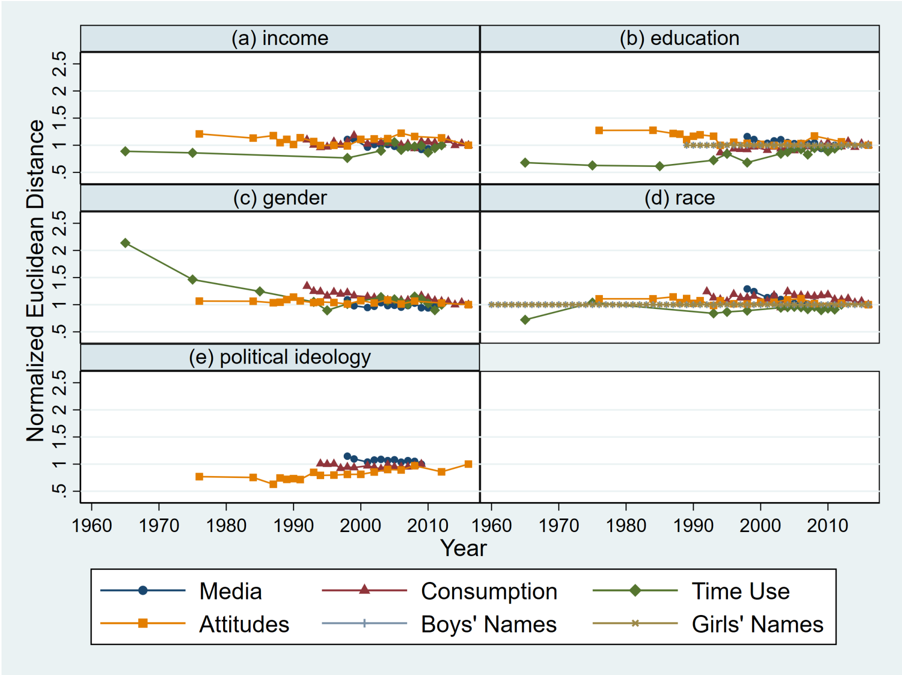

Figure B.7: Cultural distance over time: Euclidean distance

Note: Figure reports normalized Euclidean distances between groups in each year based on media diet, consumer behavior, time use, or social attitudes.

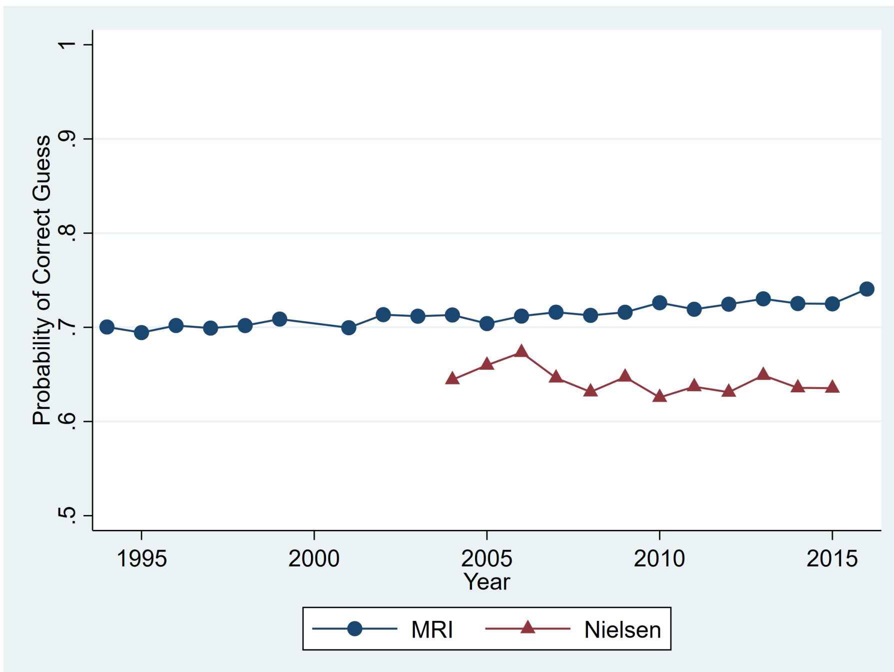

Figure B.8: Cultural distance by education over time: consumer behavior

Note: Data sources are the MRI and Nielsen. Sample sizes each year are 9,674 for MRI and 2,164 for Nielsen. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's education in the hold-out sample each year. The procedure to guess education in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these iterations.

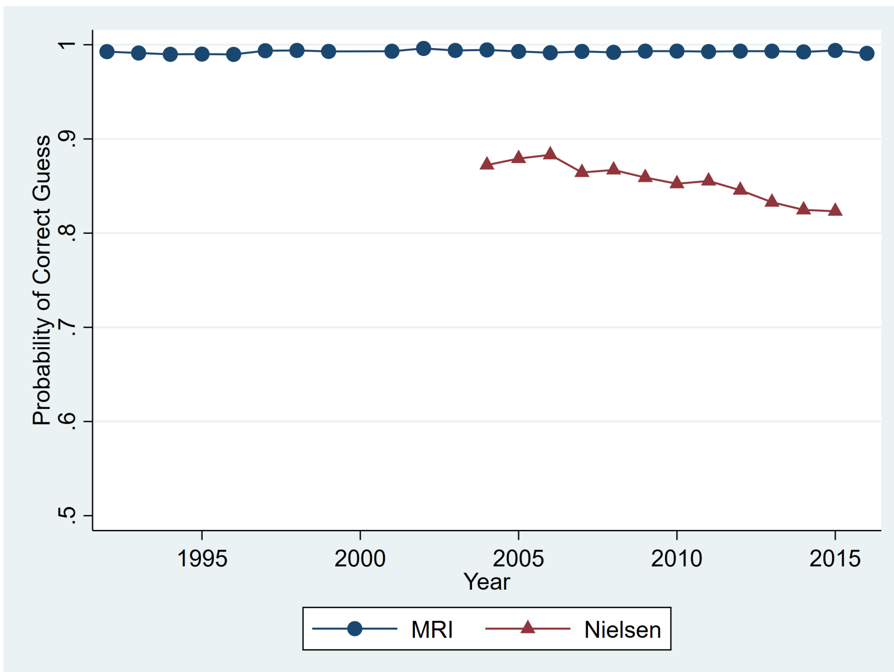

Figure B.9: Cultural distance by gender over time: consumer behavior

Note: Data sources are the MRI and Nielsen. Sample sizes each year are 15,036 for MRI and 4,566 for Nielsen. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's gender in the hold-out sample each year. The procedure to guess gender in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these iterations.

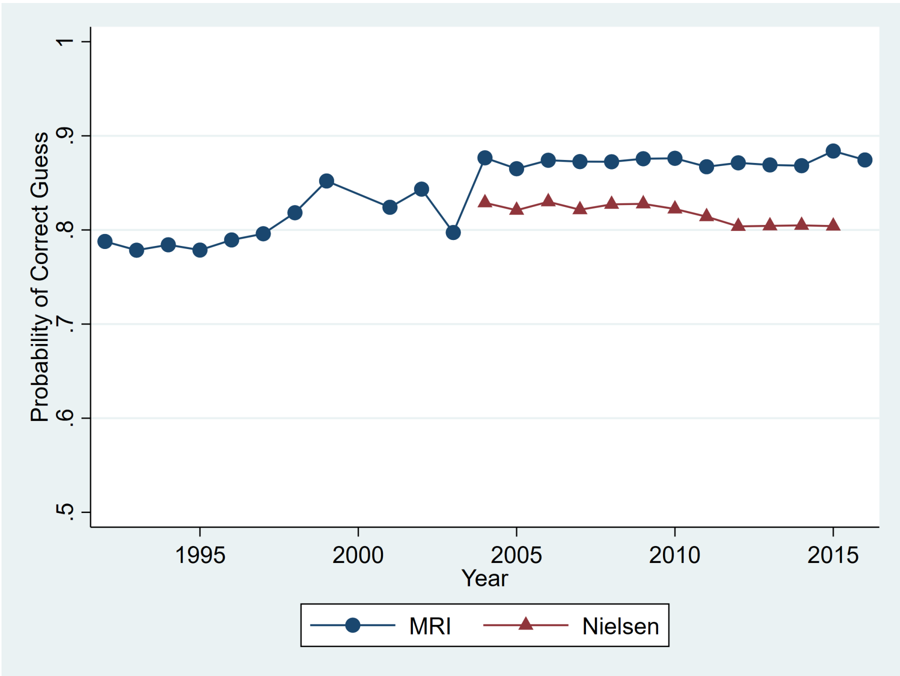

Figure B.10: Cultural distance by race over time: consumer behavior

Note: Data sources are the MRI and Nielsen. Sample sizes each year are 4,150 for MRI and 2,450 for Nielsen. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's race in the hold-out sample each year. The procedure to guess race in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these iterations.

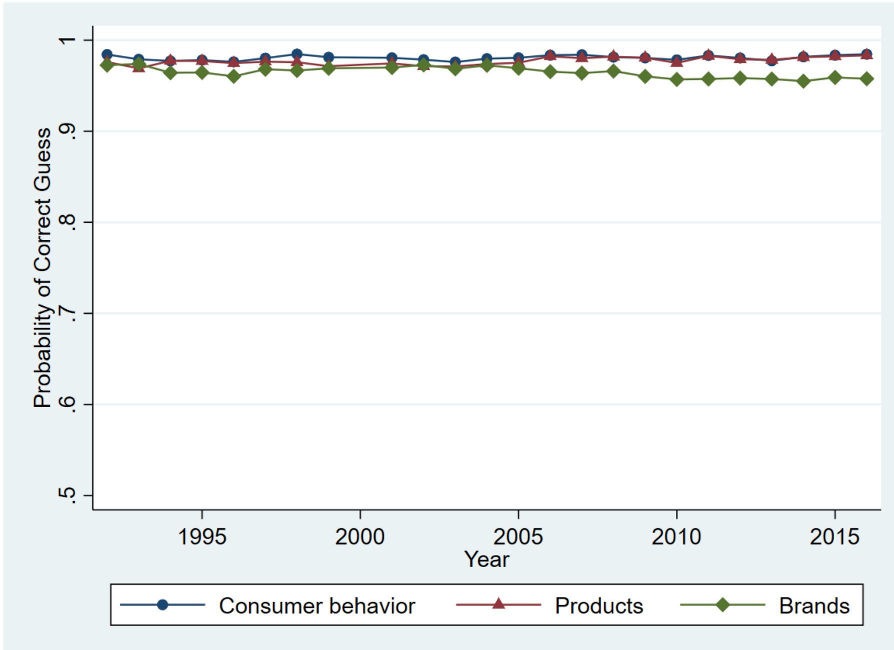

Figure B.11: Cultural distance by gender over time: consumer behavior

Note: Data source is the MRI. Sample size each year is 15,036. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's gender in the hold-out sample each year. The procedure to guess gender in the hold-out sample was repeated 5 times, and the share of guesses reported is the average of these 5 iterations. Products that allow us to guess gender correctly for over 75% of the time are dropped.

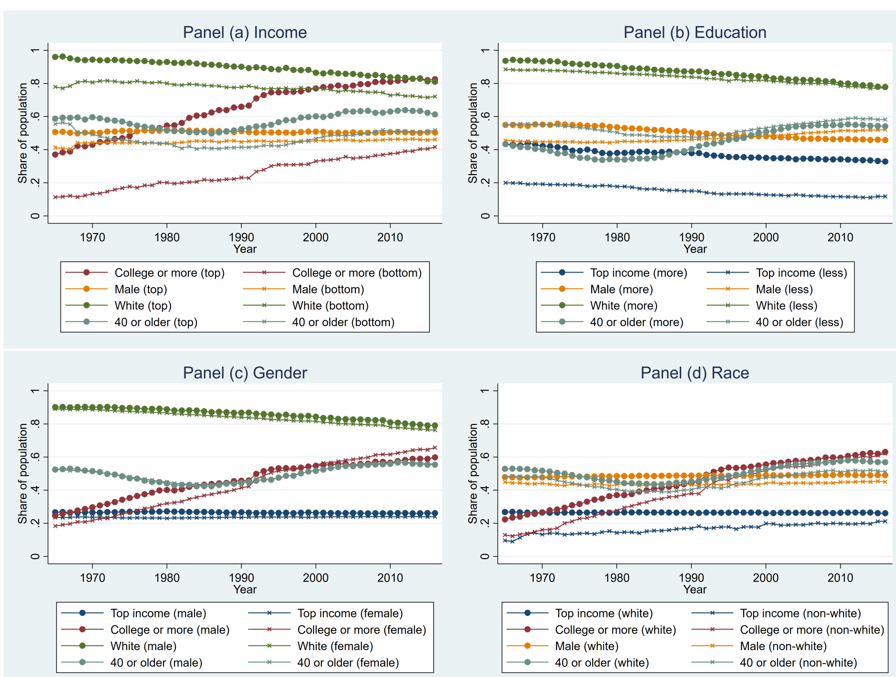

Figure B.12: Compositional changes in income, education, gender, and race

Note: Income defined by top vs. bottom quartile of household income by type.

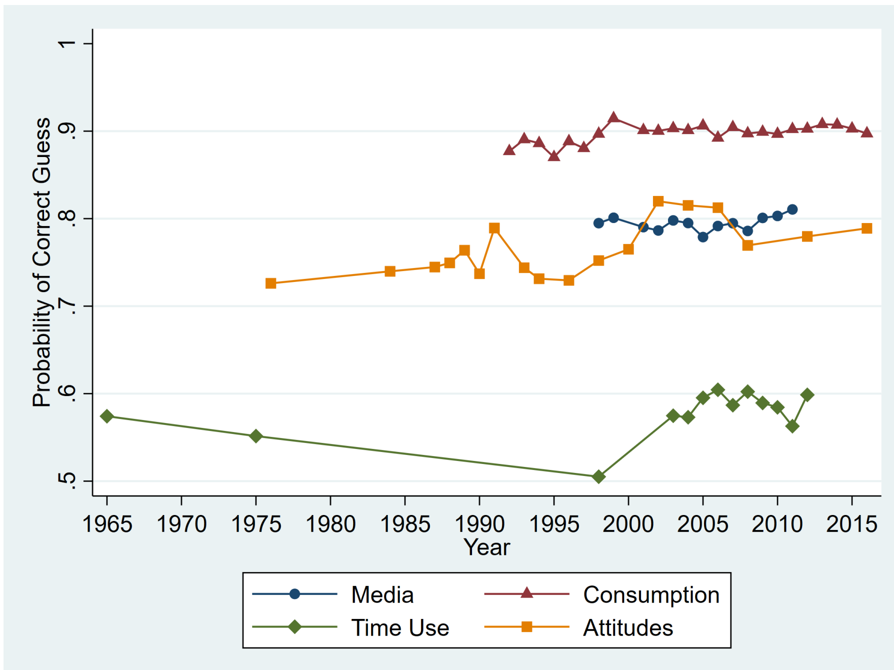

Figure B.13: Cultural distance by income controlling for age

Note: Data sources are the GSS, the AHTUS, and the MRI. Sample sizes each year are 6,472 for media and consumption, 268 for time use, and 322 for attitudes. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's income in the hold-out sample each year. The procedure to guess income in the hold-out sample was repeated 5 times for consumption, 25 times for media, and 500 times for time use and attitudes, and the share of guesses reported is the average of these iterations.

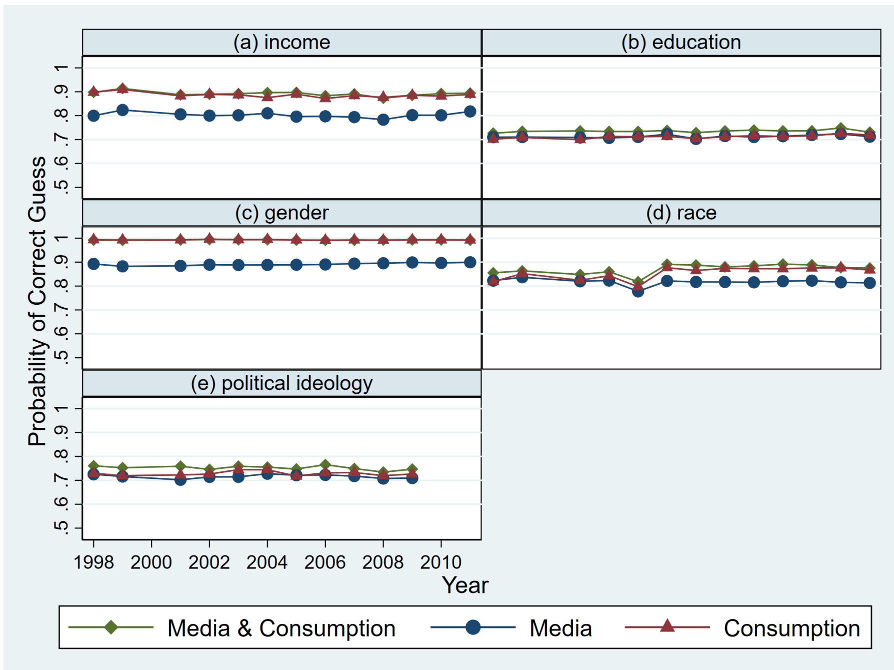

Figure B.14: Cultural distance in both media diet and consumer behavior

Note: Data source is the MRI. Sample size each year is 5,810 for income, 9,674 for education, 15,036 for gender, 4,150 for race, and 4,864 for political ideology. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's membership in a demographic group in the hold-out sample each year. The procedure to guess membership in the hold-out sample was repeated 5 times for consumption and 25 times for media, and the share of guesses reported is the average of these iterations.

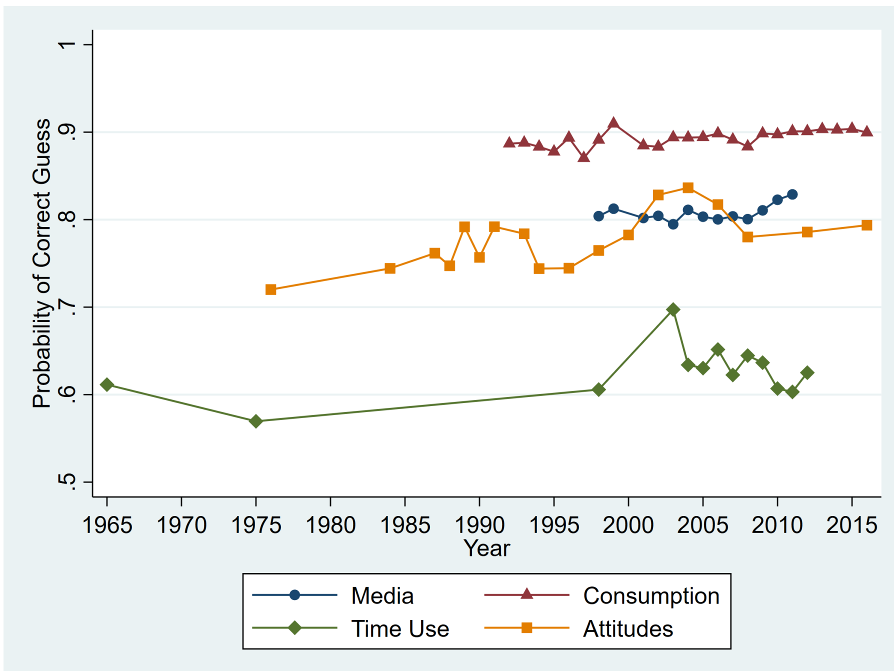

Figure B.15: Cultural distance by income, controlling for household size

Note: Data sources are the GSS, the AHTUS, and the MRI. Sample sizes each year are 5,970 for media and consumption, 422 for time use, and 386 for attitudes. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's income in the hold-out sample each year. The procedure to guess income in the hold-out sample was repeated 5 times for consumption, 25 times for media, and 500 times for time use and attitudes, and the share of guesses reported is the average of these iterations.

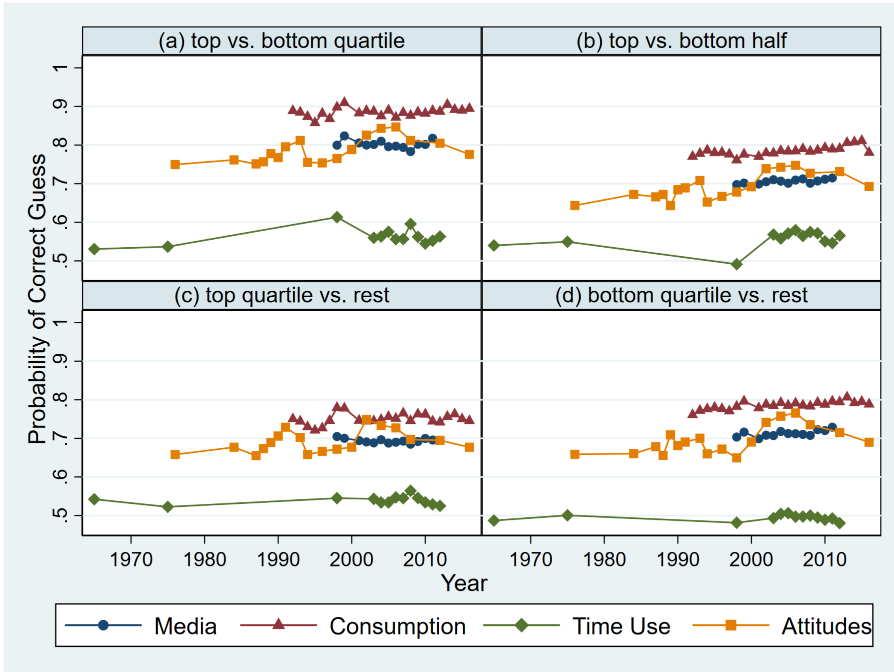

Figure B.16: Alternative income groups

Note: Figure shows the likelihood, in each year, of correctly guessing an individual's group membership based on his/her media diet, consumer behavior, time use, or social attitudes. Panel (a) is equivalent to panel (a) in 1. Panel (b) measures the cultural distance between the top half and the bottom half of the income distribution. Panel (c) measures the distance between top quartile and the rest (second, third, and fourth quartiles), and panel (d) measures the distance between the bottom quartile and the rest (first, second, and third quartiles). See text and data appendix for details on sample construction and implementation of machine-learning ensemble method.

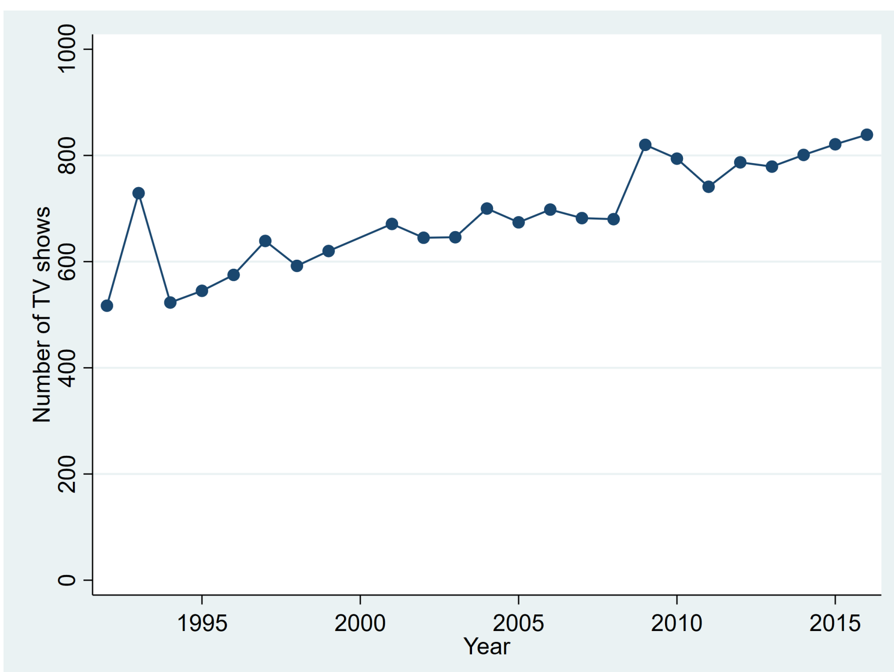

Figure B.17: Number of TV shows in the MRI data

Note: Data source is MRI. The increase in 2009 reflects addition of cable shows.

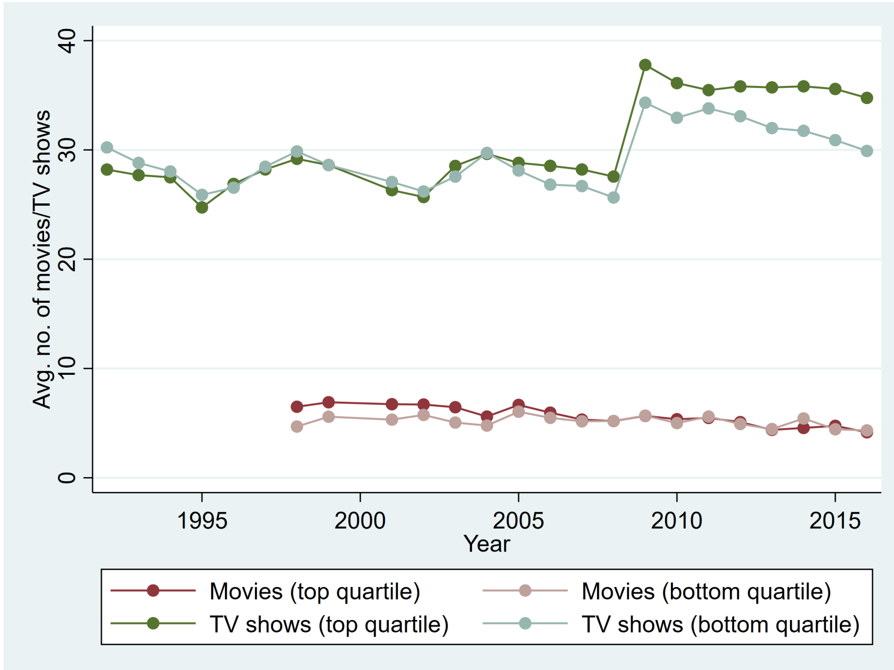

Figure B.18: Average no. of movies and TV shows watched by income in the MRI data

Note: Data source is MRI. The increase in 2009 reflects addition of cable shows.

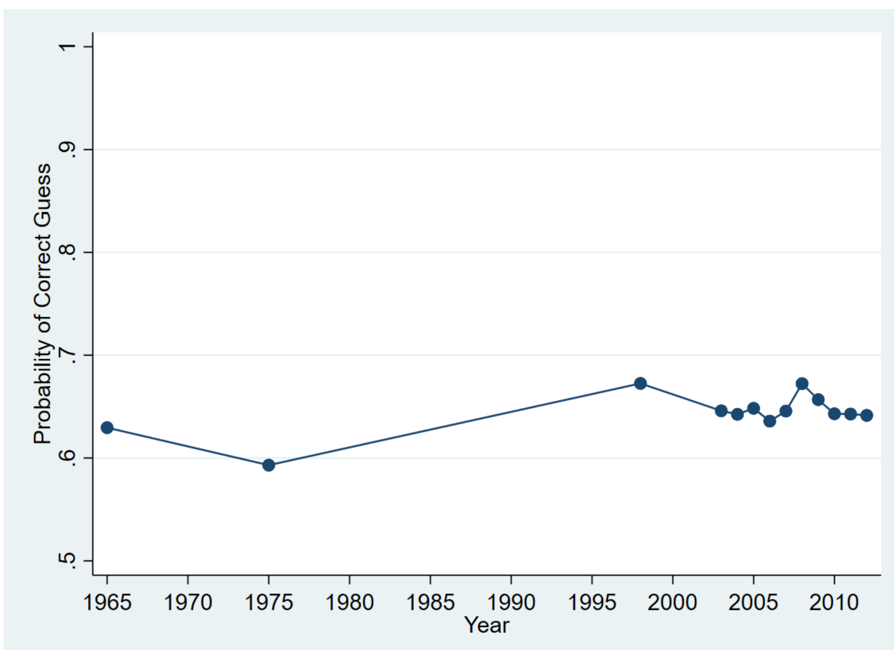

Figure B.19: Cultural distance by income in time use for the full sample

Note: Data source is the AHTUS. Sample size each year is 376. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's income in the hold-out sample each year. The procedure to guess income in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations.

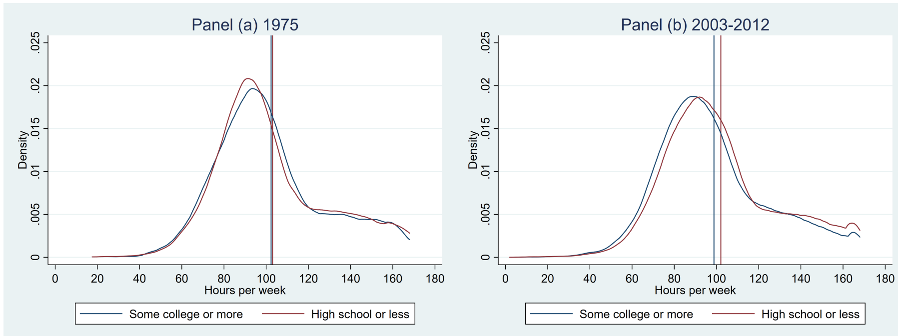

Figure B.20: Distribution of time spent on leisure by education level, 1975 vs. 2003-2012

Note: Data source is the AHTUS.

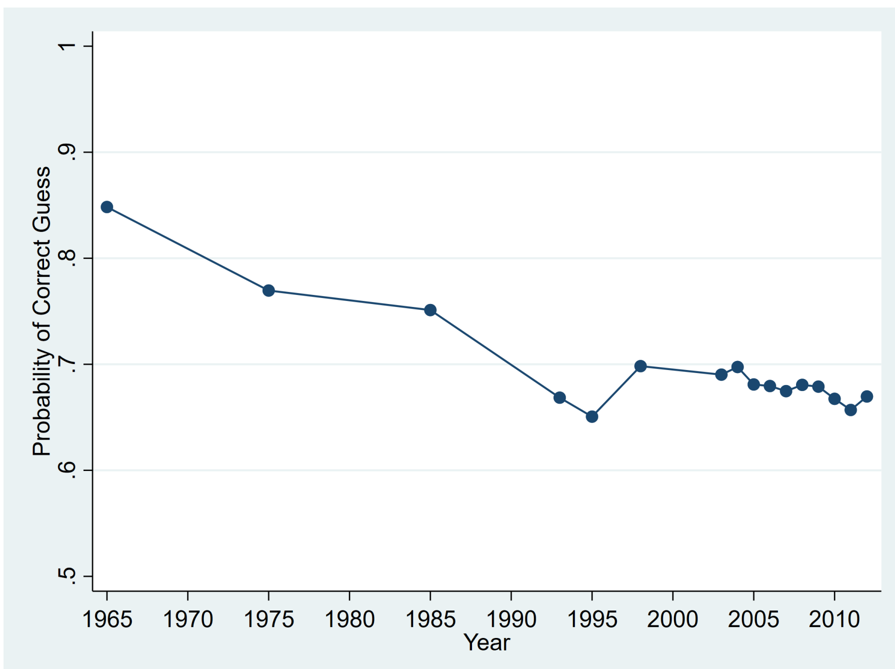

Figure B.21: Gender differences over time in allocation of non-work time

Note: Data source is the AHTUS. Sample size each year is 666. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's gender in the hold-out sample each year. The procedure to guess gender in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations.

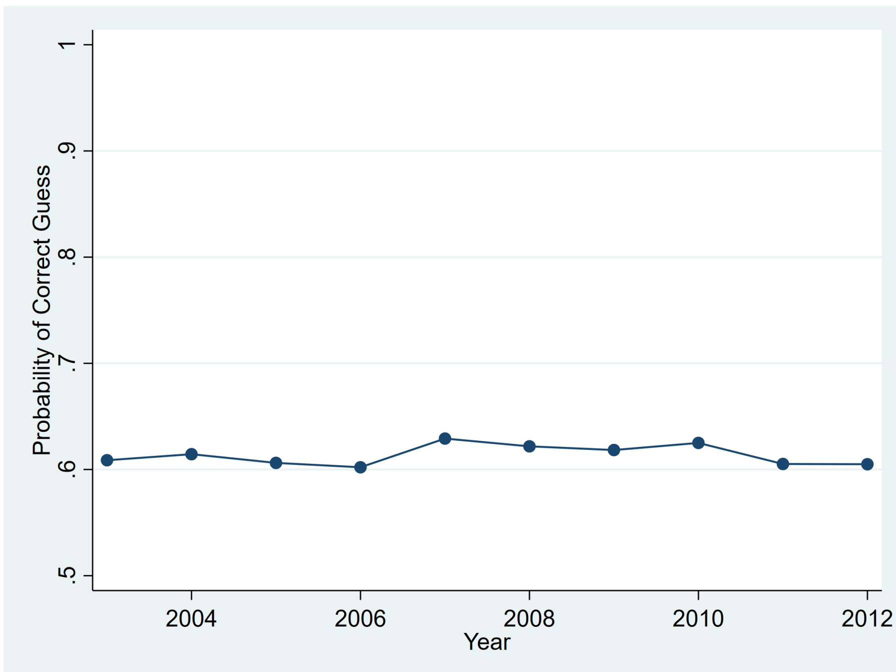

Figure B.22: Cultural distance by race in time use for the 2003-2012 sample

Note: Data source is the AHTUS. Sample size each year is 2,042. See text and data appendix for details on sample construction and implementation of machine-learning ensemble method. Presented in the figure is share of correct guesses of respondent's race in the hold-out sample each year. The procedure to guess race in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations.

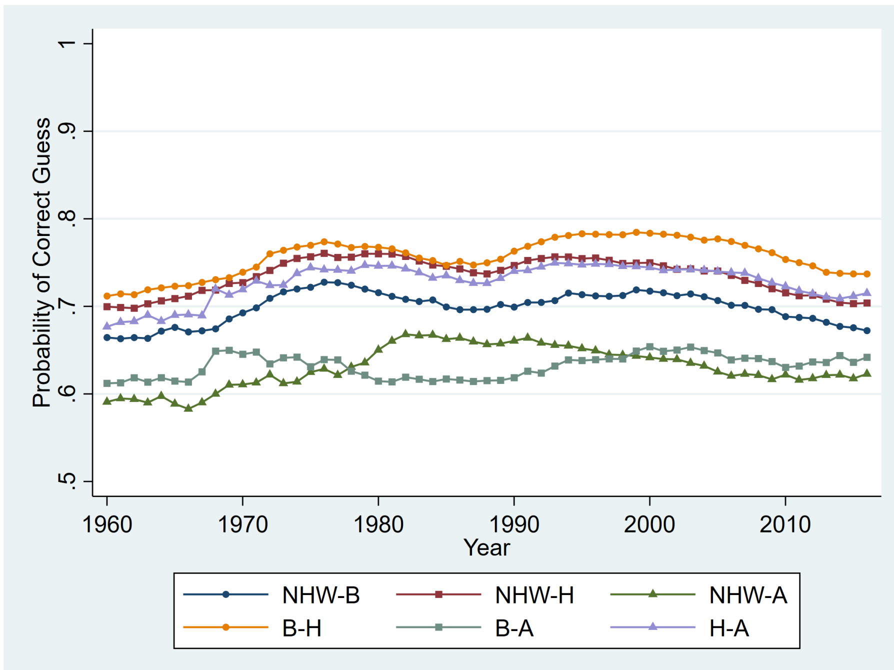

Figure B.23: Cultural distance by race and ethnicity over time (pairwise comparisons): boys' names

Data source is the CDPH. Sample size each year is 4,868. See text and data appendix for details on sample construction and implementation of the Bayesian method. Presented in the figure is share of correct guesses of mother's race in the hold-out sample each year. The procedure to guess race in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations. "NHW" denotes Non-Hispanic White, "B" denotes Black, "H" denotes Hispanic, and "A" denotes Asian.

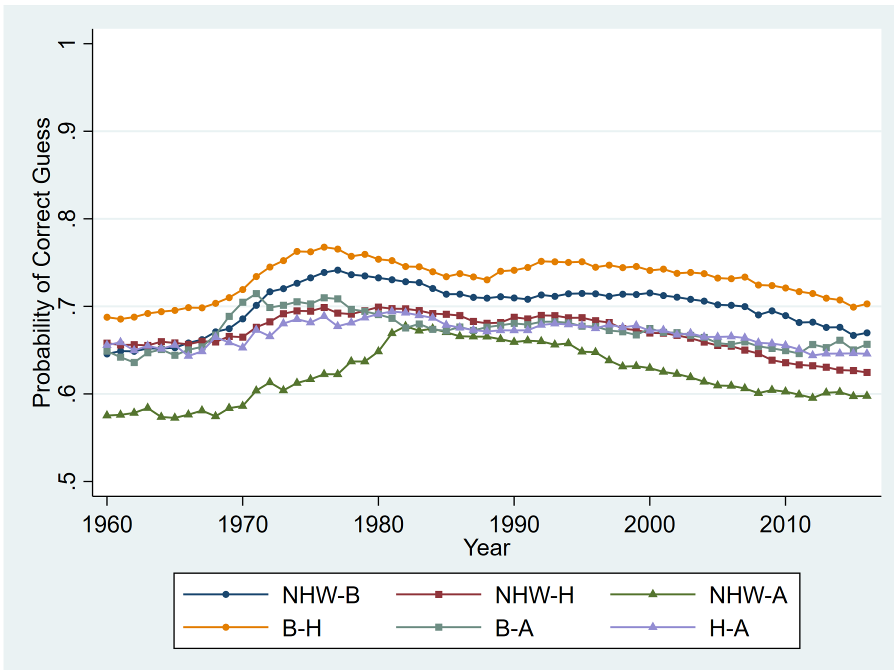

Figure B.24: Cultural distance by race and ethnicity over time (pairwise comparisons): girls' names

Data source is the CDPH. Sample size each year is 4,440. See text and data appendix for details on sample construction and implementation of the Bayesian method. Presented in the figure is share of correct guesses of mother's race in the hold-out sample each year. The procedure to guess race in the hold-out sample was repeated 500 times, and the share of guesses reported is the average of these 500 iterations. "NHW" denotes Non-Hispanic White, "B" denotes Black, "H" denotes Hispanic, and "A" denotes Asian.