

NBER WORKING PAPER SERIES

UNDERSTANDING CULTURAL CHANGE

Raquel Fernández

Working Paper 34077

<http://www.nber.org/papers/w34077>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

July 2025

I wish to thank Ben Enke, Martin Fiszbein, and participants at the Economics of Culture Kabiller Conference for many helpful comments. This paper will be a contribution to the Handbook of Culture and Economic Behavior. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Raquel Fernández. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Understanding Cultural Change
Raquel Fernández
NBER Working Paper No. 34077
July 2025
JEL No. P0, Z1

ABSTRACT

Culture's influence on economic outcomes is no longer controversial among economists even if it remains largely ignored in many areas of economics. This paper tackles a different question: why does culture change? The underlying premise adopted here is that culture changes because incentives change, transforming actions and beliefs. An idiosyncratic review of the literature follows that illustrates how the environment (e.g., the prevalence of pathogens or the suitability of land for pastoralism) and historical experiences (e.g., colonization, war, or migration) can affect relationships of power in society and shape people's beliefs. It then examines the role of new information and ideas (i.e., learning) and finally the role of policies in shaping incentives and changing culture. A second part of the paper reviews work that models the mechanisms of cultural change more explicitly, using quantitative models to examine the interplay between economic incentives and evolving beliefs or preferences and to study the importance of intermediating mechanisms. Given that one of the most profound cultural and economic transformations of the past 150 years concerns gender roles, this theme recurs throughout.

Raquel Fernández
New York University
Department of Economics
and NBER
raquel.fernandez@nyu.edu

Understanding Cultural Change*

Raquel Fernández

NYU, NBER, CEPR, IZA

July 2025

1 Introduction

Culture's influence on economic outcomes is no longer controversial among economists even if it remains largely ignored in many areas of economics.¹ A robust and growing literature has established culture's imprint on outcomes ranging from trust, entrepreneurship, corruption, violence, and education, to living arrangements, marriage markets, gender roles, and xenophobia.² This chapter tackles a different question about which relatively little has been written: *why does culture change?* That culture changes is indisputable. After all, contemporary societies with very different beliefs about fundamental issues such as gender roles, sexuality, or the moral value of hard work, presumably because they adapted to different environments, had different histories, and responded to a variety of different shocks. This has led some economists to think that culture only changes slowly, perhaps replacing one set of rules of thumb with others that do better in a changed environment. But, as the last few decades have demonstrated especially in the domains of gender roles and same-sex relationships, culture may change swiftly and profoundly.

What then drives cultural change? In virtually all theories, shocks – ecological, technological, demographic, informational, or social (e.g., war) – can cause culture to change. For example, in evolutionary models such as those of Boyd and Richerson (1988) or Giuliano and Nunn

*I wish to thank Ben Enke, Martin Fiszbein, and participants at the Economics of Culture Kabiller Conference for many helpful comments.

¹I deliberately refrain from defining culture given the wide range of definitions that already exist. This should not be problematic for understanding the literature discussed in this chapter but loosely I take culture to mean shared beliefs and values that shape behavior.

²For an early review of the literature see Fernández (2011). Many of the chapters in this handbook – particularly those in Parts IV and VI – provide reviews of the literature that study culture's influence on a variety of outcomes related to growth, gender roles, firm behavior, health, politics, and development.

(2020), or those that feature a more intentional intergenerational transmission process as in Bisin and Verdier (2000), a shock can cause new rules of thumb to be adopted, increase experimentation, or reduce parents' incentives to instill their beliefs in their children. Alternatively, in models with multiple equilibria as in Cole, Mailath and Postlewaite (1992) or Mailath and Postlewaite (2003), a shock can make an equilibrium infeasible. In those models, social or cultural norms determine, for example, men's rank in society which matters in a non-market sector such as matching in the marriage market. Despite not caring about status per se, women may be willing to marry less wealthy but higher-status men if the latter is inherited by their children, giving them a similar advantage. Shocks can destroy these norms by rendering some matching choices no longer incentive compatible and unraveling the equilibrium.

Below I discuss various papers that contribute to the literature on cultural change. Although the majority of them do not explicitly frame themselves as studying cultural change, they can nonetheless be fruitfully read in that light. The selection is necessarily idiosyncratic and intentionally broad, especially in the first section. In particular, this chapter is not a review of the vast literature that could be interpreted as potentially relevant to cultural change nor does it use surveys to study how responses have changed over time.³ Instead, I select papers across a range of topics based on my own interests, tastes, and familiarity with the literature. In particular, I focus on cultural changes that have a fairly direct economic significance. Given that one of the most profound cultural and economic transformations of the past 150 years concerns gender roles, this theme recurs throughout the chapter.

The chapter is organized in two parts. The first and longest section presents a wide-ranging set of studies that, while not always framed as analyses of cultural change, offer valuable evidence of how culture and behavior respond to shocks. These shocks are somewhat arbitrarily divided into four themes: the environment, historical experiences, learning, and policies. The second section turns to work that models the mechanisms of cultural change more explicitly, using a quantitative structure to examine the interplay between economic incentives and evolving beliefs or preferences.

³See Desmet and Wacziarg (2018) for a very interesting account of convergence/divergence of opinions across and within groups in the US using the General Social Survey (GSS).

A few themes and areas for future research emerge from this chapter. First, while in all theories cultural change is driven by changes in incentives, the *speed* of this change is an endogenous outcome that should be studied more extensively. When will culture change slowly versus rapidly? What are the factors that hinder change or accelerate it? The answers may depend on how strongly different social groups identify with particular norms, on the structure of social networks, and on the extent to which interests diverge across groups.⁴

Second, and relatedly, culture is of course a social phenomenon. As such, it is maintained via a combination of sticks and carrots. Prevailing norms are unlikely to benefit all equally and often work to actively disfavor some groups. Thus change in particular directions may be discouraged or resisted by those who gain from the status quo. Furthermore, the stability and persistence of culture is not an indication that it is or was optimal. Even from an adaptive evolutionary perspective and even in the absence of heterogeneity, an equilibrium need not be optimal even if unique (think, e.g., of payoffs that give rise to a prisoner's dilemma) and there may be multiple evolutionary stable equilibria each favoring a different set of actors. Overall, the political economy of cultural change deserves much greater attention.⁵ Case studies, theory, and particularly quantitative models that can illuminate the size and nature of intermediating mechanisms would all be valuable additions to the literature.

Third, as the papers reviewed below illustrate, not all dimensions of culture – even within the same domain, such as gender roles – respond uniformly to shocks. For example, skewed sex ratios have been shown to influence women's labor force participation, without necessarily altering beliefs about women's political participation or the acceptability of domestic violence. Understanding why some norms shift while others remain stable is central to deepening our understanding of how culture evolves.⁶

⁴See Akerlof and Kranton (2000) and the chapter by Akerlof and Kranton in this volume for a discussion of identity's role.

⁵Some recent attempts in this direction include Bisin and Verdier (2024) and Acemoglu and Robinson (2025).

⁶The catalog of oral traditions data studied by Michalopoulos and Xue (2021) could be a fruitful resource for exploring variation in the durability of different beliefs.

2 Why Does Culture Change?

The premise of this chapter is that culture changes because incentives change, transforming actions and beliefs. Incentives can change for a variety of reasons. Historical experiences such as colonization, war, or migration, can change relationships of power in society and shape people’s beliefs, giving rise to different behavior and norms. New information, ideas, or exposure to role models – which I group under “learning” – can lead to changes in behavior and beliefs as can the introduction or invention of a new technology. The introduction of a new policy can change incentives and affect people’s behavior as well as their beliefs as to what is appropriate, leading to cultural change. Below I give several examples from the literature that illustrate the importance of historical experience, learning, and policy in changing features of culture. The inclusion of a paper in one grouping or another is somewhat arbitrary. For example, a policy that changes behavior can change culture because the behavior contradicts the social norm or because people learn that their beliefs that justified the norm were incorrect.

2.1 The Environment

At an elementary level, the variety of beliefs across both space and time are direct evidence that societies responded to their environments in such a way as to distinguish them from one another in some set of values/beliefs. Thus, as humans expanded across the globe, their original beliefs must have diverged.⁷ Several papers have explored how diverse features of an environment have influenced particular aspects of culture. For example, Enke (2019) examines how variation in the strength of kinship ties, determined in part by the prevalence of pathogens, led to moral systems regulated by universal values (weak ties) versus communal values (strong ties) and Galor and Özak (2016) examine the link between the potential caloric yield of crops, determined by agro-climatic conditions, and the rate of time preference. Below I discuss two examples in greater depth that study how features of the environment affected cultural attitudes towards women.

⁷See, e.g., Falk et al. (2018) for evidence from a large variety of countries regarding variation in characteristics such as patience, risk taking, and trust.

Boserup (2007) put forward the hypothesis that societies that used plough agriculture as compared to hoe and digging stick, developed more unequal gender roles. The intuition is that the greater strength required to handle a plough resulted in men specializing in agriculture and women in more domestic activities. The influential paper of Alesina, Giuliano and Nunn (2013) studied this idea. The authors hypothesize that differences in agricultural technologies that gave a smaller role to women in agriculture led to the belief that women belonged in the home and not working outside of it. The authors show that female labor force participation and contemporary beliefs across individuals, ethnic groups, and countries vary systematically on average depending on whether their ancestors used the plough. To determine causality, plough use is instrumented by the variation in the suitability of the terrain for crops that are favorable to plough cultivation. They also use the epidemiological approach (Fernández (2008)) which allows them to separate the role of culture versus environment by studying second-generation immigrants to the same country. They show that individuals whose parents came from a country with a tradition of plough use have lower female labor force participation and more conservative gender attitudes.⁸ Interestingly, missing from these accounts is an explanation of why working at home should lead to the social inferiority of women. One might easily imagine that the same reasoning could be given to explain the opposite and were we today living in societies in which ancestral plough use were associated with social and political inferiority of men, one would argue that when the plough was used, men had to toil the fields like animals while women reigned at home. Of course, a simple explanation for the lower status of women across all these societies (independently of the ancestral use of the plough) may lie in men's greater strength. This would have been an important characteristic when physical strength mattered – whether to directly provide subsistence or indirectly via preventing the theft of crops and livestock by others both in the same society or via war and conquest. It would also allow a man to more fruitfully engage in theft and conquest himself. This feature is more important, *ceteris paribus*, when a greater proportion of subsistence was reliant on strength.

⁸This is precisely the approach used in Fernández and Fogli (2009) and Fernández (2007) to study how cultural attitudes towards women's working outside the home were transmitted by the family in a new environment.

Becker (2025) is another interesting paper that studies the the importance of pre-industrial modes of subsistence and gender-related attitudes, in this case the relationship between pastoralism and customs aimed at impeding women from being promiscuous. In pastoral societies, management of herd animals meant that men would be away for extended periods of time, leading to the fear that they might become fathers to biologically unrelated children. This led, the paper argues, to increased concern over women's chastity and a variety of methods to monitor or hinder female promiscuity ranging from patrilocality, restrictions on women's movements outside the home, harsh punishments for adultery, the association of family honor with the virginity of an unmarried daughter, and infibulation. Instrumenting historical reliance on pastoralism with the average suitability of an ethnic group's land for pastoralism relative to agriculture, the author shows that pastoral ethnic groups are more likely to restrict women's freedom of movement and for these women to have fewer sexual partners. Intriguingly, men in historically pastoral societies also have fewer sexual partners (not necessarily a mechanical finding since alternative sexual partners need not be exclusive). It would be interesting to know whether the latter is driven primarily by greater monitoring of women, thereby reducing alternatives, or whether there is some strong norm for men not to be promiscuous (which I somehow find doubtful).

2.2 Historical Experiences

Significant historical experiences leave their imprint on societies, potentially enduring for centuries. In this case, even if economic incentives have reverted to their previous state or if they no longer require particular attitudes and values, culture may be permanently changed as a result of intergenerational transmission.

As in the prior section, in most instances there are unlikely to be outcomes and attitudes measured prior to a shock that occurred a long time ago. This requires more attention to be paid to ruling out pre-existing sources of heterogeneity that may be correlated with the cultural outcome. Below I illustrate the importance of historical shocks such as slavery, wars, and migration on a variety of outcomes.

Slavery and War

The transatlantic slave trade lasted 400 years and left its mark on many dimensions of contemporary Africa. Nunn and Wantchekon (2011) note that the slaves were supplied not only by organized raids and warfare, but also by individuals turning against others in their own community in order to increase their own protection from slave raids and to obtain weapons. Slave merchants and raiders also formed strategic alliances with key groups in a community in order to extract slaves. The authors hypothesize that this aspect of the transatlantic slave trade may have had an important effect on culture in Africa, reducing trust in neighbors, relatives, and the state. Trust has been shown to affect a variety of economic outcomes (see Guiso, Sapienza and Zingales (2006) for a review), which makes it important to understand why some places/peoples exhibit lower trust.

The authors study their hypothesis using responses to trust questions in the 2005 Afrobarometer. Lacking information on the share of each ethnic group that was enslaved, they principally use the (log of the) number of slaves by ethnic group normalized by the area they occupied and include both individual and district level controls as well as country fixed effects. They show that this and related measures are correlated with several trust proxies including trust in relatives, neighbors, and local government. To further establish causality, the authors turn to an instrumental variable strategy using distance from the coast as an instrument for the number of slaves. They show that ethnic groups that historically lived closer to the coast have greater distrust, in line with these locations being easier to raid. This pattern is absent, however, in other coastal areas not subject to transatlantic slavery (e.g., in 10 Asian countries included in the Asiabarometer). Furthermore, the paper shows that while some of the effect of the slave trade on trust comes from worse legal and political institutions (as measured by survey responses indicating whether the respondent approved of their local government council and whether they thought their members were corrupt), including these variables still leaves the slave trade variable statistically significant though with its magnitude reduced by 50%. Lastly, the authors use the epidemiological approach (Fernández (2008)) in order to distinguish between an individual's beliefs and the incentives provided by their environment. They study movers to other districts and show that the their measure of the historical importance of the slave trade of an individual's ethnic group

continues to affect trust, although naturally the environment also has explanatory power.

The effect of slavery can also be found in cultural attitudes towards women. Teso (2019) is part of a growing literature that has found that changes in sex ratios, occasioned by the historical experiences of war, slavery, settlement or selective migration, is associated with changed cultural attitudes towards gender roles. Teso uses the same data as Nunn and Wantchekon (2011) to study how the slave trade affected women's labor force participation (LFP). The author starts by noting that during the centuries of transatlantic slave trade, the greater value of men in plantation work led to skewed sex ratios in the areas of origin, especially for those ethnic groups hit hardest by slavery. This led to there being around two male slaves for every female one.⁹ The literature cited by Teso argues that relative absence of men led women to engage in activities traditionally reserved for men. In particular, women took over all agricultural activity in those places in which they had traditionally participated in agriculture and expanded into commerce in those places in which they had traditionally done less agriculture. Teso hypothesizes that this led to permanently changed gender roles, at least in the arena of work.

Using the same key variable to measure the impact of the slave trade on different ethnic groups as in Nunn and Wantchekon (2011) and an even larger set of controls, the paper shows that women from high-slavery ethnic groups are more likely to work and to do so in higher ranking occupations. Furthermore, this result is not driven by these women having more education which might have followed from the same logic. The paper also follows the epidemiological approach to show that similar results hold for women who are not living in the same area as their ethnic group, thereby separating culture from institutions. Interestingly, Teso shows that, keeping constant a woman's ethnic group, women with a husband from an ethnic group with greater historical exposure to the slave trade have greater LFP. This finding echoes that of Fernández and Fogli (2009) which shows that having a husband whose parents emigrated from a country with higher female LFP in the past increases a woman's LFP in the US, *ceteris paribus*. Importantly, Teso finds that beliefs relating to gender roles, other than those pertaining to LFP and shared household decision-making, are not affected in the same way by an ethnic group's historical intensity

⁹At the peak of the transatlantic slave trade the sex ratio in West Africa was 70 men per 100 women (Teso, p.504).

of slavery. This includes whether violence against a wife is justified in particular instances or whether women should have the same rights as men to political office. This suggests that cultural attitudes towards gender roles may have many dimensions and that not all are necessarily affected even when a fundamental outcome – women’s work outside the home – is changed.¹⁰

The transatlantic slave trade took place over centuries, making it difficult to pinpoint with the available data when and how women’s role changed. By contrast, World War I was a relatively short-lived shock that changed the sex ratio by drastically reducing the number of men. Gay (2023) uses the variation across french counties (departements) in military death rates to study mechanisms of cultural change. The military death rate is defined as the ratio of deceased soldiers born in a county relative to the size of its draft. It was very high in France overall at 16%, with a cross-county variation ranging from 6% to 29% (Gay, p.2307). The paper uses the epidemiological approach to separate culture from institutional and economic features of the local environment which could also have been affected by the military death rate. In particular, it studies the effect on female LFP of women born in one county (with a given military death rate) who now live in a different county. Importantly, the paper also examines the transmission via marriage by studying women born and residing in the same counties but with husbands born in different counties, as well as transmission via parents by studying the mother to daughter transmission as well as transmission via a mother-in-law.¹¹

Using the censuses from 1962-2012, the paper finds that a simple cross-county regression (with a variety of controls for initial differences across counties) yields that women in counties that experienced a 20% relative to a 10% military death rate were 6-10% more likely to work. Restricting the sample to women who reside in counties different from the one of birth, i.e., movers (which account for around 50% of women in 2012), this number is halved, implying that 50% of the effect is due to intergenerational transmission as opposed to changes in local labor market conditions.¹² To isolate the parental transmission effect, the

¹⁰An interesting exception is whether women should have the same rights as men, agreement with which increases with the importance of slavery for women, but not for men.

¹¹Fernández, Fogli and Olivetti (2004) had found that women’s work behavior depended on whether their mother-in-law had worked when their son was growing up.

¹²It could also be that peer effects are important in maintaining culture. Thus even if local labor markets

paper examines the employment of women who reside in their county of birth but whose parents were born in a different county. Interestingly, the paper finds a quantitatively significant effect for the mother's county of birth but a much smaller effect for the father's (similar to the findings of Farré and Vella (2013) but not those of Fernández and Fogli (2009)). It also finds a significant effect for the military death rate in the mother-in-law's county of birth (controlling for own parents' county of birth). As in Fernández, Fogli and Olivetti (2004), this finding is consistent both with husbands influencing the employment decisions of their wives or with women choosing men who are more likely to support their employment.

Lastly, the author studies attitudes related to gender roles using the Generation and Gender Survey 2005. Once again separating culture from local features by using variation in the military death rates of movers to a given county, the author shows that both men and women have more gender equal attitudes associated with higher military death rates and that there do not appear to be significant gender differences in these attitudes. Importantly, the more equal gender attitudes do not form part of a more progressive general viewpoint: attitudes towards religion, marriage and the family are not affected by the military death rate.¹³ It would be very interesting to investigate this finding further, particularly because several of the questions used to construct the family index are related to gender, bringing us again to the question of which attitudes change and which not, and asking why.

Migration

Migration is an important source of cultural change. Assimilation often occurs over time as migrants face new incentives, but there are also instances in which this doesn't happen as more segregated societies can keep social incentives frozen. Alternatively, when migrants occupy socially important position they can spread their cultural beliefs to the local population instead, as shown in Bazzi, Ferrara, Fiszbein, Pearson and Testa (2023) for the case of Southern White migration to the north in the 20th century and the subsequent spread of racist and conservative beliefs. Amalgamation of social norms can also occur when large

are identical, local variation in the proportion of women who work may matter.

¹³Gay (2023), online appendix, table A14.

parts of society face new incentives.¹⁴ Below I discuss several instances relating migration to cultural change, starting with settlement patterns.

Settlers to relatively empty areas can create distinctive cultures especially if there are no recognized pre-existing institutions to tilt culture in a particular direction. I start with two historical settlement experiences during a similar time period and which also feature skewed sex ratios: Australia and the westward expansion of the frontier in the US.

An intriguing study by Grosjean and Khattar (2018) examines the consequences of skewed sex ratios favoring men in Australia in the 18th and 19th centuries. Both European settlement in Australia (which was very restricted until 1830) and especially the convict population transferred there, were very male biased. Using the first census for each of 6 Australian states in the mid 1800s, the authors show that the average sex ratio was 3 men for every woman and remained high (1.9) during the entire historical period under study (1830 to 1881). It was much higher furthermore among convicts: around 30 men for every woman. Importantly, the authors argue that European migration (both convict and otherwise) came primarily from Ireland and England and hence was ethnically homogeneous.

Starting with the historical period, the paper shows that even after including county and year fixed effects, the (time varying) county male to female ratio is positively correlated with a higher marriage rate for women, a lower one for men, significantly lower labor force participation for married women, and a smaller fraction of working women in high ranking occupations (commerce and finance).¹⁵ The lower share of married women working outside the home is the result, according to the paper's interpretation, of their greater bargaining power in the marriage market due to the relative scarcity of women. This then could have given rise to, or reinforced, the idea later on that a married woman's place is in the home.

Next, the authors study the longer-term consequences of exposure to male-biased sex ratios at the postal area level using the nationally representative survey Household, Income and Labour Dynamics in Australia (HILDA) and the Australian Census (2011). HILDA

¹⁴See the excellent chapter by Bazzi and Fiszbein (2025) in this handbook for a fuller exploration of migration and culture.

¹⁵The authors also examine cross-sectional results, controlling for a series of covariates that can influence marriage and labor patterns via, for example, income effects (e.g., mineral presence and geographic controls). They also show that their results are robust to only relying on observations before the discovery of gold in 1851.

also asks questions relating to gender roles and time allocation. They use the sex ratio in the county from the first census in the state or county as their key explanatory variable, as well as a large set of individual, geographic, and historical controls. As identification relies on the spatial distribution of the historical sex ratio being random, a concern is that there may be unobservable characteristics that influenced the latter even after controlling for proxies for economic opportunities and total population at that time. To mitigate these concerns, the authors also use an instrumental variable strategy, instrumenting the overall sex ratio with that of a sub-population that was not free to move: the convicts. In both the OLS and IV specifications, the authors show that in those areas that historically had a greater male to female sex ratio, women now work less outside the home (both in the intensive and extensive margin), women (but not men) have more leisure time, and women spend less time taking care of children. It should be noted, however, that women do not have greater leisure time than men. The latter enjoy 4 more hours of leisure per week but this gap decreases by a bit over 50% with a one standard deviation increase in the sex ratio. Those counties also have a less progressive attitude towards women working outside the home and there are fewer women in high-rank occupations (controlling for the share of men in these).¹⁶

Altogether, these are very intriguing results. They show that despite sex ratios having long reverted to their normal levels, there are effects of the latter that persist over a century later. Although the authors do not remark upon it, women have significantly more progressive gender role attitudes than men. How this varies with the historical sex ratio is not studied. It would be interesting to know whether the gender gap in attitudes is unusually large relative to countries at comparable stages of development.¹⁷ In future work, understanding why there is meaningful heterogeneity in views will add an important dimension to the area of culture.

Despite the century of westward expansion in the US also featuring skewed sex ratios favoring men on the frontier, Bazzi, Brodeur, Fiszbein and Haddad (2023) find that this is not the main causal relationship for the more conservative gender attitudes found today in counties that spent more time being on the western frontier. Instead, the authors highlight

¹⁶The authors show that these results are not driven by education.

¹⁷See Bursztny et al. (2024) for how gender role attitudes vary by sex across countries.

the isolation women faced and argue that land abundance resulted in high fertility but little in way of support given the absence of extended family and friends. While female LFP was historically low, those women who did work had higher status occupations. Present day time-use surveys show that women in counties with longer frontier experience also have lower female LFP today but do not benefit from greater leisure, instead spending more time on household work. Thus, although the marriage market historically favored women, this did not result in greater bargaining power for women. The apparent contradiction between the Australian and US results show that there is no simple relationship between sex ratios and gender roles. The geographic and social isolation in which these women lived instead may have meant that greater female empowerment requires a social and community component to be successful.

Also in the context of migration, the paper by Agte and Bernhardt (2024) illustrates the importance of social incentives to conform to cultural norms. The authors make use of the large inflow of Hindu households at the end of the 20th century to villages in central India where Adivasi's were the majority. Adivasis are not in the caste system and hence do not have to follow the caste purity norms. They also have significantly higher acceptance of women's work. The existence of a river led to patterns of Hindu migration that resulted in village-level variation in the Adivasi population share. The authors show that this share matters: Hindus that live in villages where Adivasis form a larger share of the local population have less conservative gender norms, more in line with Adivasis' culture, and are less likely to adopt caste purity norms and food taboos such as prohibitions of consumption of meat. Furthermore, Hindu women are more likely to be in the labor force and men are less likely to report attitudes disapproving of women working. Interestingly, reflecting the importance of power for cultural adaptation, the authors find similar results even when Adivasis are a small share of the population but for historical reasons hold significant positions in the village, pointing to an insufficiently explored topic in this literature – the political economy of cultural change. There may be greater incentives to adopt different cultural practices when individuals who practice them hold greater power.

Forced migration has also been shown to affect cultural attitudes. Becker et al. (2020) study the expulsion of Poles from Kresy (then part of Poland) to the Western Territories

(from which Germans were expelled) that resulted from post World War II agreements on territorial boundaries. They hypothesize that the need to leave most possessions behind led to a greater appreciation of education among the expelled Poles and that this change in attitudes was transmitted to subsequent generations.

Using a variety of datasets, the authors establish that individuals with an ancestor from Kresy are likely to have more education today than Poles without such ancestors. Furthermore, using a survey conducted in the Western Territories on ancestry, they show that having a greater share of ancestors from Kresy or even any ancestor from Kresy (as compared with descendants of voluntary migrants from Central Poland or with autochthons), is associated with higher levels of education. Bolstering their hypothesis, they find that Kresy migrants are more likely to aspire to a higher level of education for their children and believe that freedom rather than material goods is the main condition for success in life. Importantly, they show that their results are not driven by selective out migration, differential fertility, or returns to education.

2.3 Learning

Culture in large part reflects our beliefs about what works, of what is right, of our understanding of others and of our environment. Exposure to new ideas, new facts, new arguments, or simply learning that others have different beliefs and/or take different actions can lead to cultural change.¹⁸ Below I discuss several examples from the literature illustrating how what I call “learning” has led to cultural change. This section first discusses a few examples of how the media can foment cultural change. Next, it illustrates how the creation of common knowledge can change outcomes, even if private beliefs have not necessarily changed. Lastly, it discusses how various experiences have affected cultural beliefs and outcomes.

¹⁸Not surprisingly, the educational system is a way in which “learning” can influence cultural beliefs. For example, Cantoni et al. (2017) shows that a staggered reform of the curriculum for high-school students in China resulted in greater trust of the government and greater skepticism about free markets but not in any meaningful change in their behavior (e.g., voting and participating in political organizations). See Cantoni et al. for a review of the literature on the effect of education on attitudes.

The Media

The media is undoubtedly an important source of information (or disinformation) with a profound capacity to mobilize people and influence people's actions and beliefs. Below I discuss a few papers that illustrate this force in a variety of contexts.

Television's importance in spreading ideas and images, both reinforcing traditional beliefs but also often subverting them, is widely acknowledged. In a developing country context, both La Ferrara, Chong and Duryea (2012) and Jensen and Oster (2009) show that access to particular programs – soap operas in Brazil and game shows in India – changed the behavior and beliefs of people whose lives were very different from those depicted. These television programs often center on relatively affluent characters whose lives are very different than their rural counterparts but whose behavior and attitudes may nonetheless exert influence.

La Ferrara, Chong and Duryea (2012) note the importance of the telenovela (soap operas) in Brazil, where a majority of the country would watch the novela at 8pm. Despite the overall high fertility rates in Brazil, during the time period under analysis over 70% of female characters portrayed in episodes produced by Rede Globo – the network that had a virtual monopoly” over the production of Brazilian novelas between 1965 and 1999 – had no children and around 20% had only one. The portrayal of smaller, upwardly mobile, middle and upper-middle class families, may have led women, particularly poorer and more rural ones, to want to emulate them insofar as that was possible. To identify the effect of novelas on fertility, the authors use the timing of Rede Globo into different areas of the country.¹⁹ Using the 1991 Census, their main analysis examines the probability that a woman living in a particular area gave birth in a given year as a function of whether the area was exposed to the telenovela a year before. Using a variety of both time-varying and time invariant controls, they show that the telenovela decreased fertility, especially for women with less education and lower wealth. Although they are unable to determine that this is indeed a change in cultural ideas about the ideal family size, they show that naming choices for children are affected by the introduction of Globo's soap operas.

¹⁹The authors argue that the timing of entry was related to clientelistic provision of licenses over this period.

Jensen and Oster (2009) also study the power of the media to spread ideas and information, thereby influencing beliefs and outcomes. In their case, the environment is rural India where the introduction of cable television can potentially affect the information available to people outside their own village. In particular, the most popular programs offered by cable TV were game shows and soap operas that, as in Brazil, depicted urban settings where women lead less traditional lives. Using a three-year panel data set in 5 states starting in 2001, and comparing the gender attitudes of women between survey rounds according to when and whether cable TV was introduced, the authors show that exposure makes women less tolerant of a husband beating his wife and decreases preferences for sons. Furthermore, it increases reported measures of autonomy related to decision-making in the household and the ability to move more freely (e.g., go to the market or visit friends and family without requiring permission).

The media may also change attitudes simply by making an issue more salient, causing people to reexamine their beliefs. That is the hypothesis explored by Fernández, Parsa and Viarengo (N.d.) to explain what has undoubtedly been one of the most rapid changes in culture: views regarding same-sex relationships. In the US, the percentage of individuals who answered the General Social Survey (GSS) question “Is it wrong for same-sex adults to have sexual relations?” with the response “not wrong at all” went from 11% in 1973 to 61% in 2022, an amazing 50 percentage point increase in just under fifty years.²⁰ As of 2024, 38 countries have legalized same-sex marriage and many more forbid discrimination on the basis of sexual orientation.²¹

In the US context, Fernández et al. show that the process of cultural change was highly discontinuous. Approval of same-sex relationships went from around 20% in 1973 to 59% in 2016.²² This was not a process of continuous incremental change, however. On the contrary, during the first twenty years of poll data, approval remained fairly stagnant at around 20%, only to jump discontinuously in 1992-93, increasing by over 11 percentage points. The authors note that 1992 was an election year (Bill Clinton vs George H. W. Bush), and in that

²⁰Data from the GSS from <https://gssdataexplorer.norc.org/trends>.

²¹See <https://www.hrc.org/resources/marriage-equality-around-the-world>, accessed October 2024.

²²The authors code responses as approving of same-sex relationships if they answer “not wrong at all” or “wrong only sometimes” as opposed to “always wrong” or “almost always wrong” to the GSS question “Is it wrong for same-sex adults to have sexual relations?”

year for the first time both parties took opposing stands on a gay-related issue – allowing gay people to serve openly in the military – in their respective national party platforms. Throughout the electoral campaign and the entire following year when Congress debated the issue, an unprecedented level of attention to gay-related events/issues (importantly, other than AIDS) was paid both by the big evening news programs of that time as well as by newspapers than at any time prior or after.²³

Why did increased attention to gay-related issues happen only with the national election of 1992? The authors point to a shock – the AIDS epidemic. This epidemic, which began in the US in 1981, galvanized the gay community around a common cause: finding a cure for AIDS and related issues such as providing health insurance to same-sex partners, speeding up the testing of new drugs, etc. Whereas before AIDS a gay person might favor the Republican party if they wanted lower taxes and another might favor the Democratic party if they wanted more redistribution and civil liberties, the epidemic made finding a cure a priority. The Democratic party recognized the existence of a non-negligible group of voters who were highly mobilized and willing to contribute money to fund the campaigns of pro-LGBTQ candidates. In this view, political mobilization by the gay community led, via national media attention and mainstreaming of the gay community by the Democratic party, to cultural change.

The authors hypothesize that the large increase in public attention paid to gay-related issues, effectively mainstreaming their coverage, led people to reevaluate their opinion of same-sex relationships, setting off a process of cultural change that continues to this day. In keeping with both contact theory and/or simply greater diffusion of information and higher interest, Fernández et al. argue that the impact of greater attention would be larger in places with greater exposure to the gay community. In the absence of measures of the size of the gay community in different places, they proxy exposure using variables that reflect (noisily) the proportion of the population that is gay: the cumulative AIDS rate by the end of 1992 at the state level (as measured by the Centers for Disease Control and Prevention – the CDC) and 1990 Census demographic information that allowed the identification of

²³Throughout the paper and here the term “gay” is used to refer to both men and women indiscriminately. The study ends in 2002 since in 2003 the Massachusetts Supreme Court ruled that the state constitution required it to recognize same-sex marriages, opening a new phase for gay rights and for the evolution of attitudes towards same-sex relationships.

households with same-sex cohabitants both at the state and county level.²⁴

Using a differences in differences strategy, Fernández, Parsa and Viarengo (N.d.) show that individuals living in places with greater exposure to the gay community had a larger increase in approval than those with lower exposure, in the 1990s relative to the 1970s, even after controlling for a rich set of individual characteristics, including household income, age categories, race, sex, and residential controls and the interaction of these with a time dummy for each decade. Importantly, they show that their results do not stem from a growing acceptance of civil liberties. Using a wide array of civil liberties questions from the GSS to build a general civil liberties index as well as a gay civil liberties index, they show that the evolution of attitudes towards gay civil liberties followed a time series that is similar to the civil liberties of other social groups during this period. Including the gay civil liberties index in the same-sex relationship approval regression analysis, however, does not lessen the importance of exposure to the gay community to explain their greater increase in approval in the 1990s.

The authors' exposure variables are undoubtedly a proxy for a bundle of variables. Places with greater exposure to gay individuals are also more likely to have more local news devoted to gay issues, more demonstrations in defense of gay rights, and its inhabitants would be more likely to have gay friends and acquaintances. While data limitations make it hard to unpack this bundle and pinpoint the exact mechanisms, Fernández et al. show that counties with greater exposure also showed a greater increase in their county-level newspaper coverage of gay issues in the post period (including county and poll year fixed effects). Furthermore, states with greater exposure saw a larger increase in contributions from LGBTQ groups to state-level electoral campaigns.

Interestingly, however, the paper concludes that the response to exposure was driven by women. That is, although both women and men became more favorable in the 90s, only women's approval responded to the degree of exposure (and the gender gap in approval grew whereas it was essentially zero in the 1970s). Why this was the case is unanswered although the authors attempt to explore various channels such as asymmetric responses due to the presence of children or to proxies for the degree of potential macho attitudes.

²⁴To obtain opinions about same-sex relationship at the county level, they turn instead to related opinion poll data from the American National Election studies – ANES.

A few lessons can be drawn from the papers discussed in this section. First, the media can be a powerful force in changing culture, purposefully or not. By portraying different lifestyles or by mainstreaming topics and debates, they may influence people's views of what is desirable or appropriate. Second, the degree of opposition may matter. In the case of same-sex relationships, for example, while AIDS itself did not increase people's approval of same-sex relationships, it did make people more aware that being a gay person was a widespread phenomena. Therefore, those who had or could potentially one day have a gay son or daughter, brother or sister, and friends and other relatives potentially had something to gain by from a more tolerant and accepting society. For other cases where cultural beliefs are girded by economic or power benefits for particular groups, the latter may be less susceptible to the media's influence. More generally, how the characteristics of the winners and losers from a potential change in culture in terms of power, institutional gatekeeping, and social connections affect the pace and direction of cultural change is a topic that deserves much more attention.

Correcting Beliefs

The fear of censure by friends, family, neighbors, or even strangers is an important reason that individuals keep behavior aligned with social norms. Thus, social norms and behavior may remain unchanged even if many individuals no longer hold those beliefs but they think that others do. These second-order beliefs can play an fundamental role in inhibiting cultural change. If, for some reason, information about others' beliefs became public, people could take very different actions. Below I discuss some papers that illustrate this important point.

Bursztyn, González and Yanagizawa-Drott (2020) conduct an ingenious experiment that demonstrates the potential importance of misperceived social norms in the context of Saudi Arabia. This is a country with very low female employment: fewer than 15% of women were employed in 2017 and even fewer worked outside the home. Furthermore, the tradition of male guardianship until recently required a husband's or father's consent to many activities, including work. In order to disentangle the role of mistaken second-order beliefs (i.e., beliefs about the beliefs of others) versus individual beliefs on married women's employment, the

authors conducted a field experiment. They recruited married men between the ages of 20-35 with at least some college education from several different neighborhoods and conducted 30-person sessions with them, grouping them by geographical area so that they shared a common social network. Framing their surveys questions as a labor market study, they solicited an individual's views regarding (i) whether they agreed with the statement "women should be allowed to work outside the home," and (ii) how many of the other 29 men they estimated would agree with the same statement. Individuals were then randomized into a treatment versus control group, with only the former obtaining information regarding the true share of men who agreed with the statement. All men were afterwards informed about an online platform that matched job-seeking Saudi women with jobs and then asked to choose between a \$5 Amazon gift card and the ability to sign up their wives to the platform.

The results are very interesting. First, there is a significant difference between the true percentage that agreed with the statement and the individual estimate. This difference is on average 24 percentage points more favorable towards women's work. Second, while in the control group 23% signed up their wives to the online job platform, in the treatment group 32% did so, a significant difference demonstrating the importance of the information. Several months later, the authors followed up with phone calls to the participants. They found that the percentage of wives who applied for a job outside the home increased from 5.8% (in the control) to 16.2% and that interviews for this type of job went from 1.1% to 5.8%, both significant differences despite the fact that the control group could have also gained some information during this time frame as neighborhood men are likely to have discussed the experiment. Unfortunately, however, there was no significant difference in whether a wife was actually employed outside the house (less than 10% are, both before and after). While this paper is not an example of cultural change in terms of measured outcomes, it does illustrate one of the obstacles that changes in social norms face: inaccurate information about the beliefs of others.

Information indicative of others' beliefs can also allow individuals to give voice and actions to previously unexpressed views. Bursztyn, Egorov and Fiorin (2020) argue that Donald Trump's relatively unexpected electoral victory in 2016 may have increased individual

willingness to express xenophobic views and decreased others' willingness to sanction such views. Using an online panel survey company that recruited participants from the seven counties that make up the Pittsburgh Pennsylvania metropolitan statistical area (MSA), the authors conducted a revealed preference experiment to demonstrate this in 2018. The authors randomize whether they tell the participants that Donald Trump won the election in "Pittsburgh's metropolitan area" or that Hillary Clinton won the election in "Pittsburgh's county" (both true) and they offer participants the ability to have the researchers make a small donation to a strongly anti-immigration organization. This donation would cost the participant nothing and would in fact increase their payment for participating by an extra \$1. Participants are also randomized into expecting their donation decision to be made public on a website shared with all participants from their area or being guaranteed anonymity.

The authors find that those who were in the public arm of the "Clinton won" treatment were significantly more likely to forego the donation bonus payment. The authors interpret this as suggesting that these subjects wanted to avoid the stigma associated with a public xenophobic action. Those in the "Trump won" treatment, on the other hand, showed no difference in behavior in the public vs private treatment arms. The difference between the Clinton and the Trump treatments stemmed entirely from the public arm (10 percentage points); the share that was donated when guaranteed anonymity was the same under both treatments (30%). This implies that those individuals who would have been willing to donate to the anti-immigrant organization became significantly less willing to do so when their names would be made public to a group that would likely be majority Clinton supporters. This experiment shows that public expressions of beliefs depend very much on what individuals perceive to be the culture around them – the beliefs of others. When individuals perceive the local majority as sharing their own views, they are much more willing to express them thereby influencing the public strength of these views.

To examine the change in willingness to sanction xenophobic behavior, they recruited participants from the US who identified as Democrats. These were randomized into being told that, in another survey, participants came from an area where Trump had won or where Clinton had won in 2016. Next, the subjects (Democrats) were informed that they had been

randomly matched with a participant from the previous survey and that the participant had chosen to make a public or private donation. The authors did not reveal that the donation was incentivized nor the existence of multiple treatment arms. Hence there were 4 possible match treatments: Trump vs Clinton won and donation information was public or private. Thus, the subject knew, for example, that they had been matched with someone who came from an area where Trump had won in 2016 and had chosen to, say, make a donation under conditions that stipulated that this donation was public. In this set up, the subjects were asked to play a dictator game in which they would decide how to split \$2 between themselves and the other player. Furthermore, they were told that their decision regarding how much to give was anonymous and that the other player had not known that this dictator game would subsequently be played.

The authors find that the average amount given in the dictator game is similar in all treatment groups (around \$0.79) except for in the Trump won, public matches arm. In that case, on average more was given (around \$.88). The authors interpret this the following way: first note that in the other 3 cases either the information was private or, if public, Clinton had won. Therefore, in each of these 3 cases if an individual had donated to an anti-immigrant group this was a true signal of that person's beliefs. In the case of a Trump won-public information person, a donation by that person may not be a true signal of their beliefs but rather a decision to respond to social pressure and donate to the anti-immigrant group. That may make the action more excusable to the subject in the dictator game, leading them to give more to that individual. In this view, Trump may have made public xenophobic actions more acceptable/less sanctionable even when presumably some significant proportion of the subjects do not share those views.

The two papers discussed above provide evidence on the importance of information and on how behavior may diverge depending on whether actions are public versus private, both potentially contributing to cultural change. New information (e.g., about men's average approval of working wives in Saudi Arabia) or socially significant information (e.g. whether a presidential candidate was popular where one lives) affect people's willingness to take actions previously believed to be culturally frowned upon (e.g., funding a xenophobic organization) as well as potentially affecting the willingness of others to punish these acts (e.g., as shown

in the case of the dictator game).

Much more work remains to be done in this area, both quantitatively and qualitatively. For example, it remains unclear how significant the misperception of beliefs is to Saudi-Arabian female employment. The average guess is that 63% of other session participants agree with the statement whereas the true level of agreement is on average almost 70%. Although the difference is larger conditional upon being in the same session (and therefore the same geographic area), the percentage of women who actually work in Saudi Arabia is very much lower than any of these numbers, suggesting that there are other forces at work. For example, a man may feel that a woman should be *allowed* to work outside the home but still strongly prefer that his own wife or daughter not do so. This may also explain the lack of a significant employment response 5 months later. Understanding the degree to which different segments of society matter to whom and under what circumstances is also potentially very important. For example, it would have been interesting to know how much *women* themselves *want* to work outside the home and whether the wage required to persuade them to do so is greater or smaller on average than that of their husband's or than the one prevailing in the market for women with their education. Furthermore, men may be very reluctant to have a working wife outside the home when very few other men are in this position, i.e. there may be a coordination problem even if beliefs are correct. This suggests that it may take time, as in Fernández (2013), to see significant changes in women's work behavior, as women (and men) learn about about the true social acceptability of working outside the home, which may evolve endogenously over time.

Exposure to New Experiences

Policies can also change people's beliefs, not only by changing incentives as discussed below in Section 2.4, but also by exposing people to experiences that they would not otherwise have. Random allocation of roommates in college, for example, has been shown to reduce racial stereotypes in South Africa (Corno, La Ferrara and Burns (2022)) and increase the support for affirmative action in the US (Boisjoly et al. (2006)). Student exposure to successful female graduates who had majored in economics in the same university increased female enrollment in more advanced economics classes and their likelihood to also major in

the field (Porter and Serra (2020), presumably changing beliefs about the attractiveness of the field for women. While the individual exposures above may be too small to call cultural change, they show the importance of an individual's experience in changing beliefs and preferences. Below I discuss a paper that gives an important example of how experience may change cultural beliefs.

Beaman et al. (2009) study the effect on attitudes and electoral outcomes stemming from the imposition of a quota for village council positions and for the leadership of the council (pradhan) in West Bengal, India. Starting in 1998, the policy randomly reserved one third of councilor and pradhan positions for women only (in unreserved councils, around 7% of had female pradhans in 1998). Which councils would be reserved was re-randomized prior to each election. The paper shows that in 2008, in those unreserved councils where the pradhan position had been previously been reserved for women in the past two electoral cycles, women were significantly more likely to stand for and win councilor positions.²⁵ The villagers were also given a (randomized) speech adapted from one in an actual village (different from theirs) in which a leader responded to complaints. The village leader was a woman in three of the randomly chosen audio recordings and a man in three other recordings. The respondent was asked to evaluate the leader's perceived performance and effectiveness (a vignette regarding resource scarcity and whether to invest in water or irrigation was also included with both the leader's gender and investment choice varied).

The authors find that in places that never had reserved female pradhans, men had a negative evaluation of the female leader. Women did as well, but the coefficient was statistically insignificant at conventional levels and the hypothesis that it was the same as men's could not be rejected. Importantly, the experience of having a female council leader erases this bias for men, as long as the village was reserved at least once, and the erasure of this bias persists even when the village no longer has a reserved position. Thus exposure can lead to learning about the true competence of women leaders, something that might not have happened in the absence of the quota policy. Interestingly, women did not significantly update their evaluation of female leaders. The authors speculate that this could be for a variety of reasons – from women paying less attention to policy or to women

²⁵There were only 3 cycles when the paper was written.

reacting negatively to a woman having an important position and perhaps making their own traditional roles have lower stature. Implicitly, this paper once again points to the importance of understanding how a change in social norms might be potentially resisted, at least at first, by those whose position may be threatened or simply not impact others if they ignore changes that are happening in their society despite the potential relevance to their lives.

2.4 Policies

Effective policies affect incentives and generate changes in behavior, potentially influencing individual and societal beliefs. This can be a dynamic process in which initially some individuals are motivated to change their behavior which then over time may further inspire others to do the same, generating a process of cultural change. Below I discuss a series of papers in which policy leads to changes in culture. Several of these papers trace out the dynamics of change by studying the evolution of outcomes over cohorts or peers in very different contexts – ranging from the US in the 1800s to contemporary Indonesia, Ghana, Norway, and Sweden.

Changes in Family Institutions

Policies can directly affect culture by changing the incentives of individuals to respect certain norms, especially if these are being eroded anyway. Ghosh, Hwang and Squires (2023) and Bau (2021) provide two very different examples of this in a historical and contemporary context, respectively. In this context, I will think of culture as a preferred form of extended family organization as in Bau and Fernández (2023).

It has long been thought that weaker family ties may be conducive to economic development.²⁶ In a context of growing urbanization, for example, individuals with close kinship ties may be reluctant to leave their farms or villages and prefer to reside with or close to their kin. While this behavior may be beneficial by, for example, providing insurance, it can hinder economic development if it leads individuals to forgo profitable outside opportunities.

Ghosh, Hwang and Squires (2023) explore a form of close kinship ties: cousin marriage.

²⁶See, e.g., Goody (1983), Henrich, Heine and Norenzayan (2010), and Enke (2019).

They use variation in cousin marriage bans across states in the US from 1850 to 1940 to examine their effect on a variety of individual-level outcomes. It is difficult to know what exactly motivated these bans, and though Ghosh et al. review the literature, they do not find evidence for what underlies this variation. Their preferred explanation is that the timing of cousin marriage bans may simply reflect the feelings of a particularly motivated individual or group of activists in a particular state at a particular time.

In the absence of direct measures of the number of cousin marriages, the authors construct one by using marriage records from 1750 to 1940 to calculate the excess frequency of marriages that occur with individuals with the same premarital surname, using the state and decade as the marriage pool.²⁷ Couples in cousin marriages, however, are different than non-cousin ones. Within the same state they are more likely to live on a farm, less likely to move from their state of birth, and have lower-earning occupations. Classifying surnames into a “high cousin marriage” (hereafter HCM) category if the excess frequency is above 10% before the first ban in 1858, they then conduct an event study by comparing outcomes for White men born in the same state and the same decade who have HCM surnames with those who don’t, before versus after the state ban. Their regression includes birth state x birth decade as well as birth state x HCM surname and census round fixed effects. The first fixed effect accounts for other things that could be happening at the state level during the time of the ban whereas the second accounts for potentially persistent differences by cousin-marriage category at the state level. Lastly, the census year fixed effects account for time-varying changes in the economy.

The authors start by showing that HCM surnames fall disproportionately after a state ban. Next, they show that a variety of geographic mobility measures increased for the HCM surname cohorts born a few decades after the state ban, relative to those couples who are not in the HCM surnames category. These HCM surname cohorts become relatively more likely to move out of farms and into urban areas, to live in places with greater population density, and to reside in a state other than their state of birth. Furthermore, using a variety of imputations on occupation to infer income, they show that HCM surname cohorts born a few decades after the ban increase their income by more than other surnames. This is

²⁷Excess frequency is that part of the proportion of marriages with the same surname that is greater than what would be explained by random matching.

very much in keeping with the idea that weaker kinship links may be conducive to faster development (here at the individual level) by permitting greater geographic and ultimately occupational mobility. This paper shows that policies, here in a very direct fashion, change cultural practices leading to changes such as residing close to home by changing incentives.

A different form of family culture is given by matrilocality and patrilocality. These refer to the practice of a couple residing with the parents of the bride or the groom, respectively. Bau (2021) studies the effect of the introduction of pension plans on these family cultural institutions and on human capital. The paper starts by noting that, traditionally, the practices of matri/patrilocality might have arisen or at least endured because they provide parents with old-age support later in life. In addition, in the absence of complete contracts between parents and their children, this practice allows imperfectly altruistic parents to capture some of the marriage market and/or the labor market returns to parental human capital investment in their children. It follows that the introduction of pension plans may decrease parental investment in a child's education and also weaken, on both sides, the desire to follow the matri/patrilocal tradition as the benefits to both parties has decreased.²⁸ Whose human capital should fall depends on which gender was favored originally by the tradition, i.e., girls for matrilocal ethnic groups (Indonesia) and boys for patrilocal groups (Ghana).

The author uses census information on ethnicity and language and maps this data to ethnic groups using the Ethnologue (Gordon Jr (2009)) which then is matched to the "pre-modernization" practices of different ethnic groups as coded in the Ethnographic Atlas (Murdock (1967)). In Indonesia, Bau distinguishes between traditionally matrilocal ethnic groups and those that practice patrilocality/neolocality whereas in Ghana she distinguishes between patrilocal and matrilocal/neolocal.²⁹ I will focus the discussion here on the findings for Indonesia. Bau starts by showing that prior to the policy, daughters between the ages of 5-22 are more likely to be enrolled in school than sons in that age group in the same

²⁸This is not exactly Bau's argument as she follows the Bisin and Verdier (2001) tradition of viewing the transmission of cultural traits as costly, which they may very well not be. In that view, a pension plan reduces the effort parents make to inculcate their children with their traditional practices.

²⁹The paper notes that those groups that traditionally patrilocal in Indonesia use neolocality as a secondary practice whereas those that practice matrilocality also use neolocality as a secondary practice. Neolocal practices refer to the couple residing in a new location with neither set of parents.

household, controlling for relevant variables such as parental education, being Muslim, age, geographic and other historical customs controls. The paper next turns to its main concern: studying the effect of the introduction of a pension plan on both investment in education and the cultural practices of matrilocality. The paper examines these questions via a triple difference approach that compares female cohorts by their relative exposure to the pension according to their birth year relative to the reform, the intensity of the treatment since the pension plan was rolled out geographically over time, and whether the individual belonged to a matrilocal ethnic group rather than a neolocal one.

The author shows that, indeed, women from matrilocal ethnic groups became less likely to complete secondary school and university, with an especially large effect for fully treated cohorts, i.e. those below the age of six when the policy was introduced. Furthermore, culture changed. These cohorts of women became less like to practice matrilocality and reside with their parents after marriage.³⁰ The paper finds symmetric results (though with adaptations given the nature of the Ghanaian data and policy intervention), for men from patrilocal ethnic groups in Ghana. This example shows that culture can respond rapidly within one generation, given a change in incentives, and impact significant economic outcomes such as human capital.

Gender Roles in Childcare

Even in countries in which female labor force participation and other measures correlated with gender equality are very high, women do a disproportionate amount of care work and housework. Various countries, mainly in Scandinavia, have instituted policies to change the distribution of care by altering their generous parental leave policies to incentivize men to take more time with their baby. The idea is that this might then change men's participation in non-market work more generally. Below I discuss several papers that examine the effects of these policies.

Dahl, Løken and Mogstad (2014) study the importance of peers for the takeup of parental leave in the context of a reform to Norway's parental leave. In 1993, fathers who had children born after April 1st of that year were eligible for an extra month of paid parental leave that

³⁰The practice of matrilocality appeared to be in decline, but this decline was larger for the most intensively treated women.

could only be taken by them, i.e., unlike the other 38 weeks which could be split as parents wished, these were earmarked for men only.³¹ Prior to this reform, only 3% of fathers took leave; after the reform, the take-up rate jumped to 35% that same year and continued to increase over time. To examine the importance of peers, the authors implement a fuzzy RD strategy using a window around the date of the reform, and study the peer effects of both coworkers and brothers. Below I discuss the results for coworkers as these permit the study of subsequent spillovers and have interesting heterogeneous results.

Dahl et al. restrict their sample to firms in which there was only one birth to male employees in the one-year window straddling the reform (6 months on either side) and for which there are male coworkers whose first child is born both after the reform and after the original peer. The reduced form examines the differential effect of the reform for a coworker with a peer who became a father in the window after the reform relative to a coworker with a peer who became a father in the window before the reform.

The authors find that the direct effect of the itself reform is large – around 31 percentage points. The peer effect, based on the regression described, is 3.5 percentage points, implying an 11 percentage point increase due to having a peer induced to take leave. The authors also study how the peer effect propagates over time. Assuming that the effect of each subsequent peer is additive allows them to estimate the additional effect of each cohort of peers. An alternative formulation allows for an imposed uniform rate of decay of a peer’s influence on subsequent coworkers. In both cases they show that a large part of fathers’ increased parental leave uptake over time (a five year period) is due to the snowball effect of new fathers taking up parental leave in the firm and these in turn also affecting subsequent take up of parental leave by their coworkers.

Dahl et al. attribute the peer effects to individuals learning about the consequences at the firm of a male coworker taking parental leave. A possible piece of evidence supporting this explanation is that when the peer is predicted to be a manager (the person with the highest or second highest earner in the firm), the effect on the worker is 2.5 times larger than otherwise. An alternative explanation which deserves to be explored in other contexts, however, might be that when more influential or higher status individuals take actions that

³¹There are some weeks reserved only for mothers. See their paper for these and other details regarding eligibility, etc.

challenge prevailing cultural norms, these receive greater weight. Of course, both learning how a workplace reacts to a man taking parental leave and seeing a potential role model flout gender conventions may contribute to the magnitude of the peer effect and therefore to changing culture. Overall, this paper illustrates an important mechanism that drives cultural change – how observing the actions of peers influences others, leading to cultural change over time.

Kotsadam and Finseraas (2011) ask whether the same 1993 reform in Norway resulted in changes in gender roles. Using a representative national survey from 2007, they study couples who had their last child within three years prior to the reform relative to those who had their last child within three years after the reform. They find that those who were affected by the reform are less likely to report conflicts over household division of labor and that they are also more likely to report equal sharing of clothes washing (but not of cooking and cleaning). In a subsequent paper on the same reform, Kotsadam and Finseraas (2013), examine the intergenerational consequences of Norway’s parental leave policy. Comparing adolescent children born three months before to three months after the 1993 reform, they provide evidence suggesting that men’s takeup of parental leave made it less likely that girls engaged in household work without affecting the probability of boys’ involvement in household work.

Farré et al. (2023) also study the intergenerational consequences of a parental leave reform in a setting – Spain – that is less gender equal than Scandinavia. They examine a March 24 2007 reform that gave men an extra 13 days of non-transferable full pay paternity leave.³² Prior to this, men had two days of non-transferable leave whereas women had 6 compulsory weeks at full pay. An additional 10 weeks, also at full pay, were free to be used as the mother wished, including transferring any portion of it to the father. The take up rate of paternal leave in Catalonia went from close to zero to 61% that year.

The authors conduct a survey of several high-schools in Catalonia in 2019-20, asking children between the ages of 11 and 13 a variety of questions related to gender roles. They use a window of 82 days before and after the reform and employ a difference in difference strategy by comparing these responses (treated vs not-treated windows) to those

³²This was then extended over time reaching 16 weeks in 2021.

obtained by children born in the same windows but a year earlier. Thus the identifying assumption is a common date-of-birth time trend in the outcome variables. The authors use a principal component analysis to construct a gender equality index. Controlling for individual background characteristics and school fixed effects, they find that children from the post-reform cohort became more supportive of women working while they had a child below school age. Furthermore, a greater share of them perceived as “socially appropriate” that a man work part time or not at all under the same circumstances.³³ Lastly, classifying household tasks as primarily male or primarily female according to the gender gap in doing these chores, they find that both boys and girls became more likely to engage in “counterstereotypical” household tasks.

Both Kotsadam and Finseraas (2013) and Farré et al. (2023) show that gender attitudes and the gender division of household chores can change as a result of relatively small changes in policy. It would be interesting to see whether differences in attitudes translate to differences in outcomes (e.g., choices regarding college majors and occupations) as these children age. It could very well be that any cross-cohort differences disappear. A finding of this nature would be hard to interpret, however, since both parents and peers affect individual behavior and, as these children age, their peer group will expand and there will be a large group of them whose fathers were affected by this reform.

3 Quantifying Mechanisms of Cultural Change

The papers discussed in preceding sections illustrate how cultural change occurred for a variety of reasons and in different spheres. Except perhaps in the short run, it is difficult to disentangle how much of a change observed in an outcome is directly due to changed incentives vis a vis intermediating mechanisms which may also alter outcomes and which may take longer to manifest. Quantifying the contribution of these mechanisms tends to require a model. While quantitative models do not provide the same level of rigorous identification as reduced form analysis, they can be very useful in providing plausible estimates of the importance of different forces and mechanisms and in clarifying their roles.

³³The meaning of “socially (in)appropriate” was described as “a behavior that most people believe to be correct or good (incorrect or bad).”

For example, while a model is not required to understand that both parents and peers influence cultural transmission, explicitly modeling this process can provide insights and quantify the relative strength of each force, especially in the medium and long run. An early example of this is given by Bisin, Topa and Verdier (2004). That paper estimates a model of marriage and socialization of children across religious lines (Protestants, Catholics, Jews, and “Others”) in the US in which parents of the same faith desire their children to share their faith and have an advantage, relative to mixed-faith parents, in socializing their children to have their same religious beliefs. They use the model to estimate the preferences of individuals from each religion about inter-marriage with each of the other faiths and study the dynamics that result from these preferences. The authors are able to reject a “melting pot” view of the US, as all faiths are estimated to have strong preferences not to assimilate via marriage. They also estimate that all other religious denominations have a high “intolerance” towards Catholics, leading over time to a decrease in their share of the population. Importantly, they show that their predictions can differ sharply from those that would be obtained by extrapolating from the results of a linear regression model.

Unfortunately, there are very few other papers that attempt to model cultural change and quantify both the importance of the shocks that brought it about and of the intermediating mechanisms. Below I discuss two contributions to this literature, both related to gender roles.

3.1 Modeling Cultural Change towards Female LFP

A defining feature of the twentieth century has been the vast increase in women’s labor force participation, especially that of married women. While undoubtedly many factors contributed to this change, it has been accompanied and facilitated by cultural change. Polls asking virtually the same question – “Do you approve or disapprove of a married woman earning money in business or industry if she has a husband capable of supporting her?” – went from below 20% approval in 1945 to over 80% in the 1970s for White women between the ages of 25-44 (Fernández (2013) p. 491).³⁴ Starting from the observation that the evolution of married women’s labor force participation follows an “S shape” as in the

³⁴The paper studies White women as the work history of Black women is very different as a result of slavery.

case of information diffusion (e.g., learning about new seed varieties or adopting a new household technology such as a refrigerator), Fernández (2013) studies the hypothesis that women’s LFP and opinions evolved endogenously as a result of an intergenerational learning process about the true cost of working.

In the model, women have the same initial prior about the true common disutility from working which can be either high or low. This can be thought of as the cost to the marriage or to a child’s welfare from having a working wife/mother, which historically were real concerns. They also have a private idiosyncratic disutility from working, known to each individual. Women learn about the true common value of working by observing their own private signal as well as a common public but noisy signal of how much other married women worked in prior periods.³⁵ Each woman updates the initial common prior with her own private information as well as with the public signal, leading to heterogeneous views about the shared common disutility of work. Thus, in the model, whether a married woman wants to work outside the home depends not only on her own idiosyncratic disutility of work, the wage of her husband, and her potential market wage, but also on her beliefs regarding the true common payoff from working.

The author shows that this model generically generates an S-shaped curve of LFP. If women initially assign a low probability to the true cost of working being low, they will require extreme private signals, on average, to be willing to work and only a small proportion of them – those who receive such extreme signals and have a low private disutility from working – will be willing to do so. This rational decision has a negative externality, however, as it decreases learning for future cohorts of women. Intuitively, this occurs because it implies that the proportion of women who would work if the true cost of working was low does not differ very much from the proportion that would work if the true cost of working was high since in both cases relatively extreme private signals are required. Observing a noisy signal of women’s aggregate LFP, therefore, is not very informative as it is difficult to discern if this was generated by a low or high common cost of women working. Over time, as more women work (assuming the common cost is low), learning accelerates as the

³⁵Note that if the public signal were observed without noise, the true cost of learning would be learned immediately by the next generation as a result of the law of large numbers. The noise in the signal can be justified by the existence of an unobserved aggregate shock as shown in the paper.

difference in the proportion of women who would work were the true cost low versus high increases, effectively reducing the noise in the aggregate signal. This process continues until many women work, whereupon once again learning slows down as it would require an extreme private signal for a woman not to work, thus making the aggregate proportion of women working once again relatively uninformative. Ultimately, this model converges to the truth. Thus, in this model both culture – the evolving common prior as to the true cost of women working – and female LFP co-evolve endogenously. Integrating evolving beliefs into a model of women’s labor force participation allows the (extensive) elasticity of labor supply to evolve rather than stay constant, consistent with the empirical evidence.

The author calibrates the model to married women’s LFP decennial census data from 1980 to 2000 as well as own and cross-wage (husband’s wage) elasticities and the intergenerational hazard rate (the probability that a daughter works if her mother worked relative to if her mother did not work here assuming that a daughter inherits her mother’s private information). Using decennial wage data for married men and women from 1880 to 2000 and having each generation make work decisions, the model does a fairly good job of generating 120 years of married women’s LFP. It can also track the poll data on approval of married women working by examining how beliefs evolve from 1940 onwards.

An important implication of the model is that increases in female wages (resulting from changes in technology, decreases in discrimination, structural transformation, etc.) have an externality. In particular, when the common prior is low, an increase in the return to women’s work will result in more of them working because fewer require extreme signals in order to work. This will increase the informativeness of the aggregate signal for future generations, accelerating learning. This is the dynamic effect of wage increases on female LFP. The paper then asks how much of the change in female labor force participation over 120 years (1880-2000) was due to the effect of higher wages on culture – the dynamic effect – versus simply making it more attractive to work (a static effect).

To answer this question, the paper proceeds first by keeping all wages at their 1880 level and allowing beliefs to evolve endogenously. The evolution of women’s LFP is then due only to the endogenous changes in beliefs resulting from constant wages. Next the paper asks how much of the change in each decade’s LFP is due to the “static effect” of wages. In this

exercise the path of beliefs is kept as generated in the first exercise, but now wages are no longer constant but instead replaced with the ones from the data. This quantifies the static effect of wage changes on women’s LFP for a given belief path. The last exercise allows beliefs to change endogenously along with the (exogenous) changes in wages, producing the increment in LFP that comes from the “dynamic” effect of wages on beliefs (the externality). The paper finds that all three forces played quantitatively significant roles, depending on the decade, with the dynamic effect of wages being especially large in the 1980s and 1990s.

A significant lesson from this paper is that cultural change can be a process of endogenous learning that occurs over time, changing the perceived desirability of some actions (here working) and giving rise to a co-evolution of beliefs/culture and behavior (here married women’s participation in the labor market). In this process the pace of change of both beliefs and outcomes is not constant over time, but changes endogenously as the pace of learning endogenously changes. Furthermore, this process can interact with other changes in the economy (in the paper these would be wage changes), affecting outcomes and beliefs beyond the initial static impact. Without a dynamic model of cultural change, these induced effects would be very difficult to identify and disentangle.

3.2 Modeling Cultural Change and the Gender Gap in Parental Leave

Although women in developed countries now have high labor force participation levels, care work remains reserved primarily for women and the “child penalty” is estimated to account for around 80% of the gender earnings gap.³⁶ Sweden has been at the forefront of efforts to increase men’s participation in childcare, starting in 1995 with the establishment of a “daddy month” which made one of the 12 highly paid months of parental leave non-transferable between parents. This meant that one month could be used exclusively by the father and one month exclusively for the mother; the remaining 10 months could be shared as parents wished.³⁷

Albrecht et al. (2024) examine a further reform that occurred in 2002 that increased the number of daddy months from one to two while simultaneously increasing the number of

³⁶See Kleven, Landais and Leite-Mariante (2024).

³⁷In addition to the highly paid months, there are an additional 3 months of parental at very low pay. In the model this pay is set to zero.

available highly-paid parental leave months from 12 to 13. As before, these non-transferable months are an endowment of each parent and are lost if that specific parent doesn't take them. The simultaneous one month increase both in total months and in reserved months implied that a woman who before was taking the maximum number of 11 such months could still take them but now there were two months that were essentially thrown away if the man did not use them. The effect of the reform was to increase men's parental leave on average by a bit over a month and decrease women's by somewhat under a month. Men's share of parental leave went from around 12% to closer to 20% during the years around this reform.³⁸

The objective of the paper is to understand the quantitative roles of the reform, economic incentives, and cultural change in producing these changes in parental leave outcomes. To do this, the authors develop a simple model in which parents jointly decide how many months of parental leave each should take while taking into account both economic and social factors. How parents share total parental leave time should presumably depend on the wages of each parent, their desire to spend time with their newborn, and the cost of childcare.³⁹ Given that before the reform there were 12 months of highly paid leave (around 85% of monthly wages including top ups by unions and employers) and afterwards there were 13 months of the same, why wouldn't a family use all their parental leave time? In particular, why wouldn't all fathers at a minimum take the daddy months rather than lose them?

Two considerations exist, one economic and one social. First, an additional economic factor the couple may need to consider are the potential signaling consequences of different spans of parental leave. In particular, the latter may be informative to an employer if it signals something about the worker's attachment to their firm, or their degree of ambition, or their attitude towards work. This could then affect their future wages and their probability of promotion. To quantify the wage penalty associated with different lengths of leave, the paper uses panel data to estimate separately by gender and education the percentage by which a individual's wage growth (between a year prior to birth and three years after birth)

³⁸See Albrecht et al. (2024).

³⁹Childcare is very expensive in Sweden before a child becomes one year old which is when public subsidized care becomes available.

is affected as a function of the months of parental leave taken. Second, there may be social stigma and challenges to identity, both in the workplace but also more generally, associated with being the primary caretaker for different lengths of time. In fact, prior to the 2002 reform, more than 20% of fathers took no leave, willingly forfeiting the daddy month. A few years after the reform, however, the authors show that the entire distribution of men's months of parental leave is shifted to the right, with far fewer fathers taking zero months but also many more taking 3-6 months. Women also reacted, with the distribution shifting to the left and fewer women taking the low-paid months of parental leave.

The authors model the social (cultural) impact on parental leave as affecting how parents preferences. In particular, while each each parent's enjoyment from spending time with their baby is a random draw from a normal distribution whose parameters depend on gender and education, the means of these distributions are allowed to evolve endogenously. Specifically, the authors assume that the mean is a linear function of the average parental leave behavior of an individual's peers a period earlier, here defined as those individuals of the same gender and household type.⁴⁰ This captures the idea that individuals tend to interact at work with people of similar education but also to socialize with others whose education is similar to theirs or to their partner's and therefore their behavior is what most matters.

The estimation of the steady states of the model (both pre and post reform) is disciplined by having it match over 50 key moments of the data both several years prior to the reform as well as several years after it and to do so by gender and household type. These include moments of the parental leave distribution by gender and household type both pre and post reform as well as the correlation, within each household type, of the household's ratio of the man's wage relative to the woman's wage and the man's share of parental leave. The authors show that the model does a very good job of matching the data both pre and post reform.

The quantification starts by considering the only the reform without incorporating any other changes, including not allowing culture to change (i.e., keeping the means of the normal distributions constant). They find that the reform on its own played a quantitatively

⁴⁰The paper distinguishes households by the education of each parent, where education is either at least 3 years of college (standard in Sweden for college) or strictly less than that, giving rise to four types of households.

significant role, accounting for around 8% to 20% of the final change. Next they evaluate the additional contribution of the income-related changes (i.e., means, variances, and correlations of household wage distribution, childcare costs, and wage penalties) but again keeping culture (the means of the taste distributions) constant. This further increases men's parental leave by around 10% and a similar, but negative, effect on women's leave. They then return to examining the effect of the reform keeping all income-related variables at their pre-reform levels, but this time allowing culture – the means of the distributions – to evolve endogenously. They find that this combination has a large additional contribution to the final changes in parental leave. Furthermore, for the first time, the percentage of men who take zero leave falls significantly. Lastly, and similarly to Fernández (2013), allowing both endogenous cultural change and the changed income-related parameters generates a further large change in parental leave behavior which is greater than the combination of either force separately.

The authors use the model to examine several counterfactual policies such as making the heavily subsidized childcare available before the child turns one and increasing the number of non transferable months. Interestingly, they find that decreasing the price of childcare to the same low level after the child turns one does not significantly change the distribution of parental leave between women and men. Increasing the number of non-transferable months to six, however, brings the share of parental leave taken by men much closer to fifty percent.

The papers discussed above illustrate the contribution of quantitative models to understanding cultural change. Although reduced form methods provide superior identification, quantitative models allow mechanisms to be understood more deeply, analyzed, and quantified. Important questions such as how different policies or changes in culture might affect policy require a quantitative model. Other questions, such as whether policies should attempt to change culture and, relatedly, how one should think about social welfare in the context of cultural change also require an explicit theoretical framework. Overall, the literature here is scarce and more work would be very valuable.

4 Conclusion

This chapter has provided a (partial) view of the emerging literature on cultural change. The first and longest section focused mostly on studies that identify specific shocks and provide evidence of cultural change. This area has produced the most compelling causal evidence to date. The second section turned to the quantitative modeling of the mechanisms that shape the evolution of cultural norms and behaviors. While this literature remains relatively underdeveloped, it offers essential tools for answering questions that purely reduced-form approaches cannot. Unfortunately, the emphasis on rigorous identification in applied microeconomics while highly productive has also constrained the scope of questions that economists tend to ask. Many of the most important questions about cultural change – its underlying dynamics, the feedback between policy and norms, or its implications for welfare – are impossible to address without theoretical and ultimately quantitative modeling.

A particularly important but underexplored area concerns the sources of resistance to cultural change. Cultural beliefs are rarely held uniformly within a society; instead, they are constantly contested, and the speed and direction of change are themselves endogenous. Those who benefit from existing norms – or at least believe they do – are more likely to resist change and may even push in the opposite direction by suppressing dissenting ideas and actions. Understanding heterogeneity in these beliefs, and how they map onto political and economic power, is crucial for identifying the constraints and levers of change. This calls for a political economy of cultural change: Who are the winners and losers from a potential transformation? How do institutions mediate cultural conflict? How do changes in one domain, such as gender norms in the household, affect adjacent domains like political attitudes, labor markets, or educational aspirations? And why do otherwise similar societies respond differently to comparable shocks? These questions remain wide open; hopefully future work will increase our understanding in these dimensions.

References

- Acemoglu, Daron and James A. Robinson. 2025. “Culture, Institutions, and Social Equilibria: A Framework.” *Journal of Economic Literature* 63(2):637–92.
- Agte, Patrick and Arielle Bernhardt. 2024. “The Economics of Caste Norms: Purity, Status, and Women’s Work in India.” *Working Paper* .
- Akerlof, George A. and Rachel E. Kranton. 2000. “Economics and Identity.” *The Quarterly Journal of Economics* 115(3):715–753.
- Albrecht, James, Per-Anders Edin, Raquel Fernández, Jiwon Lee, Peter Thoursie and Susan Vroman. 2024. Parental Leave: Economic Incentives and Cultural Change. Working Paper 32839 National Bureau of Economic Research.
- Alesina, Alberto, Paola Giuliano and Nathan Nunn. 2013. “On the Origins of Gender Roles: Women and the Plough.” *The Quarterly Journal of Economics* 128(2):469–530.
- Bau, Natalie. 2021. “Can Policy Change Culture? Government Pension Plans and Traditional Kinship Practices.” *American Economic Review* 111(6):1880–1917.
- Bau, Natalie and Raquel Fernández. 2023. Culture and the family. In *Handbook of the Economics of the Family, Volume 1*, ed. Shelly Lundberg and Alessandra Voena. Vol. 1 North-Holland pp. 1–48.
- Bazzi, Sam, Abel Brodeur, Martin Fiszbein and Joanne Haddad. 2023. “Frontier History and Gender Norms in the United States.” *Working Paper* .
- Bazzi, Samuel, Andreas Ferrara, Martin Fiszbein, Thomas Pearson and Patrick A Testa. 2023. “The Other Great Migration: Southern Whites and the New Right*.” *The Quarterly Journal of Economics* 138(3):1577–1647.
- Bazzi, Samuel and Martin Fiszbein. 2025. When do Migrants Shape Culture. In *Handbook of Culture and Economic Behavior*, ed. Ben Enke, Paolo Giuliano, Nathan Nunn and Leonard Wantchekon. Elsevier.
- Beaman, Lori, Raghavendra Chattopadhyay, Esther Dufo, Rohini Pande and Petia Topalova. 2009. “Powerful Women: Does Exposure Reduce Bias?” *The Quarterly Journal of Economics* 124(4):1497–1540.
- Becker, Anke. 2025. “On the Economic Origins of Restricting Women’s Promiscuity.” *forthcoming, Review of Economic Studies* .
- Becker, Sascha O., Irena Grosfeld, Pauline Grosjean, Nico Voigtlander and Ekaterina Zhuravskaya. 2020. “Forced Migration and Human Capital: Evidence from Post-WWII Population Transfers.” *American Economic Review* 110(5):1430–63.
- Bisin, Alberto, Giorgio Topa and Thierry Verdier. 2004. “Religious Inter-marriage and Socialization in the United States.” *Journal of Political Economy* 112(3):615–664.

- Bisin, Alberto and Thierry Verdier. 2000. "A Model of Cultural Transmission, Voting and Political Ideology." *European Journal of Political Economy* 16(1):5–29.
- Bisin, Alberto and Thierry Verdier. 2001. "The Economics of Cultural Transmission and the Dynamics of Preferences." *Journal of Economic Theory* 97(2):298–319.
- Bisin, Alberto and Thierry Verdier. 2024. "On the Joint Evolution of Culture and Political Institutions: Elites and Civil Society." *Journal of Political Economy* 132(5):1485–1564.
- Boisjoly, Johanne, Greg J. Duncan, Michael Kremer, Dan M. Levy and Jacque Eccles. 2006. "Empathy or Antipathy? The Impact of Diversity." *American Economic Review* 96(5):1890–1905.
- Boserup, Ester. 2007. *Woman's Role in Economic Development*. Earthscan.
- Boyd, Robert and Peter J Richerson. 1988. *Culture and the Evolutionary Process*. University of Chicago Press.
- Bursztyn, L, A Cappelen, B Tungodden, A Voena and D Yanagizawa-Drott. 2024. "How Are Gender Norms Perceived?" *Working Paper* .
- Bursztyn, Leonardo, Alessandra L González and David Yanagizawa-Drott. 2020. "Misperceived Social Norms: Women Working Outside the Home in Saudi Arabia." *American Economic Review* 110(10):2997–3029.
- Bursztyn, Leonardo, Georgy Egorov and Stefano Fiorin. 2020. "From Extreme to Mainstream: The Erosion of Social Norms." *American Economic Review* 110(11):3522–48.
- Cantoni, Davide, Yuyu Chen, David Y. Yang, Noam Yuchtman and Y. Jane Zhang. 2017. "Curriculum and Ideology." *Journal of Political Economy* 125(2):pp. 338–392.
- Cole, Harold L., George J. Mailath and Andrew Postlewaite. 1992. "Social Norms, Savings Behavior, and Growth." *Journal of Political Economy* 100(6):1092–1125.
- Corno, Lucia, Eliana La Ferrara and Justine Burns. 2022. "Interaction, Stereotypes, and Performance: Evidence from South Africa." *American Economic Review* 112(12):3848–75.
- Dahl, Gordon B, Katrine V Løken and Magne Mogstad. 2014. "Peer effects in program participation." *American Economic Review* 104(7):2049–2074.
- Desmet, Klaus and Romain Wacziarg. 2018. "The Cultural Divide." *Working Paper* .
- Enke, Benjamin. 2019. "Kinship, Cooperation, and the Evolution of Moral Systems." *The Quarterly Journal of Economics* 134(2):953–1019.
- Falk, Armin, Anke Becker, Thomas Dohmen, Benjamin Enke, David Huffman and Uwe Sunde. 2018. "Global Evidence on Economic Preferences." *The Quarterly Journal of Economics* 133(4):1645–1692.
- Farré, L, C. Felfe, L. Gonzalez and P. Schneider. 2023. Changing Gender Norms Across Generations: Evidence from a Paternity Leave Reform. Working Paper 16341 IZA.

- Farré, Lídia and Francis Vella. 2013. “The Intergenerational Transmission of Gender Role Attitudes and its Implications for Female Labour Force Participation.” *Economica* 80(318):219–247.
- Fernández, Raquel. 2008. *Culture and Economics*. London: Palgrave Macmillan UK pp. 1–10.
- Fernández, Raquel. 2011. “Does Culture Matter?” *Handbook of social economics* 1:481–510.
- Fernández, Raquel. 2013. “Cultural Change as Learning: The Evolution of Female Labor Force Participation over a Century.” *American Economic Review* 103(1):472–500.
- Fernández, Raquel and Alessandra Fogli. 2009. “Culture: An Empirical Investigation of Beliefs, Work, and Fertility.” *American Economic Journal: Macroeconomics* 1(1):146–77.
- Fernández, Raquel, Alessandra Fogli and Claudia Olivetti. 2004. “Mothers and Sons: Preference Formation and Female Labor Force Dynamics.” *The Quarterly Journal of Economics* 119(4):1249–1299.
- Fernández, Raquel. 2007. “Women, Work, and Culture.” *Journal of the European Economic Association* 5(2-3):305–332.
- Fernández, Raquel, Sahar Parsa and Martina Viarengo. N.d. “Coming out in America: thirty years of cultural change.” *The Journal of Law, Economics, and Organization*. Forthcoming.
- Galor, Oded and Ömer Özak. 2016. “The Agricultural Origins of Time Preference.” *American Economic Review* 106(10):3064–3103.
- Gay, Victor. 2023. “The Intergenerational Transmission of World War I on Female Labour.” *The Economic Journal* 133(654):2303–2333.
- Ghosh, Arkadev, Sam Il Myoung Hwang and Munir Squires. 2023. “Economic Consequences of Kinship: Evidence From U.S. Bans on Cousin Marriage*.” *The Quarterly Journal of Economics* 138(4):2559–2606.
- Giuliano, Paola and Nathan Nunn. 2020. “Understanding Cultural Persistence and Change.” *The Review of Economic Studies* 88(4):1541–1581.
- Goody, Jack. 1983. *The Development of the Family and Marriage in Europe*. Cambridge University Press.
- Gordon Jr, Raymond G. 2009. *Ethnologue: Languages of the World*. Dallas: SIL International. 16 ed.
- Grosjean, Pauline and Rose Khattar. 2018. “It’s Raining Men! Hallelujah? The Long-Run Consequences of Male-Biased Sex Ratios.” *The Review of Economic Studies* 86(2):723–754.
- Guiso, Luigi, Paola Sapienza and Luigi Zingales. 2006. “Does Culture Affect Economic Outcomes?” *Journal of Economic Perspectives* 20(2):23–48.

- Henrich, Joseph, Steven J Heine and Ara Norenzayan. 2010. “The Weirdest People in the World?” *Behavioral and Brain Sciences* 33(2-3):61–83.
- Jensen, Robert and Emily Oster. 2009. “The Power of TV: Cable Television and Women’s Status in India.” *The Quarterly Journal of Economics* 124(3):1057–1094.
- Kleven, Henrik, Camille Landais and Gabriel Leite-Mariante. 2024. “The Child Penalty Atlas.” *The Review of Economic Studies* .
- Kotsadam, Andreas and Henning Finseraas. 2011. “The state intervenes in the battle of the sexes: Causal effects of paternity leave.” *Social Science Research* 40(6):1611–1622.
- Kotsadam, Andreas and Henning Finseraas. 2013. “Causal Effects of Parental Leave on Adolescents’ Household Work.” *Social Forces* 92(1):329–351.
- La Ferrara, Eliana, Alberto Chong and Suzanne Duryea. 2012. “Soap Operas and Fertility: Evidence from Brazil.” *American Economic Journal: Applied Economics* 4(4):1–31.
- Mailath, George J. and Andrew Postlewaite. 2003. “The Social Context of Economic Decisions.” *Journal of the European Economic Association* 1(2/3):354–362.
- Michalopoulos, Stelios and Melanie Meng Xue. 2021. “Folklore.” *The Quarterly Journal of Economics* 136(4):1993–2046.
- Murdock, George Peter. 1967. *Ethnographic Atlas*. Pittsburgh: University of Pittsburgh Press.
- Nunn, Nathan and Leonard Wantchekon. 2011. “The Slave Trade and the Origins of Mistrust in Africa.” *American Economic Review* 101(7):3221–52.
- Porter, Catherine and Danila Serra. 2020. “Gender Differences in the Choice of Major: The Importance of Female Role Models.” *American Economic Journal: Applied Economics* 12(3):226–54.
- Teso, Edoardo. 2019. “The Long-Term Effect of Demographic Shocks on the Evolution of Gender Roles: Evidence from the transatlantic Slave Trade.” *Journal of the European Economic Association* 17(2):497–534.