

THE CULTURAL DIVIDE*

Klaus Desmet and Romain Wacziarg

We study the evolution of cultural divides both theoretically and empirically. We propose a model of cultural change where intergenerational transmission and forces of social influence determine the distribution of cultural traits in society. We conduct an empirical investigation of the evolution of cultural heterogeneity in the USA between 1972 and 2018, using the General Social Survey. In recent decades, cultural heterogeneity between individuals has risen. Cultural divides between identity groups display contrasted patterns: increasing along lines of religion and political orientation, mildly U-shaped along income and racial lines, and flat or decreasing for gender and urbanicity. We interpret these empirical findings in light of our model, arguing that changes in modes of interaction within and across groups can explain the observed dynamics of cultural heterogeneity.

Many scholars and commentators have argued that the USA faces a growing cultural divide along lines of race, religion, income and other dimensions. These increasing disagreements go hand-in-hand with a fraying social fabric, growing dysfunction in the political arena and the disintegration of social capital. Others have argued that greater information flows and exchange between different cultures have brought about cultural convergence, so that cultural differences between groups are becoming smaller as cultural traits diffuse throughout society.¹ Which view is correct?

In this article, we study the evolution of cultural divides both theoretically and empirically. First, we propose a model of cultural change that delivers predictions on the evolution of cultural divides as a function of various forces, including the intergenerational transmission of cultural traits, the propensity to conform, the diffusion of cultural innovations, and the sorting of individuals into different identity groups based on their cultural values. Second, we conduct an empirical analysis of cultural convergence and divergence in the USA over time. We assess whether cultural values—or *memes*—have grown more or less heterogeneous across groups defined according to six identity cleavages: gender, religion, race, income, urbanicity and political party.² Third, we interpret these empirical findings using our theoretical framework.

* Corresponding author: Romain Wacziarg, Anderson School of Management, UCLA, 110 Westwood Plaza, Los Angeles CA 90095, USA. Email wacziarg@ucla.edu

This paper was received on 28 February 2020 and accepted on 16 December 2020. The Editor was Ekaterina Zhuravskaya.

The data and codes for this paper are available on the Journal website. They were checked for their ability to reproduce the results presented in the paper.

We thank Omer Ali for outstanding research assistance. We also thank Alberto Alesina, Raquel Fernández, Kai Gehring, Paola Giuliano, Ricardo Perez-Truglia, Fabio Schiantarelli, Jesse Shapiro, Francesco Trebbi, Ekaterina Zhuravskaya (the editor), three anonymous referees and seminar participants at numerous universities for useful comments.

¹ On the first view, prominent enunciations include Putnam (2000) and Murray (2013). Commentary along these lines among pundits and journalists are too numerous to list. The second view is more closely associated with modernisation theory—see for instance Inglehart (1997) and Ritzer (2011), who point to rising incomes and globalisation as powerful forces for cultural homogenisation.

² The term ‘meme’ was coined by Richard Dawkins (1989) to describe a cultural trait, much like a gene is a genetic trait. A meme can take on several variants, for instance the meme ‘belief in God’ could take on variants ‘yes’ or ‘no’. A person’s culture, in our terminology, is simply that person’s vector of memetic traits.

To understand the factors that shape changes in cultural heterogeneity over time, our model features several forces. First, an individual's memes originate, with variation, from intergenerational transmission. Second, agents tend to conform to the majority memes of their own identity group. Third, cultural innovations spread through social influence. Fourth, under certain conditions individuals can change their identity rather than their cultural values. Together, these four forces determine the dynamics of cultural change. The model shows that an increase in the propensity to conform can lead to increasing divides between groups; that differences in the susceptibility of different groups to adopt cultural innovations can also lead to deepening divisions; and finally that the ability to sort into different groups can reinforce these dynamics.

Turning to the empirical analysis, we use the General Social Survey (GSS), a US survey of norms, values and attitudes, spanning 1972 to 2018. We consider a wide range of memes covering religious beliefs and practices, confidence in institutions, preferences over public policies, moral values and attitudes, measures of trust and life satisfaction, and tolerance for alternative viewpoints and lifestyles. We begin by calculating *overall heterogeneity*, the likelihood that two randomly chosen individuals differ in terms of their cultural traits. Next we compute measures of *cultural divides between groups*, capturing the degree of fixation of memetic traits onto group identity. A high degree of fixation indicates that memes are highly group-specific, while a low degree of fixation indicates that the distribution of memes within each group closely resembles that in society overall. Rising fixation, in this context, would be associated with a growing cultural divide between groups.

The picture that emerges from our empirical analysis is not one of a generalised deepening of cultural divisions. First, *overall heterogeneity* in the USA, when averaging across all available memes, fell between 1972 and 1993, and rose thereafter. On many issues, heterogeneity changed little. Second, the extent of the *cultural divides between groups* is small: most variation in memes is within groups. Third, we do not find a generalised increase in cultural divides over time. Rather, there are differences across cleavage types: for religion and political party ID, we find pronounced increases in cultural divides since 2000; for race and income, we find a mild U-shaped pattern over the 1972–2018 period; and for urbanicity and gender, we find that the divide declined or stayed flat, respectively. In sum, the data do not justify a sweeping conclusion that there are deepening cultural divides in the USA. Instead, our findings suggest a more qualified conclusion that cultural divisions have grown only since the late 1990s, only for some cleavages, and only for some memes.

We use our model as a lens through which we interpret our empirical results. We argue that forces of cultural change operate differently depending on specific identity cleavages and specific memes.

A first distinction is whether social interactions occur mostly within or across identity cleavages: an increase in between-group social interactions tends to mitigate divides, whereas an increase in within-group interactions tends to enhance such divides. We argue that new communication technologies may strengthen or weaken between-group interactions, depending on the identity cleavage. For example, they may help bridge the geographic divide between urban and rural areas. On the other hand, they may lead to echo-chamber effects with respect to political cleavages, by facilitating interactions between members of the same group. As such, our model may help us understand why the urban–rural divide may be narrowing, whereas the political divide may be widening.

A second distinction relates to the possible sorting of people with different memes into specific identities. Sorting has the potential to deepen divides. This may explain the more pronounced

divisions along political lines relative to those along racial lines that we observe in the data—as people can more easily change their political labels than their racial identity.

A third distinction concerns the manner in which different cultural values diffuse through society. The emergence of cultural innovations such as greater social acceptance for gay sexual relations or marijuana legalisation can lead to greater cultural divides if adopted at different rates across identity cleavages. In contrast, when it comes to free speech, where large majorities of every group are already in favour, a further increase in the share of people favouring free speech implies that between-group divisions are falling.

Our article is related to a growing literature on the evolution of cultural traits. Our terminology and overall approach to culture borrow from research on cultural evolution (Cavalli-Sforza and Feldman, 1981; Boyd and Richerson, 1988; Richerson and Boyd, 2004; Bell *et al.*, 2009; Henrich, 2015). Recent work by economists also tries to better understand the causes and mechanics of cultural change (Bisin and Verdier, 2000; Kuran and Sandholm, 2008; Olivier *et al.*, 2008; Fernández, 2014; Guiso *et al.*, 2016). Another literature, originating in political science and sociology, examines cultural change arising from modernisation and globalisation (Inglehart, 1997; Inglehart and Baker, 2000; Norris and Inglehart, 2009; Ritzer, 2011). Our work is also linked to wide-ranging scholarship on cultural change and persistence (Fernández and Fogli, 2006; Giuliano, 2007; Luttmer and Singhal, 2011; Alesina and Giuliano, 2015; Giavazzi *et al.*, 2019; Giuliano and Nunn, 2020).

Drawing on the aforementioned literature on the evolution of cultural traits, social scientists have also studied heterogeneity in cultural traits, which is our main focus here. One strand studies cultural and political polarisation in the USA (DiMaggio *et al.*, 1996; McCarty *et al.*, 2006, Fiorina and Abrams, 2008; Gentzkow *et al.*, 2016; Boxell *et al.*, 2017; Enke, 2020).³ Another strand focuses on the measurement of cultural heterogeneity at the individual-level rather than at the group level, using either genetic or memetic data (Ashraf and Galor, 2013; Desmet *et al.*, 2017).

An important recent contribution by Alesina *et al.* (2017) studies cultural heterogeneity in Europe using two waves of the World Values Survey. Like us, Alesina *et al.* (2017) are interested in characterising cultural convergence or divergence. However, their focus is on the evolution of cultural differences between European countries, using heterogeneity between US states as a point of comparison. Instead, we focus on the USA, consider a wide range of six identity cleavages, use a distinct measurement framework and interpret our findings through the lens of a model of cultural evolution.

Also relevant is Bertrand and Kamenica (2020), an article that was written contemporaneously with ours. Applying a machine learning algorithm to a variety of survey data, this article analyses how well someone's culture or consumption behaviour predicts their gender, race, income, education and political ideology. There are commonalities in our approaches: both articles study the time evolution of cultural divisions between groups; both employ data on values, norms and attitudes; both look at a range of identity cleavages that include income, race, gender and political orientation; and both find that political divides have deepened in recent years.⁴ However, the two approaches to studying the cultural divide are different along a number of dimensions. First,

³ Another fast-growing strand of research examines to what extent cultural and economic factors explain the recent rise in populism: Norris and Inglehart (2019) emphasise cultural factors, while Algan *et al.* (2017) and Autor *et al.* (2020) emphasise economic factors in the USA and Europe (Guriev and Papaioannou, 2020, provide a broad survey on the determinants of populism).

⁴ We also consider religion and urbanicity as identity cleavages, while they study educational attainment. In the working paper version of this study, we considered a broader range of 11 identity cleavages (Desmet and Wacziarg, 2018).

we propose a model of cultural change to help us interpret why some divides have deepened and others have not, whereas they characterise empirical regularities. Second, their measurement framework differs from ours: they use machine learning to quantify the extent to which culture and behaviour can predict a person's identity.⁵ In contrast, we calculate how predictive identity is for cultural values. Third, we emphasise that it is difficult to understand the dynamics of the cultural divide *between groups* without paying attention to the evolution of *overall* cultural heterogeneity in society. For instance, when certain values become more acceptable in society at large, they often diffuse at different rates in different groups, giving rise to a deeper divide. Fourth, Bertrand and Kamenica (2020) also use data on consumption behaviour, including media consumption, whereas our analysis considers only data on norms, values and attitudes. Finally, some results differ. They find that since 1995, with the exception of political ideology, the cultural divide has remained stable. Instead, we find that the divide for urbanicity has declined, while at the other end of the spectrum divides along religious and party ID lines have deepened since the year 2000.

1. A Model of Cultural Change

Our model builds upon ideas from the cultural evolution literature in both anthropology and, more recently, economics. Among the earliest contributions to model culture in an explicitly evolutionary context are Boyd and Richerson (1988) and Richerson and Boyd (2004; 2005). These authors propose a range of evolutionary mechanisms explaining the dynamic paths of cultural traits that evolve through mutation and selection, much like genes, but at very different rates, partly because unlike genes, cultural traits can be transmitted horizontally.⁶ Bisin and Verdier (2000) study the intergenerational transmission of norms in an explicitly economic model where parents rationally choose which traits to pass on to their children, to derive the degree of cultural heterogeneity of a stationary population.⁷ Bernheim (1994) models conformism, assuming that it arises from social influence: social status enters the utility function, so there is a penalty for not conforming. Such conformism can lead to persistent customs as well as temporary fads. Bikhchandani, Hirshleifer and Welch (1992) contains a theory of fads and culture whereby certain values can originate from small shocks to preferences and spread through local conformism, leading to informational cascades and cultural change. Kuran and Sandholm (2008) compare the dynamics of cultural evolution in isolated and integrated societies, by analysing the role of intergroup versus intragroup socialisation and coordination. The goal is to understand the conditions under which cultural integration occurs, and circumstances under which societies can retain their original cultures.

We build on all these contributions. We define culture as a vector of values, attitudes and norms pertaining to an individual. This allows us to emphasise the evolution of cultural diversity between and within identity groups.

⁵ Suppose the single most predictive value consistently allows the machine learning algorithm to correctly classify individuals into income groups 80% of the time. Then all other cultural values will have little or no impact on the Bertrand and Kamenica measure of cultural divisions, and therefore their methodology will miss convergence or divergence along possibly important dimensions of culture. In contrast, our approach evaluates cultural convergence and divergence by treating all the memes equally.

⁶ Genes and cultural traits can also coevolve. Henrich (2015) contains further explorations in a similar vein.

⁷ Doepke and Zilibotti (2008) also explicitly model parents' choices of values to impart to their children as a function of economic incentives. Lazear (1999) models an individual's choice to learn languages, gain familiarity with other cultures, and assimilate culturally, as a function of economic incentives to trade.

1.1. Indicators of Overall and Between-Group Heterogeneity

To capture cultural heterogeneity and the cultural divide between identity cleavages, we start from the measurement framework in Desmet *et al.* (2017).⁸ Consider $c = 1, \dots, C$ identity cleavages that each consist of groups $k_c = 1, \dots, K_c$. Consider also $m = 1, \dots, M$ memes that each can take on values $i_m = 1, \dots, I_m$. For instance, c could be gender ($k_c =$ male, female) and m could be belief in God ($i_m =$ yes, no). We denote by s^{i_m} the share of the total population that holds variant i_m of meme m , and by s_{k_c} the share of group k_c in the total population. We denote by $s_{k_c}^{i_m}$ the share of group k_c (defined over cleavage c) that holds variant i_m of meme m . For instance, this could be the share of males that believe in God.

Overall heterogeneity is simply memetic fractionalisation over the whole population. For meme m :

$$CF^m = 1 - \sum_{i_m=1}^{I_m} (s^{i_m})^2. \quad (1)$$

Averaging over memes, we get average memetic fractionalisation—the probability that two randomly chosen individuals from the entire sample hold a different variant of a randomly drawn memetic trait:

$$CF = \frac{1}{M} \sum_{m=1}^M CF^m = 1 - \frac{1}{M} \sum_{m=1}^M \sum_{i_m=1}^{I_m} (s^{i_m})^2. \quad (2)$$

CF is a measure of memetic heterogeneity in the entire population, regardless of identity cleavages.

To derive a measure of the cultural divide between groups, we calculate F_{ST} measures of memetic fixation. Heuristically, F_{ST} captures the share of heterogeneity that occurs between groups defined by identity cleavages (Wright, 1949; Cavalli-Sforza *et al.*, 1994; Desmet *et al.*, 2017). We start by defining heterogeneity in meme m within group k_c :

$$CF_{k_c}^m = 1 - \sum_{i_m=1}^{I_m} (s_{k_c}^{i_m})^2. \quad (3)$$

Taking the weighted average over groups for a given identity cleavage c , we obtain the average within-group heterogeneity for meme m , CF_c^m :

$$CF_c^m = \sum_{k_c=1}^{K_c} s_{k_c} CF_{k_c}^m. \quad (4)$$

Finally F_{ST} for meme m defined over cleavage c is simply the share of the total heterogeneity that is not attributable to within-group heterogeneity:

$$(F_{ST})_c^m = 1 - \frac{CF_c^m}{CF^m}. \quad (5)$$

$(F_{ST})_c^m$ takes on values between 0 and 1. When $(F_{ST})_c^m = 0$, group identity carries no information concerning an individual's cultural value. When $(F_{ST})_c^m = 1$, knowing a person's identity is

⁸ Beyond the common measurement framework, Desmet *et al.* (2017) deal only with ethnic divisions, and do so in a cross-section of countries. Instead, the present study covers the evolution of multiple cultural divides over time for a single country, the USA.

equivalent to knowing their value, i.e., the meme is perfectly fixated on groups.⁹ As was the case for CF^m , $(F_{ST})_c^m$ can be averaged over all memes m to obtain the expected cultural divide between groups defined over cleavage c .

F_{ST} is subject to a well-known theoretical drawback: it relies on an additive decomposition of total fractionalisation into within-group and between-group components, in a context where the sum of these measures is bounded above by 1. This introduces dependence between overall heterogeneity and between-group heterogeneity (Jost, 2008). An alternative measure of between-group heterogeneity not subject to this drawback is the χ^2 index used by Desmet *et al.* (2017).¹⁰ In that article, the correlation between F_{ST} and χ^2 is 98%, so that in practice it matters little which one is used. Due to its more widespread use in the literature and its intuitive interpretation, we focus here on F_{ST} .

We now turn to a model aimed at understanding the dynamic evolution of these main indicators of cultural heterogeneity.

1.2. Setup of the Model

Our model focuses on the simplified case of a society with one identity cleavage (e.g., gender) and one cultural meme (e.g., belief in God). In addition, the identity cleavage comprises only two identity traits k and $-k$ (e.g., male or female) and the cultural meme can take only two values i and $-i$. Time is discrete, $t = 1, 2, \dots$. Identity groups are of equal size, and for now we assume that an individual cannot choose her trait. Each agent has one child, so that each generation is as large as the previous one. Cultural values are imperfectly transmitted from parent to child. As an agent socialises, she may change her cultural value in two situations. First, if she was born with the minority value and is sensitive to conforming to her group's majority value, she may switch to the majority value. Second, we allow for the emergence of cultural innovations, meaning that one of the values becomes more socially acceptable. If an agent has a taste for adopting cultural innovations, she may switch to the value that has become more acceptable. Before stating an agent's decision problem, we describe in more detail the different determinants of his culture: vertical transmission, pressure to conform and the adoption of cultural innovations.

Vertical transmission and innate values. There is imperfect vertical transmission between a parent and a child. In particular, a share α of children inherits the value of their parent, and a share $(1 - \alpha)$ is born with the other value. The parameter α therefore measures the intensity of vertical transmission.¹¹ We refer to the value an agent is born with as his innate value. In the absence of conformism and innovation, the utility an agent derives from his innate value is normalised to one.

Pressure to conform. As an agent socialises, he may perceive a benefit from conforming to the majority value of his group. One benefit from conformism may be that agents who frequently interact gain from coordinating on the same value; another reason may be that some agents do not like to stand out by being different from their group's mainstream view. The benefit from conforming increases in the share of the own group that holds the majority value, but is heterogeneous across individuals. In what follows, we denote by s_k^i the share of group k that

⁹ $(F_{ST})_c^m = 1$ can only happen when the number of identity groups K_c is at least as large as the number of possible cultural values I_m , and there is no within-group heterogeneity in values.

¹⁰ χ^2 captures the information content of a person's identity in terms of that person's cultural values (Cover and Thomas, 2006). This index will take on high values when cultural traits are very group-specific.

¹¹ We do not endogenise α , in contrast to the approach in the classic paper by Bisin and Verdier (2000), where the intergenerational transmission of culture results from purposeful decisions by parents.

holds the majority view (and by s_k^{-i} the share that holds the minority value, where obviously $s_k^{-i} = 1 - s_k^i$). When born, an individual draws a random variable p from a uniform distribution with support $[0, 1/\bar{p}]$. The utility he gets from conforming to the majority value is then $\frac{1}{p}s_k^i$ if he was born with the minority value and $\frac{1+\gamma}{p}s_k^i$ if he was born with the majority value, where $\gamma \geq 0$ is a utility premium from having been born with the majority value. A higher \bar{p} indicates a higher expected level of intragroup conformism in society overall.

Our setup does not allow for intergroup conformism per se. However, when discussing comparative statics on \bar{p} , we will argue that a weakening of within-group conformism (a lower \bar{p}) can be interpreted as a strengthening of between-group conformism.

The diffusion of cultural innovations. We define a cultural innovation as an existing value that becomes socially more acceptable or fashionable. A cultural innovation is simply a label attached to a given value that makes that value more attractive to hold. Some agents may find it attractive to adopt this value, and switch from the non-innovating to the innovating value. Suppose that j is the innovating value. For an agent of group k , the benefit of holding the innovating value is increasing in the share of agents of group k that hold this value, but is heterogeneous across agents. When a cultural innovation occurs, each agent draws a random variable r from a uniform distribution with support $[0, 1/\bar{r}_k]$. This determines an agent’s utility from imitating the innovating value, $\frac{1}{r}s_k^j$. A higher \bar{r}_k indicates a higher expected level of sensitivity to imitating cultural innovations, i.e., a higher sensitivity to fads, fashions or social trends.

Cultural innovations diffuse within groups, but may evolve differently in the two groups if \bar{r}_k and \bar{r}_{-k} are very different from each other. We discuss below situations under which \bar{r}_k and \bar{r}_{-k} may be more or less similar to each other.

1.3. Decision Problem

We now analyse an agent’s value choice at a given time t . Denote by i the value held by the majority of the agent’s group and by j the value experiencing an innovation, where j could be equal or different from i . An agent born with value x in group k , after drawing variables p and r , decides which value x' to adopt by solving the following discrete choice problem:

$$u(x, k) = \max_{x' \in \{j, i, x\}} \left\{ I(x), I(i) \frac{1 + I(x)\gamma}{p} s_k^i, I(j) \frac{1}{r} s_k^j \right\}, \tag{6}$$

where

$$I(x) = \begin{cases} 1 & \text{if } x' = x \\ 0 & \text{otherwise} \end{cases}$$

$$I(i) = \begin{cases} 1 & \text{if } x' = i \\ 0 & \text{otherwise} \end{cases}$$

$$I(j) = \begin{cases} 1 & \text{if } x' = j \\ 0 & \text{otherwise} \end{cases} .$$

To give an example, consider someone born with the majority value in a society where there is a cultural innovation to the minority value. If she holds on to her majority value, she will get a utility equal to $\max\{\frac{1+\gamma}{p}s_k^i, 1\}$, whereas if she switches to the innovating value she will get utility $\frac{1}{r}s_k^j$.

Laws of motion. Since individuals do not always keep the value they are born with, we denote by $z_k^i(t)$ the share of people of group k born in period t with innate value i and by $s_k^i(t)$ the share of people of group k with value i after solving the discrete choice problem. Our assumption on the imperfect vertical transmission of values between a parent and a child implies that:

$$z_k^i(t + 1) = \alpha s_k^i(t) + (1 - \alpha)(1 - s_k^i(t)) = (2\alpha - 1)s_k^i(t) + (1 - \alpha). \tag{7}$$

Of course if $\alpha = 1$, vertical transmission is perfect so that $z_k^i(t + 1) = s_k^i(t)$.

To derive the laws of motion that determine cultural evolution, we solve the discrete choice problem (6), assuming that the random draws of p and r are independent. We start by analysing the case where the innovation occurs to the minority value $-i$. Depending on their draws of p and r , agents of group k born with value $-i$ may want to switch to value i . Similarly, agents of group k born with value i may prefer value $-i$. In the Theoretical Online Appendix, we show that the *ex ante* probability that an individual of group k born in period $t + 1$ with value $-i$ prefers value i is $\bar{p}s_k^i(t + 1) - \frac{1}{2}\bar{p}\bar{r}_k s_k^i(t + 1)(1 - s_k^i(t + 1))$. Because of the law of large numbers, this is the same as the share of agents of group k born with value $-i$ that switch to value i . As for individuals of group k born in period $t + 1$ with value i , the share that prefers to switch to value $-i$ is $\bar{r}_k(1 - s_k^i(t + 1)) - \frac{1}{2}\bar{r}_k\bar{p}(1 + \gamma)s_k^i(t + 1)(1 - s_k^i(t + 1))$. These results yield the following law of motion for the share of the population holding value i when the innovation occurs to value $-i$:

$$\begin{aligned} s_k^i(t + 1) &= z_k^i(t + 1) + \bar{p}s_k^i(t + 1)(1 - z_k^i(t + 1)) - \bar{r}_k(1 - s_k^i(t + 1))z_k^i(t + 1) \\ &\quad + \frac{1}{2}\bar{r}_k\bar{p}s_k^i(t + 1)(1 - s_k^i(t + 1))(1 + \gamma)z_k^i(t + 1) \\ &\quad - (1 - z_k^i(t + 1)) \text{ if } j = -i. \end{aligned} \tag{8}$$

Next we turn to the case where the innovation occurs to the majority value i . The share of individuals of group k born in period $t + 1$ with value $-i$ who prefer to switch to value i can be shown to be $\bar{p}s_i^k(t + 1) + \bar{r}_k s_i^k(t + 1) - \bar{r}_k\bar{p}(s_i^k(t + 1))^2$. The law of motion for the share of people holding value i when the innovation occurs to value i then becomes:

$$\begin{aligned} s_k^i(t + 1) &= z_k^i(t + 1) + \bar{p}s_k^i(t + 1)(1 - z_k^i(t + 1)) + \bar{r}_k s_k^i(t + 1)(1 - z_k^i(t + 1)) \\ &\quad - \bar{r}_k\bar{p}(s_k^i(t + 1))^2(1 - z_k^i(t + 1)) \text{ if } j = i. \end{aligned} \tag{9}$$

The above two laws of motion are difference equations that describe the evolution of the majority value. Of course, the two laws of motion of the minority value are the complements of the laws of motion of the majority value. The laws of motion of the other group $-k$ can be written down by analogy. The Theoretical Online Appendix gives further details. Note that if nobody is sensitive to cultural innovations (i.e., $\bar{r}_k = 0$), or if there is no conformism ($\bar{p} = 0$), then these difference equations simplify considerably and become linear.

Choice of identity trait. Until now we have assumed that agents cannot choose their identity trait. Of course, for some identity cleavages (e.g., party ID) an individual can freely choose identity trait k or $-k$. In that case, at a given time t , the discrete choice problem of an agent born with value x becomes:

$$u(x) = \max\{u(x, k), u(x, -k)\}, \tag{10}$$

where $u(x, k)$ and $u(x, -k)$ are the outcomes of maximisation problem (6) for an agent who, respectively, chooses identity trait k and $-k$. We postpone the discussion of the laws of motion under this scenario until Proposition 3.

1.4. Patterns of Cultural Evolution

In this section, we derive three propositions that describe salient patterns of cultural evolution generated by our model. In doing so, we do not mean to be exhaustive in representing all possible dynamics coming out of the model. Rather, we seek to illustrate the dynamics of cultural evolution under some salient assumptions about initial conditions and the main mechanisms of cultural change.

Conformism. We start by exploring a society with no diffusion of cultural innovation and with no choice of identity traits. We are interested in understanding how the steady-state value shares, and hence CF and F_{ST} , depend on the intensity of vertical transmission and the intensity of conformism. As we will now see, the results depend crucially on whether the majority value is the same across groups or not.

PROPOSITION 1. Conformism. *Consider a society with no diffusion of cultural innovations (i.e., $\bar{r}_k = 0$). Then, in steady state:*

- (1) *The majority share in each group is weakly increasing in the strength of vertical transmission (α) and conformism (\bar{p});*
- (2) *If the majority value is the same in both groups, F_{ST} is zero and CF is weakly decreasing in the strength of vertical transmission and conformism;*
- (3) *If the majority value is different in both groups, F_{ST} is weakly increasing in the strength of vertical transmission and conformism, and CF is maximised (and equal to 0.5).*

PROOF. See the Theoretical Online Appendix. □

This proposition is intuitive. The steady-state share of the majority value is increasing in the pressure to conform (\bar{p}) and in the strength of the intergenerational vertical transmission of values (α). With stronger pressure to conform, individuals have a greater incentive to switch to the majority value. As a result, the steady-state majority share becomes larger. With stronger intergenerational transmission of values, the constraint on how high the majority share can become is weakened. Taken together, there is less intragroup heterogeneity when \bar{p} is larger and/or α is larger.

By increasing the steady-state share of the majority value, larger values for \bar{p} and α reduce within-group cultural fractionalisation. If both groups conform to the same majority value, this also reduces overall cultural fractionalisation. Since, in that case, there are no differences between groups, F_{ST} is zero in steady state. If the two groups conform to different majority values, then a higher α and/or a higher \bar{p} leave the society's overall cultural fractionalisation unchanged, because the two groups are assumed to be of equal size. In this case, the cleavage between groups deepens, thus increasing F_{ST} .

How can we extend this discussion to a consideration of between-group conformism? Individuals from one group may be sensitive to the majority value of the other group. Of course, the importance of this force would depend on the importance of interactions between groups. For example, if the intensity of interactions between groups declines, individuals become less

sensitive to the majority view of the other group. If the majority values differ across groups, then in our interpretation becoming less sensitive to the other group is akin to becoming more sensitive to one's own group. This translates into an increase in \bar{p} , and hence a higher F_{ST} . If, on the other hand, the majority value is the same across groups, then allowing for intergroup conformism does not affect the steady-state cultural divide, since F_{ST} remains zero.

Proposition 1 has a simple corollary which states that if an exogenous shock switches the majority value of one of the groups, the cultural divide between groups will increase.

COROLLARY 1. *Switching of Majority Values.* *Consider a society with no diffusion of cultural innovations. Starting off in a steady state where both groups conform to the same majority value, assume the value of the majority switches in one of the two groups. In that case, society converges to a new steady state with higher F_{ST} and higher CF .*

This result is immediate. If initially both groups have the same majority value, their steady-state value shares are identical, so that F_{ST} is zero. Consider a shock that turns the majority value of one of the groups into the minority value. Irrespective of the magnitude of this initial shock, the steady-state value shares of that group will switch. For instance, if the two values had shares of 0.2 – 0.8 in both groups, these now switch to 0.8 – 0.2 in one of the two groups. As a result, the steady-state *aggregate* value shares are 1/2, so CF is maximised. Given that both groups now conform to different majority values, there is a growing divide between groups, so F_{ST} increases. This result can be applied to a situation where shifting circumstances disrupt the existing consensus enough to make the majority view change in one of the groups.

Cultural innovations. We now turn to analysing the diffusion of cultural innovations, while still assuming that individuals cannot choose their identity trait. We focus on a situation in which both groups start off holding the same majority value and where the innovation affects the minority value.¹²

PROPOSITION 2. *Diffusion of Cultural Innovations.* *Starting from a situation in which both groups have the same majority value and the same majority share, suppose an innovation occurs to the minority value.*

- (1) *If conformism is sufficiently weak and diffusion is sufficiently strong, the majority value switches in both groups. During the transition, CF exhibits a hump-shaped path.*
- (2) *If conformism is sufficiently strong and diffusion is sufficiently weak, the majority value stays the same in both groups. During the transition, CF increases.*

PROOF. See the Theoretical Online Appendix. □

Once again, this proposition is intuitive. If diffusion is strong, and hence \bar{r}_k and \bar{r}_{-k} are high, individuals have a strong propensity to adopt innovations. Fads diffuse easily, and eventually take over, becoming the new majority norm. As the original consensus breaks down, there is initially growing disagreement between individuals. However, as the old majority norm is replaced by a new majority norm, agreement between individuals once again increases. This translates into

¹² In practice, for many memes, the majority value is the same across groups, so focusing on the case where both groups have the same majority value is reasonable. The Theoretical Online Appendix analyses what happens if initially both groups hold different majority values. As for the cultural innovation, the more interesting case is when it occurs to the minority value. If, instead, it occurs to the majority value, then it simply reinforces the share of people holding the majority view.

a hump-shaped transition path for cultural fractionalisation. If cultural diffusion is weak in both groups, the cultural innovation increases CF . In both cases, if the strength of diffusion of a particular cultural innovation differs across groups, this will lead to a growing divide across groups since the steady state shares of each value will be different across groups, and F_{ST} will rise.

How can this proposition inform our understanding of intergroup cultural diffusion? In our model, the innovation affects the same value in both groups. However, the adoption pattern may be different in the two groups, to the extent that \bar{r}_k and \bar{r}_{-k} are different. For instance, if \bar{r}_k is low and \bar{r}_{-k} is high, then group k will be much less sensitive to the innovation than group $-k$. If interactions between groups are frequent and intense, the sensitivity to cultural innovations in the two groups is likely to be more similar.¹³ Hence we can interpret differences between \bar{r}_k and \bar{r}_{-k} as having effects on cultural diversity akin to those of interactions between groups. If \bar{r}_k and \bar{r}_{-k} are the same, cultural innovations are adopted to the same degree in both groups, leaving F_{ST} unchanged. Correspondingly, if the intensity of intragroup diffusion is different across groups, a cultural innovation will lead to a growing divide between groups. These insights are summarised in the following corollary.

COROLLARY 2. Differences in Intragroup Diffusion. *Starting off in a steady state where \bar{r}_k and \bar{r}_{-k} are different, if this difference becomes smaller, then F_{ST} falls.*

Choice of identity trait. We now let individuals choose their identity trait. Consider an individual born with the minority value in her identity group. In addition to holding on to the minority value in her group or adopting the majority value of her group, she now has one more option: she can also switch identity groups. This may be an attractive option if she is a conformist and her value is held by the majority in the other identity group. The following proposition summarises this insight.

PROPOSITION 3. Choice of Identity Trait. *In a society with no diffusion of cultural innovations where the majority value of one group is the minority value of the other, then as long as the majority shares are smaller than one,*

- (1) F_{ST} is larger if individuals can choose their identity trait than if individuals cannot choose their identity trait;
- (2) The greater the degree of conformism, the larger the difference in F_{ST} between a situation where individuals can choose their identity trait and one where they cannot.

PROOF. See the Theoretical Online Appendix. □

This proposition says that the cultural divide between groups increases if individuals can freely choose their identity trait. Moreover, the increase in the cultural divide is larger if within-group conformism is stronger. The intuition for these two results is straightforward. Take an individual who holds the minority value in the group she is born into. If it is costless to switch groups, then she would rather change to the group where her innate value is held by the majority, as opposed to changing her value. That is, if changing identity trait is free, then it is better to change identity trait than to change value. This leads to sorting of values along identity traits, and hence to a rising cultural divide between groups. The average payoff from sorting into the identity trait where one's innate value is held by the majority is especially high if within-group conformism

¹³ By interactions we mean communication, contact and cooperative exchange between groups, not unlike the meaning of 'contact' in Intergroup Contact Theory in social psychology (Allport, 1954).

is strong. Hence, the incentive to sort on the majority value is greater in societies where people care a lot about conforming to the group.¹⁴

In sum, our model predicts different evolutions of CF and F_{ST} as a function of two main sets of theoretical parameters: \bar{p} , the degree of intragroup conformism, and the difference between \bar{r}_k and \bar{r}_{-k} , which captures how similar the diffusion of cultural innovations is in the two groups. In addition, the possibility of choosing one's identity trait increases the cultural divide between groups, and more so if within-group conformism is high.

2. The Evolution of Cultural Divides in the USA

2.1. Data

Selection of questions. We use survey data from 32 waves of the GSS between 1972 and 2018. The universe of all GSS questions across all waves includes 6,108 fields, but many of these questions were asked only once, either in special modules of the GSS appearing only in a single wave, or as time-specific questions (e.g., about a given presidential election). The first filter that we apply is therefore to require that a given question be asked in at least two different waves, in order to obtain *some* time-series variation. This leaves us with 2,408 questions.

Among these, questions fall into various types. To capture a respondent's vector of memes, we need to consider the universe of questions that refer to values and attitudes. To this end, we classified each question into one of three types: (i) questions clearly about the respondent's attitudes and values, (ii) factual questions that can be considered to reflect the values of the respondent, for instance, 'do you have a gun in your home?' or 'how often do you attend religious services?', and (iii) Questions not related to the respondent's values, including those that relate to facts about other people (the respondent's spouse, parents, etc.), the respondent's education when younger, as well as identity or demographic questions. For the purpose of determining the set of cultural memes used in the analysis, we only retain questions of the first two types. This results in a set of 1,115 questions.

Baseline set of questions. Among these 1,115 questions, some were asked more or less continuously across all waves, while others were asked for only a small subset of waves. For our baseline exercise, we require as long a time series as possible over a common set of questions, to ensure the comparability of the measures of cultural diversity across time. The questions that are asked regularly in the GSS are also more likely to reflect important, salient societal and cultural issues—trust, life satisfaction, attitudes toward fundamental civil rights like freedom of speech, etc.

There is a trade-off: the higher the frequency over which the measures are computed, the smaller the set of common questions across successive observations. To achieve balance between these considerations, we group the survey data into either two-wave sets or five-year sets and keep questions that were asked at least once in each grouping (Online Appendix Table A1 displays these groupings). This amounts to keeping questions asked at least once every two waves, or at least once in any five-year period. In the end, we are left with 76 memes when requiring questions to be asked every two waves, and 96 memes when requiring questions to be asked at least once in each five-year interval. These questions are listed in the Online Appendix Table A2. We use the

¹⁴ For a recent contribution that also models sorting of individuals into different political parties, see Gennaioli and Tabellini (2019) and Gethin *et al.* (2019). Both papers study how a shock to the salience of redistributive issues versus identity-based issues may have led to a realignment of individuals across parties.

76 questions obtained from the two-wave groupings as our baseline set, since it provides higher frequency for the heterogeneity measures, i.e., 16 groupings computed from 32 waves.¹⁵ We use the expanded set of questions obtained from the five-year groupings for robustness checks presented in the Online Appendix (this gives nine time periods). It is important to emphasise that both the wording of each question and the number of possible answers in the GSS does not vary across time.

Question entry and exit. The analysis of cultural heterogeneity over questions that enter or exit the survey at a given point in time could also be interesting. Many of these fleeting questions are asked only episodically in special GSS modules devoted to deeper investigations of topical subjects. But some questions may also enter or exit the survey depending on the degree of social consensus. Of particular concern is the exit of questions for which a social consensus has developed, and the entry of questions that are characterised by emerging divides. Entry and exit of questions along those lines cause opposing biases on CF (it is hard to form priors on the direction of the bias on F_{ST}). Since our baseline exercise considers a common set of questions across time, it may ignore changes in dimensions of cultural heterogeneity that are particularly salient at different points in time.

To address these issues, we conduct a systematic analysis of question entry and exit. In an extension to our baseline exercise, we include questions that enter and exit in our measures of overall and cross-group heterogeneity, to assess the effect they have on the dynamics of the cultural divide. The analysis proceeds in two ways. First, we calculate the heterogeneity indices over the full set of 1,115 questions. Of course, the indices are based on sets of questions that vary greatly through time, so this exercise is the polar opposite of our baseline analysis based on a time-invariant set of questions. Second, we focus more specifically on questions that appeared repeatedly in the survey and then were permanently removed, and conversely questions that did not appear and then were consistently included.

To do so, we apply a simple algorithm: we divide the sample period into two subperiods (1972–1989, i.e., 16 waves, and 1990–2018, i.e., 16 waves). Next, we identify questions asked at least five times in the first subperiod and never in the second ('exit' questions). There are 21 such questions. For instance, a question on whether birth control information should be available (*pill*) is asked in five waves in the 1970s and early 1980s, and is then permanently dropped. Conversely, we identify questions never asked in the first subperiod, and asked at least five times in the second ('entry' questions). There are 96 such questions. For instance, a question about affirmative action in hiring and promoting women (*fehire*) appears first in 1996 and is asked in almost every wave thereafter.

We observe that in the set of questions on cultural values that are asked at least twice, the number of questions that enter or exit as defined above is rather limited (10.5% of the questions). Most questions appear only episodically, without a systematic pattern of sustained entry or exit. Having identified the set of questions that persistently enter and exit the survey, we can examine if their inclusion in our indices of average heterogeneity affects the dynamics we describe.

Question categories. We calculate our indices of heterogeneity question by question. However, to report the results it is useful to average the indices across all questions or across subsets of these questions. To create these subsets, we use factor analysis and let the data inform us on sets of questions for which respondents tend to give similar answers. For instance, to the extent respondents tend to answer questions on free speech similarly, factor analysis will lead us to

¹⁵ Of these 76 questions, 64 are unambiguously about values and attitudes, while 12 are factual questions that we classify as reflecting the respondent's values, such as those on gun ownership or church attendance.

create a subset of questions that contain all the free speech questions.¹⁶ The resulting grouping of questions into distinct categories helps capture the fact that culture is a system of interrelated values, norms, and attitudes, and that as one meme changes, this may trigger changes in other values as well.

To describe more formally how we proceed, we conduct factor analysis on 76 questions across all 64,814 individuals who answered the GSS between 1972 and 2018. We retain seven factors that explain 45.8% of the variance in answers. Each one of these factors loads heavily on a subset of questions: factor 1 loads heavily on questions related to free speech and personal freedoms, such as marijuana legalisation and homosexual sex; factor 2 on questions related to abortion and church attendance; factor 3 on confidence in various institutions; factor 4 on gun ownership; factor 5 on life satisfaction; factor 6 on government spending priorities; and factor 7 on questions related to attitudes toward the police. We use factor analysis as a way of grouping questions into different categories. More specifically, for each factor we identify the questions with factor loadings above 0.3 in absolute value. That gives us seven groups of questions, over which we compute average CF and average F_{ST} .¹⁷

Question types. Questions are either binary or answered on a scale. In our baseline set of 76 questions, 26 are binary (yes/no, agree/disagree) and 50 admit answers that can be ordered on a scale. In 35 cases, the scale admits three answers, and in the remainder, four or more possible answers.

Identity cleavages. We consider six identity cleavages to compute the F_{ST} indices.¹⁸ These cleavages are characteristics of the respondents also observed in the GSS survey waves. They are labelled family income, gender, party ID, race, religion and urbanicity. These cleavages admit anywhere from two values (gender) to seven values (party ID), with the modal number of categories equal to five. Comparisons of the level of cultural divisions across cleavages are more reliable when the number of categories is relatively similar. However, most of our focus is on the dynamics of the measures of heterogeneity through time, not on comparisons of their average levels across cleavages. Online Appendix Table A3 displays the cleavages and corresponding categories.

2.2. The Evolution of Overall Heterogeneity

Figure 1 displays the time path of CF , averaged over the 76 baseline questions, and the second column of Table 1 show the underlying numbers. We find that average CF varies between 0.482 (in 1993) and 0.500 (1976).¹⁹ There is a U-shaped pattern over the sample period: overall

¹⁶ An alternative would be to rely on preexisting question categories and subcategories provided by the GSS. Broad categories include civil liberties, current affairs, gender & marriage, politics, and religion & spirituality. These are further divided into finer subcategories. For instance, gender and marriage includes questions on children and working, on marriage, and on sex and sexual orientation. It turns out there is significant overlap between the data-driven bunching delivered by factor analysis, and the pre-defined categories provided by the GSS. Results based on the GSS classification appear in the Online Appendix.

¹⁷ Online Appendix Table A2 lists the questions and the factor to which each of them is assigned. 59 questions, or 78%, are classified as part of one of the factors. As an alternative, we could do our analysis directly on the factor scores. However, that would involve arbitrarily discretising the factor scores in order to compute our measures of heterogeneity (see Desmet *et al.*, 2017).

¹⁸ The working paper version of this study considered an additional five cleavages: ethnicity, region, age, education and work status (Desmet and Wacziarg, 2018).

¹⁹ To place these numbers in a broader, cross-country context, we refer to results in Desmet *et al.* (2017). In that paper, using cross-sectional data from the World Values Survey (WVS), the USA has an average CF of 0.551, about half a standard deviation above the World mean of 0.529.




Fig. 1. Cultural Heterogeneity over Time (76 questions, Two-wave Grouping).

Table 1. Cultural Heterogeneity by Time Period and Cultural Divides by Time Period and Cleavage (Two-wave Grouping, 1972–2018 Waves).

Time period	Cultural heterogeneity	Family income divide	Gender divide	Party ID divide	Race divide	Religion divide	Urbanicity divide
1972	0.495	0.020	0.006	0.016	0.012	0.024	0.015
1974	0.494	0.017	0.006	0.013	0.009	0.022	0.016
1976	0.500	0.017	0.006	0.011	0.009	0.019	0.012
1979	0.499	0.018	0.006	0.011	0.010	0.022	0.013
1982	0.497	0.020	0.006	0.014	0.015	0.019	0.013
1984	0.496	0.018	0.007	0.013	0.010	0.020	0.011
1986	0.494	0.020	0.007	0.012	0.015	0.018	0.011
1988	0.486	0.018	0.007	0.010	0.008	0.019	0.010
1990	0.484	0.016	0.009	0.010	0.009	0.019	0.011
1993	0.482	0.015	0.006	0.011	0.009	0.018	0.009
1997	0.489	0.014	0.006	0.013	0.009	0.020	0.009
2001	0.491	0.017	0.007	0.013	0.011	0.021	0.008
2005	0.495	0.018	0.006	0.019	0.015	0.023	0.009
2009	0.499	0.019	0.006	0.018	0.012	0.024	0.009
2013	0.496	0.019	0.005	0.021	0.012	0.025	0.008
2016	0.489	0.016	0.005	0.023	0.013	0.026	0.008
Average	0.493	0.018	0.006	0.014	0.011	0.021	0.011

Notes: Time period refers to two-wave groupings. So for instance 1972 refers to pooled data over the 1972 and 1973 waves of the GSS. Cultural heterogeneity is measured by CF and cultural divides are measured by F_{ST} .

heterogeneity declined between the early 1970s and the mid-1990s and grew back to its initial level by the end of the period.²⁰ This average over all questions masks some underlying heterogeneity.

²⁰ The overall variation can reflect a substantial change in the underlying shares of respondents giving a specific answer to a question. For instance, consider a binary question. With a CF of 0.5, response shares would be equally

Table 2. *Dynamics of Cultural Heterogeneity and Cultural Divides by Cleavage (Two-wave Grouping, 1972–2018, 76 Questions).*

	U-shaped	Hump-shaped	Increasing	Decreasing	Flat
<i>Panel A: Cultural heterogeneity</i>					
Cultural heterogeneity	11.84%	10.53%	25.00%	28.95%	21.05%
<i>Panel B: Cultural divides by cleavage</i>					
Family income	3.95%	1.32%	15.79%	27.63%	51.32%
Gender	5.26%	10.53%	13.16%	14.47%	55.26%
Party ID	46.05%	2.63%	7.89%	7.89%	34.21%
Race	14.47%	1.32%	25.00%	15.79%	40.79%
Religion	26.32%	0.00%	10.53%	17.11%	44.74%
Urbanicity	10.53%	2.63%	6.58%	34.21%	46.05%
Panel B average	17.76%	3.07%	13.16%	19.52%	45.39%

Notes: This table displays the fraction of questions, among the 76 in our baseline sample, for which CF or F_{ST} follows the types of dynamics listed in the first row, i.e., U-shaped, hump shaped, increasing, decreasing or flat. To assess these dynamics, we regress for each question its CF/F_{ST} on a time trend and its square. If both the linear and quadratic terms are statistically significant at the 5% level, and the vertex of the fitted quadratic curve is between 1980 and 2005, we characterise the dynamics as either U-shaped or hump-shaped. In all other cases, we run a linear regression of CF/F_{ST} on a time trend, and classify the dynamics as increasing, decreasing or flat depending on whether the coefficient on the time trend is significantly positive, significantly negative, or insignificant, respectively.

Table 3. *Dynamics of Cultural Heterogeneity, by Question Category (Two-wave Grouping, 1972–2018, 76 Questions).*

	# of Questions	Type of Dynamics
Factor 1—Free speech and personal freedoms	14	Decreasing
Factor 2—Abortion and church attendance	10	Increasing
Factor 3—Confidence in institutions	12	Flat
Factor 4—Gun ownership	4	Decreasing
Factor 5—Life satisfaction	8	Increasing
Factor 6—Government spending priorities	8	Flat
Factor 7—Attitudes towards the police	3	Increasing

Notes: This table displays the types of dynamics of CF for different categories of questions, as determined by factor analysis. The possible types of dynamics are U-shaped, hump shaped, increasing, decreasing or flat. To assess these dynamics, we first averaged CF across questions within each question category. For each of the resulting averages, we regressed average CF on a time trend and its square. If the quadratic term is statistically significant at the 5% level, and the vertex of the fitted quadratic curve is between 1980 and 2005, we characterise the dynamics as either U-shaped or hump-shaped. In all other cases, we run a linear regression of CF on a time trend, and classify the dynamics as increasing, decreasing or flat depending on whether the coefficient on the time trend is positive, negative, or insignificant, respectively. We display these dynamics by reporting the type of dynamics observed for each category of questions.

Panel A of Table 2 breaks down the dynamics by question. We find that 11.8% of the questions display a significant U-shaped pattern. Heterogeneity is declining for 29% of the questions and increasing for 25% of them. The rest is either hump-shaped or flat. These contrasted dynamics across questions are not consistent with a generalised increase in the level of disagreement over cultural issues.

Table 3 characterises the dynamics of CF by question category, as determined by factor analysis. We find a lot of variation in the dynamics of CF across question categories. Heterogeneity is increasing for questions relating to abortion and church attendance, life satisfaction and attitudes

divided between both possible answers. Then a change in CF to 0.482 represents a shift in answer shares of 9.5% (shares of 40.5% to 59.5% for each possible answer). More generally, given the specific distribution of the number of possible answers among our baseline set of 76 questions, the theoretical maximum average level of CF is 0.63. The USA appears to be quite culturally diverse overall, but there is room for that diversity to grow.

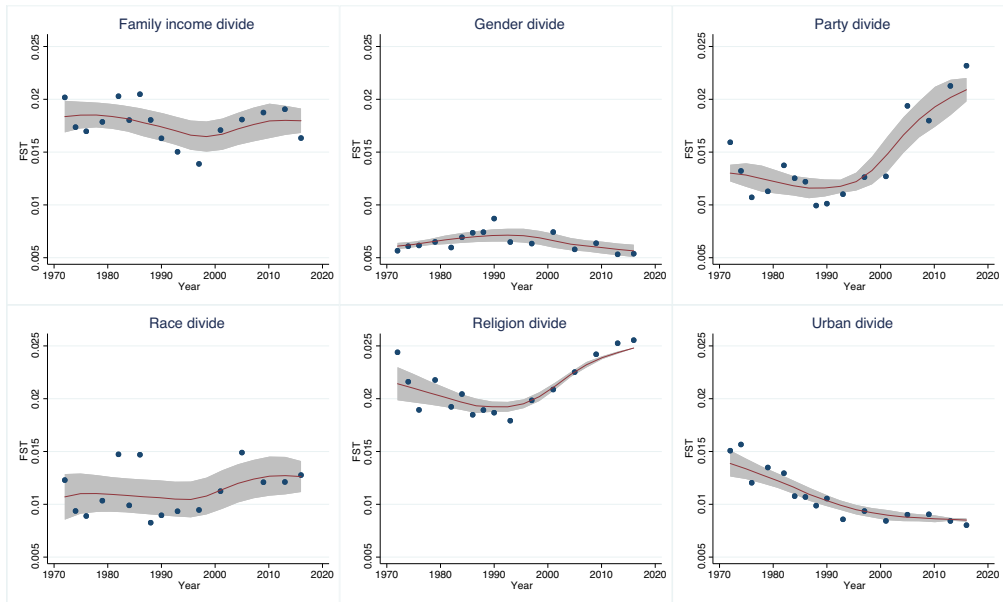


Fig. 2. Evolution of the Cultural Divides for Each of Six Cleavages, over Time.

toward the police. On the other hand, heterogeneity is decreasing when it comes to free speech (with respondents increasingly agreeing with a pro-free speech stance) and gun ownership (with an increase in the majority share that does not own guns). For government spending priorities and confidence in institutions, we fail to uncover significant trends in either direction.²¹

2.3. The Evolution of Cultural Divides

Types of dynamics by cleavage. For each of the 6 cleavages, Table 1 lists F_{ST} at different points in time and Figure 2 displays these data graphically. In both cases, F_{ST} is averaged over all 76 questions. The level of the divides is generally low, with F_{ST} ranging between 0.5% (for gender in 2016–2018) and 2.6% (for religion in 2016–2018).²² The high level of cultural pluralism in the USA, then, is not primarily due to diversity *between* identity cleavages, but mostly due to diversity *within* identity categories. Of course, even a small level of cultural division could have large effects on political economy outcomes (Desmet *et al.*, 2017).²³ On average across the sample period, we find that the level of the cultural divide is highest for family income and religion and lowest for gender. The level of the divide is at an intermediate level for race, party

²¹ Online Appendix Table A4 carries out the same breakdown using the GSS-defined question categories and sub-categories. The results broadly confirm our observations about the heterogeneity in the dynamics of CF across question categories based on factor analysis.

²² Random errors in answers to questions would tend to cause an upward bias in CF and a downward bias in F_{ST} . Take the extreme case of everyone in the sample answering a question randomly—then CF would be maximised and F_{ST} would be zero. Simple simulations show that this intuition holds in less extreme cases. However, these biases in the levels of F_{ST} and CF do not qualitatively affect the shape and direction of their time paths.

²³ In the same spirit, Tesei (2017) shows that while between-race income inequality is a small share of inequality in the USA, it is strongly negatively associated with generalised trust.

Table 4. *Dynamics of Cultural Divides, by Question Category and by Cleavage (Two-wave Grouping, 1972–2018, 76 Questions).*

	Family income	Gender	Party ID	Race	Religion	Urbanicity
Factor 1—Free speech and personal freedoms	Decreasing	Hump-shaped	U-Shaped	Increasing	U-Shaped	Decreasing
Factor 2—Abortion and church attendance	Decreasing	Flat	U-Shaped	U-Shaped	Increasing	Decreasing
Factor 3—Confidence in institutions	Flat	Flat	Flat	Flat	U-Shaped	Flat
Factor 4—Gun ownership	Increasing	Hump-shaped	U-Shaped	Increasing	Decreasing	U-Shaped
Factor 5—Life satisfaction	Increasing	Flat	Flat	Decreasing	Decreasing	Decreasing
Factor 6—Government spending priorities	Flat	Increasing	Increasing	Flat	U-Shaped	Decreasing
Factor 7—Attitudes towards the police	Increasing	Hump-shaped	Increasing	U-Shaped	Flat	Flat

Notes: This table displays the types of dynamics of F_{ST} for different categories of questions, as determined by factor analysis. The types of dynamics are: U-shaped, hump shaped, increasing, decreasing or flat. To assess these dynamics, we first average F_{ST} values across questions within categories, as well as across all six cleavages. We then regress these average F_{ST} 's on a time trend and its square. If the quadratic term is statistically significant at the 5% level, and the vertex of the fitted quadratic curve is between 1980 and 2005, we characterise the dynamics as either U-shaped or hump-shaped. In all other cases, we run a linear regression of F_{ST} on a time trend, and classify the dynamics as increasing, decreasing or flat depending on whether the coefficient on the time trend is positive, negative, or insignificant, respectively.

ID and urbanicity. It is, however, important not to overstate these level differences, as they are affected by the number of groups characterising each cleavage.²⁴

Figure 2 shows that these levels of cultural fixation change substantially through time, in ways specific to each cleavage. There are three categories of cleavages: (i) Cleavages with increasing time-paths: average F_{ST} for party ID is relatively flat through the mid-1990s, starts to gradually increase in the late 1990s, and then accelerates in the 2000s, reaching its maximum in 2016–2018. A similar pattern is found for religion. (ii) Cleavages with mildly U-shaped time-paths: family income and race. (iii) Cleavages with flat or declining levels of cultural fixation: gender and urbanicity. The latter patterns are not in line with many commentators' priors on rising divides across urban categories and genders.

We also find substantial heterogeneity in the dynamics of the cultural divides across questions: for each cleavage, Table 2 Panel B classifies the types of dynamics of F_{ST} across questions. For instance, F_{ST} based on party ID displays a significant U-shaped pattern over the sample period for 46.1% of the questions, with an additional 7.9% of the questions displaying a strictly increasing trend. The remainder display mostly flat dynamics. As another example, take urbanicity: F_{ST} is falling for 34.2% of the questions, but it is flat for 46.1% of the questions. This heterogeneity casts doubt upon sweeping pronouncements about rising cultural divides in the USA.

Types of dynamics by question category and cleavage. Table 4 classifies the types of dynamics by question category and cleavage. The breakdown by cleavage confirms our previous observations (mostly U-shaped or increasing for party ID and religion, flat or decreasing for gender and urbanicity, and weakly U-shaped or flat for family income and race). The breakdown

²⁴ It is possible to partial out the effect of the number of groups. Stacking our data on F_{ST} across 76 questions, 16 time periods and six cleavages (7,296 data points), we regressed F_{ST} on a variable capturing the number of groups per cleavage (that variable is constant within cleavages). We then averaged the residuals from this regression across questions, for each cleavage and each period. Plotting these partialled-out F_{ST} s for each cleavage across time, we find that differences in levels are compressed: the curve for gender moves up while the curve for party ID shifts down. We also find that correcting for the effect of the number of groups preserves the ranking of the level of F_{ST} across cleavages.

by question category reveals a high degree of heterogeneity. For questions on abortion and church attendance, the rural–urban and income divides are declining, but the party ID and religion divides are increasing. On the other hand, the time paths of divisions over confidence in government are more consistent across cleavages (mostly flat).²⁵

2.4. Further Exploration and Robustness

The distribution of F_{ST} across questions. Online Appendix Figure A1 graphically displays the distribution of F_{ST} across questions for each cleavage and for different wave groupings. Three observations are in order. First, the modal level of F_{ST} across cleavages is consistent with the averages discussed above—for instance the F_{ST} for gender tends to be low across all questions and the F_{ST} for religion tends to be higher. Second, for each cleavage, most of the density is concentrated tightly around the mode. The F_{ST} rarely exceeds 0.1, indicating that for most cleavage-meme pairings, within-group variation is a very large share of total cultural variation. A rare exception is for the religious cleavage, where a single question registers a high level of F_{ST} . That specific question is *fund* ('how fundamentalist is the respondent currently'—in terms of religion). This question is an outlier, and not a particularly surprising one given the close link between the meme (fundamentalism) and the cleavage (religion). A detailed look at this question reveals that Catholics tend to answer 'moderate', Jews and those with no religion tend to answer 'liberal' and Protestants tend to be divided between 'fundamentalist' (about 50%), liberal (about 25%) and moderate (about 25%).²⁶ Third, the time evolution of the distributions confirms our previous findings. For example, the distribution of F_{ST} for the political party ID cleavage has been moving rightward since the beginning of the 2000s.

The most divisive questions. Online Appendix Table A6 displays the most divisive questions for each cleavage. Panel A displays the five questions with the highest F_{ST} averaged over all 16 available wave groupings. The most divisive questions have F_{ST} levels on the order of 10–12% (one exception is *fund*, already discussed above, with an average religion F_{ST} of 42.9%). The pattern of most divisive questions is also interesting, and is largely in line with what one would expect a priori. For instance, along the party ID cleavage, the most divisive questions have to do with government spending policies—toward blacks, health care and welfare—as well as confidence in the federal government. These clearly reflect differences in political platforms across parties, so it makes sense that they would be the most divisive. Along the rural–urban divide, three of the top-five most divisive questions concern gun ownership. Finally, the single most divisive issue between genders is the question about whether the respondent is afraid to walk at night in their neighbourhood ($F_{ST} = 11.3\%$).

Panel B of Online Appendix Table A6 shows that the single most divisive question for each cleavage tends to remain relatively stable over time. For instance, for religion the question about the respondent's degree of fundamentalism is the most divisive question in every single year. Another example concerns the racial cleavage: questions about government spending to improve the lives of black people (*natrace*) and attitudes toward police violence (*polhitok*) tend to alternate as the most divisive over the years.

²⁵ See also Online Appendix Table A5 for a breakdown using the GSS-defined question categories and subcategories, rather than those based on factor analysis.

²⁶ However, excluding the question *fund* when calculating average F_{ST} for the religious cleavage does not in any way change the dynamics of the religious divide, which has been on the rise since the mid-1990s.

Panel C displays the question categories that are the most divisive: for each cleavage and period, we calculate the average F_{ST} across all questions belonging to a given question category (as obtained from factor analysis), and list the category with the highest value. Once again, we see a lot of stability in the most divisive question categories, and these make intuitive sense: questions on life satisfaction tend to be the most divisive for the income cleavage, and questions on abortion and church attendance are the most divisive for religion. However, for the political divide we see a clear shift over time in the most divisive categories: free speech and freedoms in the early period, government spending priorities in the middle of the period, and gun ownership in the 2000s.

Further analysis of the level of F_{ST} . Online Appendix Table A7 carries out a regression analysis of variation in the level of F_{ST} . We pooled all of the F_{ST} measures across cleavages, questions and periods (with 76 questions, 16 periods and six cleavages, this gives us 7,296 observations). Each panel reports results on each of three sets of regressors (all entered simultaneously): cleavage type, question category or subcategory, and time period. Panel A looks at time effects by including a dummy for each of the 16 periods. We find a U-shaped pattern: cultural fixation across all questions and cleavages tends to fall until the late 1990s, and to rise in the 2000s. Panel B replicates the ordering of F_{ST} magnitudes across cleavages: average F_{ST} is elevated for family income and religion, and is low for race and gender. Panel C analyses the level of fixation by question category. We find that across all cleavages, F_{ST} tends to be high for abortion and church attendance, and low for gun ownership and government spending.

Alternative frequency. Online Appendix Tables A8–A12 and Figures A2 and A3 replicate our baseline analysis with the set of questions that appear in the GSS at least once every five years. The frequency of observations is correspondingly coarser (nine time periods instead of 16), but the number of questions is expanded (96 rather than 76). We uncover dynamics that are unchanged compared to the baseline exercise: CF displays a U-shaped pattern over time. We also find the same patterns for the dynamics of F_{ST} : a pronounced increase in the later part of the period for party ID and religion, a more mitigated increase for race and family income, and flat or decreasing cultural divides between genders and urban categories. The analysis of the dynamics of cultural divides across question categories (Online Appendix Tables A10 and A11) reveals no new insight compared to the baseline. We conclude that expanding the set of questions by reducing the frequency of observations leads to findings that are very similar to the baseline.

The role of changing group shares. Our measures of cultural fixation could change through time either because the distribution of memes within groups changes, or because the relative sizes of the groups themselves change. Indeed, in equation (4), each group's within-group heterogeneity is weighted by the group's share (s_{k_c}) in order to calculate a measure of average within-group variation. Variation in s_{k_c} could drive changes in F_{ST} . For some cleavages, like gender, this is not an issue: group shares tend to be stable through time by construction. For others, such as religion, there is more scope for time variation due to demographic change. To isolate the role of changing within-group meme distributions from that of changing group shares, we compute alternative measures of fixation. Our algorithm, outlined formally in the Empirical Online Appendix, keeps the weight of each group constant across time at a baseline level corresponding to the initial distribution of group shares.

The time paths of these alternative F_{ST} measures are displayed in Online Appendix Figure A4. The dynamics are similar to those in Figure 2. This indicates that most of the variation in cultural fixation is not due to changes in group composition. There is one notable exception: religion. For this cleavage, changes in group shares over time do affect the dynamics of the cultural divide.

When correcting for changes in group shares, the religion F_{ST} displays L-shaped rather than U-shaped dynamics. The main change in that period is the decline of the share of respondents who are Protestant and the concurrent increase in the share of those with 'no religion'. To the extent that the cultural divide between religious and non-religious people is quantitatively significant, ignoring the increase in the share of non-religious respondents (as we do in Figure A4) amounts to flattening the dynamics of the religious divide in the second half of the sample period. We return to this observation in Subsection 3.1, when discussing sorting.

Question entry and exit. Online Appendix Figures A5 and A6 display the dynamics of CF and F_{ST} incorporating questions that are not asked uniformly throughout the sample period. The first panel of each row replicates results using the baseline set of 76 questions asked at least once every two waves. The second panel shows these series obtained from the most expansive set of questions (1,115 questions asked at least in two waves of the GSS).

Figure A5 shows that findings regarding CF are different across the first two panels: with the expanded set of questions, the average level of CF is higher, indicating that questions asked only episodically tend to display more heterogeneity. The dynamics of CF are also different: in the second panel, the series rises monotonically from the start of the sample period, going from about 0.5 to about 0.6. Overall, we confirm our finding of growing overall heterogeneity in memes after the mid-1990s when considering an expansive set of questions.

In contrast, our findings for F_{ST} in the second panel broadly confirm the baseline results. We expected the F_{ST} series to display more volatility than those based on a common set of questions, because the averages are constructed on a constantly changing set of questions, most of which are only asked episodically. However, this was not the case: comparing the first and second panels of each row of Figure A6, we see very similar levels and dynamic paths for F_{ST} across most cleavages. The only exception is for gender where we see a more pronounced rise in F_{ST} early in the period (still from a very low level), and a stabilisation rather than a fall in more recent times.

The baseline set of 76 questions and the expansive set of 1,115 questions represent polar opposite choices along a spectrum. The third panels of Figures A5 and A6 represent a compromise between the two extremes: they display the evolution of CF and F_{ST} for the baseline set of 76 questions, augmented with 21 questions that permanently exited the survey at some point, and 96 questions that were at first never asked, and then asked consistently. Here, we find that the dynamics of both cultural heterogeneity and the cultural divides are almost the same as in the baseline. One exception is for the religion cleavage, where the dynamics are a bit flatter than in the baseline. But overall, our inferences about the evolution of cultural heterogeneity do not seem greatly affected by the fact that exiting questions could be more consensual, and entering questions more divisive.

Cleavage intersections. So far we have considered one-dimensional cleavages, where groups are defined by their identity with respect to a single cleavage. It is straightforward to extend our analysis to a consideration of intersections between two or more cleavages. For instance, a group could be defined by race and gender. Groups would consist of white females, black males, etc. The availability of sufficient data to calculate meaningful shares is the only limit to how many intersections can be considered. Here we consider the intersections of any two cleavages. With six cleavages, this implies 15 intersections. We calculate F_{ST} for all these intersections and present the results in Figure A7.

We find that when the underlying cleavages display different dynamics, the resulting cleavage intersection tends to feature dampened trends. For instance, gender displays a flat cultural divide, religion displays an increasing divide, and the dynamics for the gender–religion intersection are

weakly increasing. We also find that any intersection involving party ID tends to display rising divides after 2000—this is the case even for intersections defined by urbanicity and party ID, where the intersection involves a cleavage with a declining divide. This exercise highlights that the rising cultural divide along party ID is a salient and robust finding. Another expected result is that F_{ST} levels based on intersections tend to be larger, on the order of 2–5%. That is because groups are now more finely defined, for example, the income–religion intersection involves 25 groups (in the limit, if every individual is a group, F_{ST} is equal to 1).

2.5. Alternative Approach Based on Regressions

As an alternative approach to assessing the dynamics of the cultural divide, we use a regression approach. For each meme m at time t , we run a regression of the following form:

$$y = \alpha + \sum_{c=1}^{11} \sum_{k_c=1}^{K_c-1} \beta_{k_c} D_{k_c} + \varepsilon, \quad (11)$$

where y is the response of an individual to meme m at time t , α is a constant, D_{k_c} is a dummy variable taking on a value of 1 if the individual is in cleavage category k_c , and ε is an error term. With 16 time periods and 76 memes, this means we are running 1,216 regressions. We record the total R^2 from each of these regressions, a measure of the informativeness of all identity cleavages together, in terms of cultural memes. We then calculate the incremental R^2 due to each set of cleavage dummies. To do so, we rerun the above regressions excluding the set of dummies for the cleavage of interest (this is an additional 1,216 regressions for each of six cleavages). We refer to the R^2 from these regressions as the restricted R^2 . For each meme i at time t , we then take the difference between the total R^2 and the restricted R^2 , giving us the incremental R^2 for the corresponding cleavage. For each time t , we then average the total and the incremental R^2 over all memes.

This approach is related to measuring fixation using F_{ST} . The greater the explanatory power of an identity cleavage for cultural values, the higher the corresponding incremental R^2 in the above regression. Similarly, the F_{ST} for that cleavage will tend to be relatively high. One advantage of the R^2 approach is that all identity cleavages are entered jointly, so we are controlling for the effect of other cleavages when assessing the explanatory power of a particular cleavage.

The results are presented in Table 5 and displayed graphically in Figures 3 and A8.²⁷ The overall R^2 , i.e., the joint explanatory power of all cleavages, displays a U-shaped pattern and is minimised for the 1993–1994 wave grouping (Figure A8). The level of the R^2 itself is modest, going from 9.3% in 1972–1973 to 7.3% in 1993–1994 and back to 9.5% in 2016–2018. The ability of cleavages overall to explain answers to these 76 GSS questions has increased since the late 1990s, indicating growing cultural divides in the last two decades.

However, this average pattern masks interesting differences by cleavage. These differences largely replicate those found for F_{ST} , confirming that the average incremental R^2 reflects a similar concept (the informativeness of cleavages for memes) as does F_{ST} . This is most easily seen by comparing Figures 2 and 3: the dynamics of incremental R^2 are broadly similar to those of F_{ST} , cleavage by cleavage. Overall, incremental R^2 values for party ID are relatively flat until

²⁷ Online Appendix Table A12 and Figures A9–A10 replicate these results using five-year frequency data, expanded to 96 questions. The results are very similar to those described here for the baseline exercise using a two-wave frequency and 76 questions.

Table 5. Overall and Incremental R^2 , Over Time.

Time period	R^2 overall	Family income	Gender	Party ID	Race	Religion	Urbanicity
1972	9.35%	1.85%	0.51%	1.16%	0.93%	2.01%	0.95%
1974	8.42%	1.84%	0.51%	1.07%	0.78%	1.82%	0.94%
1976	7.65%	1.68%	0.52%	0.94%	0.83%	1.58%	0.89%
1979	8.49%	1.71%	0.58%	1.00%	0.83%	2.00%	0.93%
1982	8.48%	1.78%	0.50%	0.96%	0.97%	1.58%	0.82%
1984	8.01%	1.82%	0.57%	1.17%	0.65%	1.69%	0.74%
1986	8.12%	1.75%	0.56%	0.93%	0.86%	1.50%	0.69%
1988	7.67%	1.97%	0.62%	1.01%	0.62%	1.57%	0.66%
1990	7.48%	1.64%	0.68%	0.95%	0.61%	1.43%	0.71%
1993	7.30%	1.75%	0.54%	1.02%	0.70%	1.54%	0.60%
1997	7.36%	1.49%	0.50%	1.10%	0.63%	1.60%	0.58%
2001	7.92%	1.76%	0.54%	1.04%	0.73%	1.71%	0.60%
2005	9.27%	1.93%	0.48%	1.86%	0.96%	1.72%	0.61%
2009	9.02%	1.80%	0.50%	1.55%	0.82%	1.98%	0.58%
2013	9.43%	2.06%	0.41%	1.81%	0.75%	2.02%	0.54%
2016	9.45%	1.83%	0.44%	1.96%	0.83%	2.16%	0.46%
Average	8.34%	1.79%	0.53%	1.22%	0.78%	1.74%	0.71%

Notes: Time period refers to two-wave groupings. So, for instance, 1972 refers to pooled data over the 1972 and 1973 waves of the GSS.

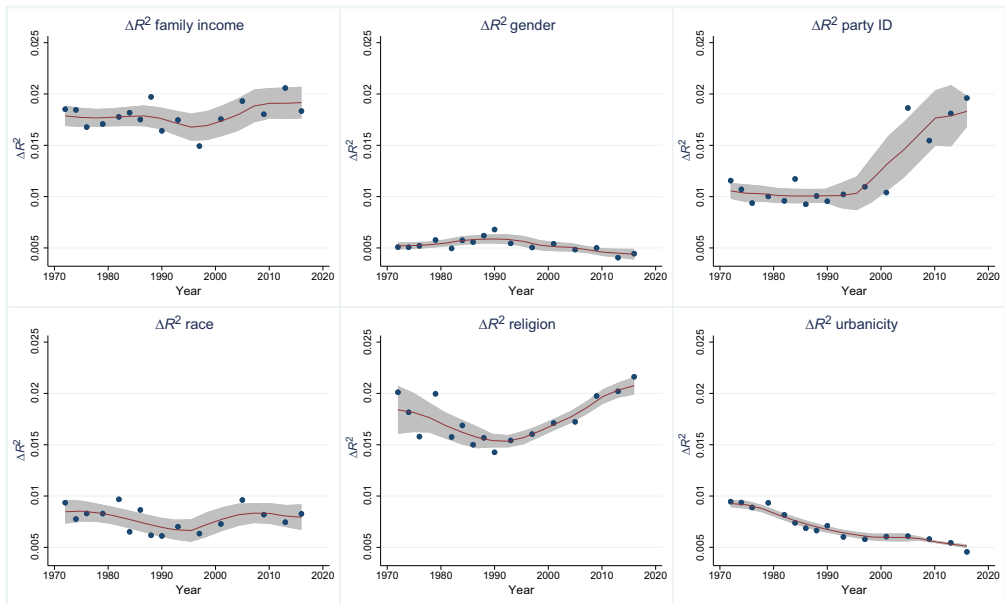


Fig. 3. Incremental R^2 Across Each of Six Cleavages, Over Time (76 Questions, Two-wave Groupings).

the early 2000s, after which they increase rapidly, doubling in the span of 15 years. We uncover a similar pattern for religion, with an acceleration starting slightly earlier, in the second half of the 1990s. We find a flat incremental R^2 for gender, and a falling incremental R^2 for urbanicity, as we did when capturing cultural divides with F_{ST} .

3. Interpretation of the Empirical Findings

How can we understand the patterns of change in heterogeneity and the cultural divides described in the previous section? Here, we use the model as a lens through which to interpret these empirical patterns. To do so, we rely on the specific theoretical configurations captured by Propositions 1 through 3. We begin by relating changes in the main model parameters to the dynamics of cultural divisions across different cleavages. Next, we analyse the role of conformism and cultural innovations in explaining differences across specific memes.

3.1. *Changes in Interaction Technologies and Sorting*

In the context of our conceptual framework, the main parameters of interest are the average degree of conformism (\bar{p}) and the differential susceptibility of groups to cultural innovations ($\bar{r}_k - \bar{r}_{-k}$). For a given cleavage, an increase in \bar{p} can be interpreted as either a strengthening of intragroup conformism or a weakening of intergroup conformism. According to Proposition 1, this would lead to an increase in F_{ST} . An increase in the difference between \bar{r}_k and \bar{r}_{-k} can be interpreted as a weakening of the intergroup diffusion of cultural innovations. According to Corollary 2, this would also lead to an increase in F_{ST} .

What forces might lead to shifts in the level of \bar{p} and the difference between \bar{r}_k and \bar{r}_{-k} ? These parameters reflect the relative extent of interactions within and between groups, which depends on media and communication technologies. In other words, the manner in which individuals interact with each other affects the level of \bar{p} and the difference between \bar{r}_k and \bar{r}_{-k} , i.e., whether social influence occurs mostly within groups or also between groups. In this context, the rise of new forms of digital communication in the late 1990s and early 2000s may have led to differential changes in our model's main parameters depending on the specific cleavage under consideration. We consider three examples.

Urban cleavages. In the case where cleavages imply geographic separation between groups, such as urbanicity, new interaction technologies are expected to facilitate cross-group interactions by reducing the effect of geographic barriers. In the absence of any cross-group interaction technology, geographic barriers to interactions were strongly operative. The early introduction of communication technologies like newspapers, the telephone and television already served to reduce geographic barriers to interactions, reducing \bar{p} and the difference between \bar{r}_k and \bar{r}_{-k} . More recently, new forms of digital interactions furthered the potential for cross-location interactions. Hence, for urbanicity, after the mid to late 1990s we expect that \bar{p} and the difference between \bar{r}_k and \bar{r}_{-k} both went down, reducing geography-based echo chambers. Indeed, for urbanicity, F_{ST} has tended to fall throughout our sample period.

Party ID and religion cleavages. For party ID or religion, new interaction technologies create the *potential* for people to seek out those of the same group, to interact with them disproportionately more, and to reduce the intensity of their interactions with those of a different group. In this case, the initial condition is one where there is inevitably some day-to-day interaction with people from a different group, but digital interaction technologies allow individuals to reduce or minimise such cross-group interactions. For instance, on social media, people of a specific political orientation or religion may seek each other out and interact mostly with each other, creating stronger echo chambers with respect to these cleavages. Thus, for these kinds of cleavages, we expect that \bar{p} increased and the difference between \bar{r}_k and \bar{r}_{-k} also increased since the mid to late 1990s. Consistent with this, for party ID and religion, we find strong increases in F_{ST} , particularly

after the expansion of access to the Internet in the second half of the 1990s. We would argue that similar considerations may hold for races and income groups, where new interaction technologies and tailored sources of information may have mostly facilitated within-group interactions in the past two decades.

The gender cleavage. In the case of the gender cleavage, the introduction of new forms of social interactions is likely to have had a more ambiguous effect. There is no reason to expect that Internet-based interaction technologies acted to intensify within-gender rather than between-gender interactions. In this case, the initial condition is one of intentional day-to-day interaction between men and women in offline contexts. Given the intentionality of such contact between genders (*heterophily*), we would not expect new interaction technologies to be used for one gender to systematically avoid the other gender. As a result, it is unlikely that new interaction technologies would have created a pronounced increase in the cultural divide across genders. Indeed, our empirical analysis shows no tendency for much of a change in the gender F_{ST} across time.

Sorting and the choice of identity traits. For identity cleavages with scope for an echo chamber effect (party ID, religion, family income and race), Proposition 3 suggests that we should expect the effect to be particularly important for cleavages along which individuals can freely choose their trait. For example, individuals can choose their party ID or their religious affiliation. The payoff from changing party ID or religious affiliation in order to align individual values with those of the majority is greater if within-group conformism is stronger (i.e., when \bar{p} is higher). By introducing a complementarity between within-group echo chambers and sorting, this increases the cultural divide along party ID and religion.

Indeed, the religion divide is a good example of Proposition 3 at work: our evidence suggests that much of the increase in F_{ST} may be due to increased sorting, with a declining share of Protestants and a growing share of non-religious individuals: as noted earlier, when comparing Figures 2 and A4, we no longer see a growing religious divide when keeping group shares constant. This indicates that sorting of individuals with secular cultural values into the category ‘non-religious’ explains in large part the growing religious divide that is observed in the USA since the late 1990s.

In contrast to party ID and religion, individuals can typically not choose their race. Although the Internet and social media make it easier for isolated individuals of a certain race to interact with others of the same race, it does not increase the sorting of races on particular values. Hence, fixation on race only varies because of changes in the distribution of values within racial categories, or because of differential demographic trends across racial groups, not because of sorting.

Summary: the dynamics of F_{ST} across identity cleavages. To summarise, the above discussion suggests that two dimensions matter for the dynamics of F_{ST} : the scope of the echo chamber effect, and the ease of sorting into identity trait. Figure 4 shows this graphically in a two-dimensional matrix with three regions. We can distinguish between three categories of identity cleavages:

A first category consists of identity cleavages for which there is little scope for new interaction technologies to lead to an echo chamber effect: urbanicity and gender. For the identity traits in the left half of Figure 4 we would therefore expect no increase in F_{ST} —it could fall or stay flat.

A second category consists of identity cleavages with scope for a more pronounced echo chamber effect, but identity traits cannot be freely chosen: race, and to a lesser extent, family

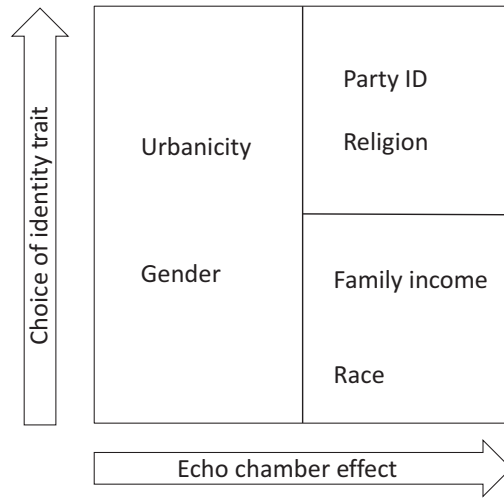


Fig. 4. *Classification of Identity Cleavages.*

income. For the identity cleavages in the bottom-right quadrant of Figure 4 we would therefore expect the introduction of modern media to have a moderately positive effect on F_{ST} .

A third category consists of identity cleavages with echo chamber effects for which the complementarity between echo chambers and sorting is at work: party ID and religion. For the identity cleavages in the top-right quadrant of Figure 4 our conceptual framework therefore predicts a steeper increase in F_{ST} following the introduction of modern media.

These theoretical predictions are largely consistent with the empirical patterns seen after the introduction of modern media and communication technologies. Since the late 1990s, fixation is mostly flat or decreasing for gender and urbanicity; it is mildly increasing for race and family income; and it is strongly increasing for religion and party ID. While these patterns are suggestive of a role for new interaction technologies, a formal empirical test of the forces at work in the theory is left for future research.

3.2. *Cultural Innovations and Conformism*

We documented substantial heterogeneity in cultural evolution across question categories and within question categories. In what follows, we discuss three examples.

Crime. For questions pertaining to crime, CF and F_{ST} both exhibit U-shaped patterns over time. What might account for these patterns? One obvious candidate is the evolution of the violent crime rate and the property crime rate, both of which peaked in 1991. There are many candidate explanations for the decline in crime rates since then. They include more and better policing, mass incarceration, the end of the crack epidemic, the introduction of legalised abortion, and the decline in lead exposure, among others.

To see how the rapid decline in crime rates might have changed people's attitudes towards crime issues, it is useful to focus on a particular example. Take, for instance, the question in the GSS that asks respondents whether courts deal too harshly or not harshly enough with criminals. In 1991, of those surveyed by the GSS, 4% answered courts were dealing too harshly with

criminals, compared to 80% who said courts were not harsh enough. By 2018, those numbers had changed to 18% and 54%, respectively. In our model, we would view this as a cultural innovation that increases the minority view that courts are too harsh. In other words, there is an innovation to the minority value. Through the parameter \bar{r}_k , this leads to a changing cultural consensus in the direction of a growing minority believing that courts are too harsh. In that case, cultural heterogeneity increases, since the overall consensus that courts are not harsh enough is waning. Hence, according to Proposition 2, we should expect CF to increase, because of a cultural innovation to the minority value.

At the same time, views on crime have become more divided across identity groups. Going back to the question on the harshness of courts, consider the changing racial divide. In 1991, there was a broad consensus across racial groups: only 3% of whites and 12% of blacks answered that courts were treating criminals too harshly. By 2018, these shares had increased to 15% and 35%, respectively. One way of interpreting these facts is that whites have a lower \bar{r}_k for this particular value than blacks. The sensitivity of each group to the cultural innovation differs, because different groups may be affected differently by, say, the increase in mass incarceration. Consistent with Proposition 2, if \bar{r}_k differs from \bar{r}_{-k} , the divide between groups increases when an innovation occurs, and F_{ST} increases.

Another interesting question in the crime category relates to the legality of marijuana use. Between 1972 and the early 1990s there was a growing consensus that it should be illegal, reaching a maximum of 83% in favour of keeping it illegal in 1990. Since then, the consensus has completely shifted, and by 2018 only 35% were still in favour of keeping marijuana illegal. As with the question on the harshness of courts, this has led to an increase in CF since the early 1990s. In terms of F_{ST} , here as well the susceptibility to the innovation differs across groups. For example, blacks were less in favour of legalisation than whites in 1990; this had switched by 2018.

These examples illustrate that when circumstances change, in a way that affects different groups differently, the pre-existing consensus may weaken (showing up as increasing cultural heterogeneity) and there may be growing divides across identity groups (showing up as increasing F_{ST}). Looking ahead, whether in the long run the pre-existing consensus is replaced by a new consensus or whether the new steady state is a lack of consensus depends on the specific question. For example, in the case of marijuana the growing majority in favour of legalisation is such that in recent years CF has started to decline, suggesting that a new consensus might be emerging. Indeed, when the old consensus is replaced by a new consensus, Proposition 2 predicts a hump-shaped path for CF .

Free speech. For questions related to freedom of speech, CF and F_{ST} exhibit decreasing patterns over time. As an example, consider the question whether an atheist should be allowed to make a speech against religion in your community. In 1972, 67% of those surveyed answered positively; by 2018, this percentage had increased to 80%. This points to a long-term growing consensus in favour of free speech, thus leading to a falling CF over time. In general, this increasing agreement happened across all groups. As an illustration, consider how the question on free speech for an atheist changed across the rural–urban divide. In 1972, 80% of those living in locations of more than 1 million favoured free speech for atheists, compared to 58% of those living in locations of fewer than 10,000. In 2018, those numbers were 75% and 79%, respectively. Hence, for this particular question on free speech, the rural–urban divide all but disappeared. As a result, in this case F_{ST} converged to a number very close to zero.

In the context of our model, this can be viewed as the diffusion of a cultural value across groups. The end of McCarthyism, the civil rights movement, and the increasing level of education might have led to a renewed commitment to the First Amendment. Not all groups took this change on board simultaneously, but eventually it diffused to all groups. This led to a decrease in the difference between \bar{r}_k and \bar{r}_{-k} . According to Corollary 2, this should lead to a decrease in F_{ST} . This is an example of cultural convergence. Why do some changing values diffuse across groups and others do not? One reason is that the issue at stake may affect different groups very differently. For example, the harshness of courts may affect African Americans differently from whites, whereas the issue of free speech does not have a strong racial element.

Same-sex relations. Within question categories, specific questions exhibit strong dynamics that are worth highlighting in the context of our model. For example, the percentage of people answering that homosexual relations were always wrong peaked at the end of the eighties, with 78% in 1987; by 2018, that figure had come down to 33%. The decline was especially rapid in the early 1990s. Between two consecutive GSS waves, 1991 and 1993, the percentage dropped by nearly 10 percentage points.²⁸ This increasing tolerance towards same-sex relations translated into an increasing CF . This is consistent with Proposition 2: as the original consensus disintegrates, we initially see rising disagreement in society, and hence an increase in CF . This has happened across groups, but not at the same rate. Compare Strong Republicans to Strong Democrats, two groups along the party ID cleavage. In 1990, the share answering homosexual relations were not wrong at all was 7% and 18%, respectively. These figures stood at 34% and 60% respectively in 2018. Hence, both saw an increase, but the increase was more pronounced for Democrats. In the context of our model, this is a cultural change going from one consensus to a different consensus, but at differing rates across groups. Thus, F_{ST} increases in the transition.

4. Conclusion

In this article, we conducted a systematic analysis of the evolution of cultural heterogeneity in the USA. We sought to assess growing concerns about deepening cultural divides between groups defined along a wide range of identity cleavages. We considered six such cleavages: gender, income, political orientation, race, religion and urbanicity. Using answers to questions on values, attitudes and norms, we characterised the time paths of cultural divides. The picture that emerges from this analysis is not one of a generalised deepening of cultural divisions. First, overall cultural heterogeneity declined from 1972 to the mid-1990s and has increased since, but remains roughly at the level it reached at the beginning of the period. Second, the extent of the cultural divides is very small, as between-group variation represents between 0.6% (for gender) and 2.1% (religion) of total variation: most variation in memes is within groups. Third, the time path of the cultural divides differs across cleavages: in recent decades, the divide has grown steeply across political and religious lines, more weakly across income quintiles and races, but has declined for urbanicity and remained flat for gender. Many commentators have focused on the cultural divide across political lines, ignoring trends across other divides and ignoring heterogeneity across memes. Our article in contrast took a more systematic approach of looking at a wide range of cleavages and memes. This broader approach does not warrant a pessimistic conclusion that the USA is experiencing cultural disintegration. The data suggests a more qualified conclusion that

²⁸ Fernández *et al.* (2020) explore the origins of changing attitudes toward gay people in the USA. In our terminology, they study the origins of a cultural innovation. They argue that the AIDS epidemic explains the emergence of a new cultural consensus of tolerance toward homosexual rights and lifestyles.

cultural divisions have grown only since the late 1990s, only for some cleavages and only for some memes.

We also provided a theoretical interpretation for the heterogeneity in the dynamics of cultural divides across cleavages and memes. In our model, agents are born with cultural traits inherited with variation from their parents. Social influence then triggers potential changes in these inherited traits, because agents conform to the majority of their own group and also because they respond to cultural fads and innovations initially espoused only by a minority: social influence is a major force explaining cultural change. The degree to which cultural change is group-specific determines the evolution of cultural divides between groups.

The model suggests that the manner in which agents obtain information and interact with each other has important effects on the evolution of cultural divisions. If the predominant mode of interaction is between groups, cultural change will occur in a similar manner across group identities, keeping F_{ST} low. If instead most interactions are within groups and information is group-specific, it becomes more likely that cultural fixation increases as a result of a cultural innovation. We argued that new interaction technologies have contrasted effects on within- versus between-group social interactions depending on the nature of the cleavage. When the preexisting baseline is one in which interactions occur mostly within-groups (such as for urbanicity), the introduction of new communication technologies can lead to an increase in cross-group interactions, and hence declining divides. Instead, when new technologies such as tailored cable TV channels and online social media allow individuals to seek out others with the same identity (for instance people with the same political views), they have the potential to create echo chambers.

The dynamics of cultural divides also depend on characteristics of the cultural traits under consideration. For instance, since the mid-1990s, there is an increasing view that the justice system is too harsh on crime, but this change has occurred differentially across races. African Americans are more likely to find the judicial system too harsh than whites. In terms of our model, this happens because the susceptibility of each group to this specific cultural innovation is different, creating a growing divide.

Our work can be extended in several directions. First, for each cleavage, we have considered all groups jointly, but this may mask interesting patterns for specific group pairs. For instance, the average divide between all races may follow a certain time path, but the specific divide between Hispanics and whites may follow a different pattern. Second, given that individuals' identities are multidimensional, we could go beyond considering single cleavages, or simple interactions of two cleavages, and instead define groups in more complex ways. Third, our methodology can easily be applied to other societies, provided survey data on attitudes, norms and values are available and gathered consistently across time. Fourth, this article has characterised the evolution of cultural divides, and provided a theoretical framework to make sense of these patterns conceptually. Future work should aim to test whether these hypothesised mechanisms are operative in the data.

Ultimately, we are interested in the evolution of cultural heterogeneity because of its potential effects on social cohesion, social capital and the ability of different groups to reach agreements on public policy. In this article, we have characterised the evolution of cultural divides, but the question of their impact on political economy outcomes such as public goods provision, voting, inequality and economic interactions between groups remains an important topic for future research.

*Southern Methodist University, USA, NBER and CEPR, Europe
University of California, Los Angeles, USA and NBER, USA*

Additional Supporting Information may be found in the online version of this article:

Online Appendix Replication Package

References

- Alesina, A. and Giuliano, P. (2015). 'Culture and institutions', *Journal of Economic Literature*, vol. 53(4), pp. 898–944.
- Alesina, A., Tabellini, G. and Trebbi, F. (2017). 'Is Europe an optimal political area?', *Brookings Papers on Economic Activity*, vol. 48(1), pp. 169–234.
- Algan, Y., Guriev, S., Papaioannou, E. and Passari, E. (2017). 'The European trust crisis and the rise of populism', *Brookings Papers on Economic Activity*, vol. 48(2), pp. 309–82.
- Allport, G.W. (1954). *The Nature of Prejudice*, Cambridge, MA: Addison-Wesley Publishing Company.
- Ashraf, Q. and Galor, O. (2013). 'The out of Africa hypothesis, human genetic diversity, and comparative economic development', *American Economic Review*, vol. 103(1), pp. 1–46.
- Autor, D., Dorn, D., Hanson, G. and Majlesi, K. (2020). 'Importing political polarization? The electoral consequences of rising trade exposure', *American Economic Review*, vol. 110(10), pp. 3139–83.
- Bell, A., Richerson, P.J. and McElreath, R. (2009). 'Culture rather than genes provides greater scope for the evolution of large-scale human prosociality', *PNAS*, vol. 106(2), pp. 17671–4.
- Bernheim, B.D. (1994). 'A theory of conformity', *Journal of Political Economy*, vol. 102(5), pp. 841–76.
- Bertrand, M. and Kamenica, E. (2020). 'Coming apart? Cultural distances in the United States over time', University of Chicago.
- Bikhchandani, S., Hirshleifer, D. and Welch, I. (1992). 'A theory of fads, fashion, custom, and cultural change in informational cascades', *Journal of Political Economy*, vol. 100(5), pp. 992–1026.
- Bisin, A. and Verdier, T. (2000). 'Beyond the melting pot: cultural transmission, marriage, and the evolution of ethnic and religious traits', *Quarterly Journal of Economics*, vol. 105, pp. 955–88.
- Boxell, L., Gentzkow, M. and Shapiro, J.M. (2017). 'Greater Internet use is not associated with faster growth in political polarization among US demographic groups', *PNAS*, vol. 114(40), pp. 10612–17.
- Boyd, R. and Richerson, P.J. (1988). *Culture and the Evolutionary Process*, Chicago: University of Chicago Press.
- Boyd, R. and Richerson, P.J. (2005). *The Origin and Evolution of Cultures*, New York: Oxford University Press.
- Cavalli-Sforza, L.L. and Feldman, M.W. (1981). *Cultural Transmission and Evolution: A Quantitative Approach*, Princeton: Princeton University Press.
- Cavalli-Sforza, L.L., Menozzi, P. and Piazza, A. (1994). *The History and Geography of Human Genes*, Princeton, NJ: Princeton University Press.
- Cover, T.M. and Thomas, J.A. (2006). *Elements of Information Theory*, Hoboken, NJ: John Wiley & Sons.
- Dawkins, R. (1989). *The Selfish Gene*, (2 ed.), Oxford: Oxford University Press.
- Desmet, K., Ortuño-Ortín, I. and Wacziarg, R. (2017). 'Culture, ethnicity and diversity', *American Economic Review*, vol. 107(9), pp. 2479–513.
- Desmet, K. and Wacziarg, R. (2018). 'The cultural divide', NBER Working Paper No. 24630, National Bureau of Economic Research.
- DiMaggio, P., Evans, J. and Bryson, B. (1996). 'Have Americans' social attitudes become more polarized', *American Journal of Sociology*, vol. 102(3), pp. 690–755.
- Doepke, M. and Zilibotti, F. (2008). 'Occupational choice and the spirit of capitalism', *Quarterly Journal of Economics*, vol. 123(2), pp. 747–93.
- Enke, B. (2020). 'Moral values and voting', *Journal of Political Economy*, vol. 128(10), pp. 3679–729.
- Fernández, R. (2014). 'Cultural change as learning: the evolution of female labor force participation over a century', *American Economic Review*, vol. 103(1), pp. 472–500.
- Fernández, R. and Fogli, A. (2006). 'Fertility: the role of culture and family experience', *Journal of the European Economic Association*, vol. 4(2–3), pp. 552–61.
- Fernández, R., Parsa, S. and Viarengo, M. (2020). 'Coming out in America: AIDS, politics, and cultural change', Working Paper, New York University.
- Fiorina, M.P. and Abrams, S.J. (2008). 'Political polarization in the American public', *Annual Review of Political Science*, vol. 11, pp. 563–88.
- Gennaioli, N. and Tabellini, G. (2019). 'Identity, beliefs, and political conflict', Working Paper, July.
- Gentzkow, M., Shapiro, J. and Taddy, M. (2016). 'Measuring polarization in high-dimensional data: method and application to congressional speech', NBER Working Paper No. 22423, National Bureau of Economic Research.
- Gethin, A., Martínez-Toledano, C. and Piketty, T. (2019). 'Political cleavages and inequality: evidence from electoral democracies, 1950–2018', work in progress.

- Giavazzi, F., Petkov, I. and Schiantarelli, F. (2019). 'Culture: persistence and evolution', *Journal of Economic Growth*, vol. 24, pp. 117–54.
- Giuliano, P. (2007). 'Living arrangements in Western Europe: does cultural origin matter?', *Journal of the European Economic Association*, vol. 5(5), pp. 927–52.
- Giuliano, P. and Nunn, N. (2020). 'Understanding cultural persistence and change', *Review of Economic Studies*, forthcoming.
- Guiso, L., Herrera, H. and Morelli, M. (2016). 'Cultural differences and institutional integration', *Journal of International Economics*, vol. 99, pp. S97–113.
- Guriey, S. and Papaioannou, E. (2020). 'The political economy of populism', *Journal of Economic Literature*, forthcoming.
- Henrich, J. (2015). *The Secret of Our Success: How Culture Is Driving Human Evolution, Domesticating Our Species, and Making Us Smarter*, Princeton, NJ: Princeton University Press.
- Inglehart, R. (1997). *Modernization and Postmodernization: Cultural, Economic, and Political Change in 43 Societies*, Princeton, NJ: Princeton University Press.
- Inglehart, R. and Baker, W.E. (2000). 'Modernization, cultural change and the persistence of traditional values', *American Sociological Review*, vol. 65, pp. 19–51.
- Jost, L. (2008). 'GST and its relatives do not measure differentiation', *Molecular Ecology*, vol. 17(18), pp. 4015–26.
- Kuran, T. and Sandholm, W.H. (2008). 'Cultural integration and its discontents', *Review of Economic Studies*, vol. 75(1), pp. 201–28.
- Lazear, E. (1999). 'Culture and language', *Journal of Political Economy*, vol. 107(S6), pp. S95–126.
- Luttmer, E.F.P. and Singhal, M. (2011). 'Culture, context, and the taste for redistribution', *American Economic Journal: Economic Policy*, vol. 3(1), pp. 157–79.
- McCarty, N., Poole, K.T. and Rosenthal, H. (2006). *Polarized America: The Dance of Ideology and Unequal Riches*, Cambridge, MA: MIT Press.
- Murray, C. (2013). *Coming Apart: The State of White America, 1960–2010*, New York, NY: Crown Forum.
- Norris, P. and Inglehart, R. (2009). *Cosmopolitan Communications: Cultural Diversity in a Globalized World*, Cambridge: Cambridge University Press.
- Norris, P. and Inglehart, R. (2019). *Cultural Backlash: Trump, Brexit, and Authoritarian Populism*, Cambridge: Cambridge University Press.
- Olivier, J., Thoenig, M. and Verdier, T. (2008). 'Globalization and the dynamics of cultural identity', *Journal of International Economics*, vol. 76(2), pp. 356–70.
- Putnam, R.D. (2000). *Bowling Alone: The Collapse and Revival of American Community*, New York: Simon & Schuster.
- Richerson, P.J. and Boyd, R. (2004). *Not by Genes Alone: How Culture Transformed Human Evolution*, Chicago: University of Chicago Press.
- Ritzer, G. (2011). *Globalization - The Essentials*, London: Wiley-Blackwell.
- Tesei, A. (2017). 'Trust and racial income inequality: evidence from the U.S.', Working Papers 737, Queen Mary University of London.
- Wright, S. (1949). 'The genetical structure of populations', *Annals of Human Genetics*, vol. 15(1), pp. 323–54.