

DISCUSSION PAPER SERIES

No. 7965

DOES CULTURE MATTER?

Raquel Fernández

*DEVELOPMENT ECONOMICS,
INTERNATIONAL MACROECONOMICS
and PUBLIC POLICY*



Centre for **E**conomic **P**olicy **R**esearch

www.cepr.org

Available online at:

www.cepr.org/pubs/dps/DP7965.asp

DOES CULTURE MATTER?

Raquel Fernández, New York University, CEPR, NBER and IZA

Discussion Paper No. 7965
August 2010

Centre for Economic Policy Research
53–56 Gt Sutton St, London EC1V 0DG, UK
Tel: (44 20) 7183 8801, Fax: (44 20) 7183 8820
Email: cepr@cepr.org, Website: www.cepr.org

This Discussion Paper is issued under the auspices of the Centre's research programme in **DEVELOPMENT ECONOMICS, INTERNATIONAL MACROECONOMICS and PUBLIC POLICY**. Any opinions expressed here are those of the author(s) and not those of the Centre for Economic Policy Research. Research disseminated by CEPR may include views on policy, but the Centre itself takes no institutional policy positions.

The Centre for Economic Policy Research was established in 1983 as an educational charity, to promote independent analysis and public discussion of open economies and the relations among them. It is pluralist and non-partisan, bringing economic research to bear on the analysis of medium- and long-run policy questions.

These Discussion Papers often represent preliminary or incomplete work, circulated to encourage discussion and comment. Citation and use of such a paper should take account of its provisional character.

Copyright: Raquel Fernández

CEPR Discussion Paper No. 7965

August 2010

ABSTRACT

Does Culture Matter?*

This paper reviews the literature on culture and economics, focusing primarily on the epidemiological approach. The epidemiological approach studies the variation in outcomes across different immigrant groups residing in the same country. Immigrants presumably differ in their cultures but share a common institutional and economic environment. This allows one to separate the effect of culture from the original economic and institutional environment. This approach has been used to study a variety of issues, including female labor force participation, fertility, labor market regulation, redistribution, growth, and financial development among others.

JEL Classification: D01, D1, O10 and Z1

Keywords: beliefs, culture, norms and preferences

Raquel Fernández
Department of Economics
New York University
19 West 4th Street
New York, NY 10012
USA

Email: raquel.fernandez@nyu.edu

For further Discussion Papers by this author see:
www.cepr.org/pubs/new-dps/dplist.asp?authorid=116433

* The author wishes to thank Joyce Cheng Wong for excellent research assistance and the NSF and the Russell Sage Foundation for financial support. This is a draft chapter for of the *Handbook of Social Economics*.

Submitted 06 August 2010

1 Introduction

Societies differ markedly in their economic outcomes. This is evidenced in a variety of ways: from different choices of redistributive policies and social security provisions to differences in aggregate outcomes such as average savings rates, fertility rates, or women's participation rate in the formal labor market. As shown in cross-country opinion polls, social attitudes also vary. On average, across countries people hold different views of, for example, the role that luck versus merit plays in generating income, the degree of social obligation one has towards others, or the importance of thrift as a moral virtue. These differences in social attitudes tend to be correlated with the differences in cross-country economic outcomes. For example, countries in which people value thrift also tend to have higher savings rates. Guiso, Sapienza, and Zingales (2006) find that a one-standard-deviation increase in the share of people who value thriftiness is associated with an increase in the national saving rate of 1.8 percentage points.¹ Similarly, countries that hold a more traditional view of women's role tend to have lower female labor force participation and higher fertility. For example, using attitude data from the World Value Survey (WVS), one finds that the percentage of individuals in a country that think that housework is as fulfilling as having a job is negatively and significantly correlated with female labor force participation (LFP) across countries.² Lastly, countries in which people tend to think that luck plays a fundamental role in the income process also have higher redistribution. Alesina and Angeletos (2005) show that the share of respondents in each country who believe that luck determines income is highly correlated with that country's spending in social welfare as a proportion of GDP.

Is the correlation between social attitudes and economic outcomes due entirely to economic and institutional differences across societies or are potentially systematic differences in social beliefs playing a causal role? More generally, what

¹This is calculated from answers to survey questions from the World Value Survey.

²These calculation use data from the WVS and from the OECD as reported in Fernández (2007b).

role do differences in the distribution of social preferences and beliefs (what I will henceforth call *culture*) play in explaining the variation in economic outcomes be it at the level of countries, social groups (e.g., ethnic or socioeconomic groups), or over time?

For a long period of time, questions regarding the role of culture in economic outcomes were largely absent in economic research. This was primarily the result of the absence of an empirical methodology that would allow one to investigate this issue. In particular, it reflected the difficulty in finding an approach that was capable of distinguishing the effects of culture from those of the economic and institutional environment in which economic decisions are taken. Did differences in aggregate outcomes across countries, for example, arise because they had different economic and institutional environments or because social attitudes were different? Standard approaches to this question, such as the use of cross-country regressions on a large variety of variables that are meant to capture economic and institutional differences across countries, identify culture with the regression residual. This approach, however is fraught with problems of omitted variables and endogeneity, compounded by mismeasurement.

In the last decade there have been a variety of new approaches that provide more persuasive evidence that culture matters. Some of the evidence comes from historical case studies that have attempted to use “natural experiments” to identify the effect of culture (e.g., Botticini and Eckstein (2005) or Greif (1994)). Some evidence has been provided by experiments showing that, on average, individuals from different social groups play different strategies in games such as the dictator game or public goods game (e.g. Henrich, Boyd, Bowles, Camerer, Fehr, Gintis, and McElreath (2001)). Better instruments for culture have also strengthened the case in favor of culture’s impact on economic outcomes (see, e.g. Tabellini (2010) and Guiso, Sapienza, and Zingales (2004)). Finally a large portion of evidence has come from following what I have called “the epidemiological approach” (see Fernández (2008)) to which this chapter is mostly devoted. The epidemiological approach attempts to separate culture from the environment by studying the out-

comes of individuals whose cultures potentially differ, but in a common economic and institutional setting.

This chapter will primarily focus on the epidemiological approach to culture although some of the experimental and historical evidence will also be reviewed. A chapter by Guiso, Sapienza and Zingales in this handbook provides a more thorough review of the literature that uses instrumental variables, particularly for understanding social capital, and Fernández (2008) reviews several of the historical case studies. This chapter is organized as follows: the next section provides a definition as well as some historical evidence for how cultures differ, and reviews some of the experimental literature. Section three develops a theoretical framework for the epidemiological approach and discusses the empirical challenges in the context of an example. Section four reviews the epidemiological literature and the last section concludes.

2 Some Preliminaries

Before proceeding with a review of the literature on culture and economics, a definition of culture is useful, even if it is left somewhat vague.³ In general terms, we may think of culture as a body of shared knowledge, understanding, and practice. According to the Merriam Webster dictionary, culture is: “the integrated pattern of human knowledge, belief, and behavior that depends upon the capacity for learning and transmitting knowledge to succeeding generations;” and “the customary beliefs, social forms, and material traits of a racial, religious, or social group; (and) the set of shared attitudes, values, goals, and practices that characterizes an institution or organization.”

Economists model individuals as economic agents who make choices in an economic and institutional environment, given their preferences and beliefs. Consider two hypothetical societies faced with identical institutional and economic settings.

³There is no agreed upon definition. By 1950, Kroeber and Kluckhohn (1952) provided over 150 definitions.

Suppose that, despite these identical environments, these societies end up with different outcomes, reflecting the fact that their inhabitants made different choices. We would like to say that these choices differed because these societies possessed different cultures, i.e., because they differ in their distributions of preferences and beliefs across individuals. Thus, for the purposes of what I will be discussing, a more useful working definition is to consider *differences in culture* as systematic variation in beliefs and preferences across time, space, or social groups.

Why should societies differ in their distributions of preferences and beliefs? This can happen for a variety of reasons. One possibility is that differences arise because actions are taken in an environment that resembles a game with multiple equilibria. In this case, non-identical outcomes are simply the result of the different strategies chosen by individuals reflecting their different expectations about the equilibrium outcome. Alternatively, the agents across the two societies could possess different priors about, for example, the payoffs to various actions, which could have resulted from different histories (obtained, for example, from different realization sequences of aggregate shocks).⁴

It may be useful to explicitly note here that nothing in this conception of culture considers it as either irrational, static, or slow changing. In particular, a definition of culture that considers the latter to be slow-moving (see, e.g. Guiso, Sapienza, and Zingales (2006)⁵) is rejected. The speed of cultural change depends on how quickly social beliefs and preferences change over time, which in turn depends on the environment broadly speaking, including the opportunities which determine individuals' learning pace, their interactions with others, and particular historical experiences. A salient example of a cultural change that began slow and accelerated considerably is seen in the social attitudes towards married women working. As shown in Figure (1) below (from Fernández (2007a)), beliefs in the US of the propriety of a married woman working if she had a husband "capable

⁴See Fernández (2007a) for a model of cultural change as a process of endogenous intergenerational learning. If societies obtained different shocks, that would lead them to learn at different rates and would give rise to different actions on average.

⁵Guiso, Sapienza, and Zingales (2006) define culture as "those customary beliefs and values that ethnic, religious, and social groups transmit fairly unchanged from generation to generation."

of supporting her” evolved dramatically over the 20th century, going from under 20% of the population being in favor of this in 1936 to less than 20% being against it in the 1990s.




Figure 1: Fraction that approves of wife working if husband can support her. (Data Source: WVS.) Picture from Fernández (2007).

Different historical experiences have important repercussions on individuals’ beliefs and preferences. As shown by Alesina and Fuchs-Schundeln (2007), for example, Communism had a significant effect on the beliefs of people who lived under it. The authors study German attitudes towards the role of the state in two main areas of social security: the extent of state provision desired in case of unemployment or illness and the extent to which the state should provide financial security for families, old-age or for people needing care.⁶ They find that if an individual lived in East Germany prior to reunification, she/he is much more likely to favor government provision for financial security for all of the areas mentioned above, after controlling for traits such as age, education level and type, gender,

⁶The answers for each question ranged from 1 to 5 which correspond to “only the state,” “mostly the state,” “state and private forces,” “mostly private forces,” and “only private forces.”

number of children, marital status, occupation, income, among others. This is independent of whether the responder lived in the former East or West Germany at the time of the survey.

The allocation of land titles to squatters in Argentina in 1989, as shown in DiTella, Galiani, and Schargrotsky (2006), likewise provides vivid testimony to the power of past experience. Hundreds of squatter families occupied an area of wasteland in the outskirts of Buenos Aires which belonged to many different private owners. The government attempted to redistribute the land to the squatters by buying it, but not all private owners were willing to sell. The squatters who had settled on tracts bought by the government obtained full property rights. The authors argue that this can be viewed as a case of random assignment and they provide evidence that the family heads in the group that received land titles are similar in age, gender, education levels, and ethnic origin to the ones in the group that did not. Despite the economic similarities across the two groups (those with land titles and those without), answers to survey questions regarding individualism, materialism, and trust differed markedly across them, with the group that received property rights demonstrating beliefs which are more aligned with those of the general Buenos Aires population. Namely, the squatters who were granted property rights were more likely than their counterparts without these rights to believe that success can be achieved alone, that money is important to happiness and that one can, in general, trust other people. Interestingly, however, the beliefs of the two groups regarding the role of merit do not differ significantly (perhaps reflecting the role of luck in determining who obtained property rights).

A last example is provided by Giuliano and Spilimbergo (2009) who use survey questions to show that an individual's (regional) location at age 16 affects her/his adult attitudes. In particular, individuals who grew up in an area affected more severely by recession were more likely to believe in luck and redistribution and to have less confidence in institutions such as Congress and the executive branch of the federal government.⁷

⁷Specifically, the authors use the answers to questions in the GSS which asked the individual

In addition to history, there is abundant evidence that attitudes are transmitted from parents to children. For example, Dohmen, Falk, Huffman, and Sunde (2008) use German data to show that a child's propensity to trust and her/his attitudes towards risk (as measured in answers to survey questions at age 23) is strongly positively correlated with parental attitudes.⁸ Farré and Vella (2007) use a sample of mother-child pairs to show that mothers transmit their attitudes regarding women's role in the labor market to their children. On the other hand, Cipriani, Giuliano, and Jeanne (2007) do not find a significant correlation in the way in which parents and their children play public goods games, but their sample is quite small.⁹

As noted in the introduction, there is plenty of evidence that economic outcomes and social beliefs are correlated. At the national level, for example, the extent to which executives believe that labor relations are good is correlated with union density across countries as shown in Figure (2) from Aghion, Algan, and Cahuc (2008).¹⁰ A different example is provided by Figure (3) from Alesina and Giuliano (2007) that shows a negative correlation across between the strength of family ties and the ratio of girls to boys in tertiary education.¹¹

A significant correlation between attitudes and outcomes is also found at the individual level within the same national environment. For example, Vella (1994), using Australian data, shows that attitude variables are correlated with the extent

whether she believed in government intervention to reduce income inequality and improve standards of living, as well as questions that asked individuals to express the degree of confidence in various government branches and whether luck is a driver of success.

⁸The results for risk remain significant and strong even after controlling for region where individual lived for the last 15 years, religion, ethnicity, subjective health status, income, and years of schooling, as well as a dummy for whether the family lived in East Germany before 1989. Interestingly, the mother's attitudes towards trust have a much stronger correlation with those of the child and the effects of both parents' attitudes decrease with birth order.

⁹The authors conduct the experiment with 38 parent-children pairs recruited from the same public elementary school in Washington, DC.

¹⁰The authors use executives' responses across more than 50 countries to the statement "Labor/employer relations are generally cooperative" from the *Global Competitiveness Reports*.

¹¹To measure the strength of family ties, the authors use answers to a series of question in the World Value Survey that attempt to assess how important the family is in a person's life, the degree to which one should love and respect one's parents regardless of their characteristics, and whether parents have a duty to do their best for their children even at the expense of their own well-being.




Figure 2: Correlation between union density and executives' beliefs that the labor relations are cooperative. (Data Sources: OECD and GRC 1999 database.) Picture from Aghion, Algan and Cahuc (2008).

of a woman's involvement in market work and Farré and Vella (2007) find that a mother's attitudes towards working women is correlated with her daughter's labor market decisions as well as those of her son's spouse. Dohmen, Falk, Huffman, Sunde, Schupp, and Wagner (2005) show that risk attitude measures from survey questions in Germany are correlated with a variety of risky behavior including traffic offenses, portfolio choice, smoking, risk in occupational choice, participation in sports, migration and overall life satisfaction.

Of course, correlation does not imply causation. Before turning to the epidemiological approach, the next section reviews some of the experimental literature that attempts to show the effect of culture by comparing the decisions of individuals from different societies who face identical controlled environments.




Figure 3: Relationship between strength of family ties and the ratio of girls to boys in tertiary education. (Data Source: WVS.) Picture from Alesina and Giuliano (2007).

2.1 Some Experimental Evidence

Experiments constitute an obvious methodological choice to investigate cultural differences as they can be transposed to various geographical locations and conducted with locally recruited samples.¹² Overall, however, the prevalence of small sample sizes and the fact that many experiments are conducted with college students makes it difficult to control for individual characteristics that may potentially differ in important ways across various groups. Below I give a brief review of some of the work in this area.

Evidence suggestive of cultural differences in players' choices of strategies in a given game is found by many authors. For example, Chuah, Hoffmann, Jones, and Williams (2007) and Chuah, Hoffmann, Jones, and Williams (2009) use the ultimatum game to investigate whether UK and Malaysian subjects exhibit differential behavior when bargaining within and across their national groups. They find stronger evidence of "home country bias" on the part of Malaysians. Namely,

¹²Roth, Prasnikar, Okuno-Fujiwara, and Zamir (1991) point out, however, various problems with experimental design in multinational experiments, namely, how to control differences in languages, currencies and experimenters.

Malaysian students offered higher shares to their countrymen than UK proposers gave to theirs, whereas Malaysians gave lower offers to UK nationals than to their own countrymen. UK proposers, on the other hand, did not change their offers when bargaining with Malaysians. Neither nationality was punished or rewarded for using different strategies, as the authors found that the rejection rates were not different for the two groups.

Roth, Prasnikar, Okuno-Fujiwara, and Zamir (1991) compared how individuals in four international cities play market games versus bargaining games. Interestingly, they found no cultural differences in the behavior of individuals in market games, whereas there were significant differences in the way bargaining games were played, giving greater credence to a cultural explanation.¹³ Henrich (2000) finds that the Machiguenga tribe of the Peruvian Amazon, when playing the ultimatum game, makes significantly smaller offers than a control group in Los Angeles and that the former also has a lower rejection rate. In post game interviews, tribal members explained that they accepted low offers because they did not want to reject any money. The proposers also expected their low offers to be accepted.

Henrich, Boyd, Bowles, Camerer, Fehr, Gintis, and McElreath (2001) summarizes the results of experiments conducted in 15 small-scale societies in various countries. They had individuals play various games, including the ultimatum, dictator and public goods game and found important differences across societies in the average outcomes. They argue that this may reflect differences in culture that arise from different structures of production requiring a smaller or greater degree of cooperation among individuals.

Does the fact that different societies play these games differently reflect different cultural attitudes? Even leaving aside the (critically) important issue of whether these results are driven by systematic differences in individual characteristics, a meta-analysis of 37 papers conducted by Oosterbeek, Sloof, and van de Kuilen (2004), which includes 75 results from ultimatum game experiments, finds

¹³They study market behavior using first price auctions, where all buyers have the same valuation. As predicted by standard theory, they found that the seller obtained the entire surplus.

that differences in game outcomes are not reflected in variations in attitudes. The authors use answers to several questions in WVS to construct a measure of average attitudes across countries that reflect the respect for authority, trust, and competition. They regress the outcomes (e.g. the share offered and the rejection rate) on variables such as the amount offered, regional dummies, the Gini coefficient in the country, GDP per capita and the average attitude as constructed from the WVS. They tend to find that attitudes are insignificant in explaining the variation. Of course, it is quite possible that the attitudes chosen by the authors are not capturing the cultural features that are relevant for these outcomes or that the demographic groups from which the experimental subjects are drawn do not have the average attitudes of their countries. Nonetheless, this finding suggests that one must be cautious about the cultural interpretation of experimental results based on small samples and on subjects whose individual characteristics are not controlled for.

3 The Epidemiological Approach

The essence of what I call the epidemiological approach is the attempt to identify the effect of culture through the variation in economic outcomes of individuals who share the same economic and institutional environment, but whose social beliefs are potentially different. Very often, the focus is on the economic behavior of immigrants or their descendants, but this need not always be the case (see, e.g., Fisman and Miguel (2007) and Miguel, Saiegh, and Satyanath (2008)).¹⁴ This approach is reminiscent of that used by epidemiologists (hence the name) who, in order to attempt to distinguish the genetic contribution to disease from the physical (including cultural, e.g. diet) environmental contribution, study various

¹⁴ Studying outcomes for the second generation rather than the first-generation immigrants offers some advantages. It avoids some of the confounding difficulties that first-generation immigrants are more likely to suffer to varying degrees such as the ability to speak the host country language and the prevalence of ties with non-immigrating family members. These factors are likely to be less important for the second generation.

health outcomes for immigrants and compare them to outcomes for natives.¹⁵

To understand the strengths and weaknesses of an epidemiological approach to medical issues, suppose that the incidence of, say, heart disease differs markedly between two countries (the source and host countries). If the incidence of heart disease in immigrants converges to that of the natives in the host country, the difference between the two countries is unlikely to be driven by genetics and instead results from the environment. Failure to find convergence, on the other hand, does not imply the opposite. Even when the environment is the sole responsible, there are still many ways to sustain differential levels of heart disease. For example, cultural assimilation may occur slowly (for instance, if immigrants maintain the same dietary patterns as in the source country), or living in the source country at a young age may confer some degree of immunity, or selection into immigration may be correlated with a particular health outcome.

In economics, unlike in medicine, the epidemiological approach attempts to distinguish between cultural versus environmental factors contributing to individual variation (and thus the environment now includes the economic and (formal) institutional settings that may affect outcomes, but excludes culture). The reasoning underlying this strategy is that (i) parents transmit their cultural beliefs to their children; (ii) cultural beliefs vary across (immigrant) groups in a systematic fashion reflecting culture in the country of origin; (iii) individuals who live in the same country or in the same appropriately defined geographical area, face similar economic and formal institutional environments. The idea is thus that individuals from different cultures will take different actions despite facing identical environments.

The basic empirical exercise uses data on individuals that live in one given country but whose parents were born in some other country – the country of ancestry. With this data one can estimate the probability that an individual i

¹⁵See, for example, the classic study by Marmot, Syme, Kagan, Kato, Cohen, and Belsky (1975). The methodological basis for this approach to culture in economics is developed in Fernández (2008) and the explanation offered here follows closely the one laid out there.

from country-of-ancestry c takes some action, y_{ic} ,

$$y_{ic} = \beta_0 + \beta_1 X_i + \beta_2 Y_c + \varepsilon_i \quad (1)$$

in which X_i is a vector of individual characteristics, Y_c is a proxy for culture in country c , and ε_i is an error term. Thus, X_i can consist of demographic information such as gender and age as well as measures of household income, education, etc. The primary variable of interest is the one that attempts to capture culture in the country of ancestry. Although it is possible to simply use a country-of-ancestry dummy for this variable, a superior strategy is to use a variable that more directly reflects the cultural attitudes of interest. For example, if y_{ic} is a labor force participation decision for a woman whose parents were first generation immigrants, then Y_c could be the female LFP in her parents' home country.

One may question the epidemiological approach described above for a variety of grounds. First, parents are not the only (nor necessarily even the most important) transmitters of culture; the relationships and institutions of the local environment (schools, local institutions, neighborhood, etc.) will also impact an individual's beliefs. Culture, furthermore, is socially constructed: to be replicated, the behavior may require the incentives – rewards and punishments – provided by a larger social body.¹⁶ Second, although studying the descendants of immigrants rather than immigrants directly allows one to avoid some potential problems (see footnote 14), it also means that the impact of culture from the source country is likely to have been attenuated over time. Both of these factors will lead to an underestimation of the effect of culture on economic outcomes in the above specification. Nonetheless, for a wide variety of issues, there appears to be a significant correlation between attitudes in the home country and attitudes expressed by immigrants and their descendants. This can be seen, for example, in the attitudes towards trust as shown in Figure (4) from Guiso, Sapienza, and Zin-

¹⁶Fernández and Fogli (2009) show that, in fact, the impact of culture appears to be greater for the descendants of those immigrant groups that have a greater tendency to cluster in the same neighborhood.

gales (2006), or in the attitudes towards redistribution in Figure (5) from Luttmer and Singhal (2010).¹⁷ Third, immigrants (and their descendants) from different countries may face different economic and institutional environments within the host country. Fourth, immigrants are not a random sample of a source-country's population. I will discuss the potential problems that the last two concerns pose for the estimation strategy in greater detail below, in the context of an example.




Figure 4: Correlation between trust level of country of origin and trust level of immigrants relative to Great Britain. (Data Sources: World Values Survey, General Social Survey.) Picture from Guiso, Sapienza and Zingales (2006).

¹⁷ As a measure of trust, Guiso, Sapienza, and Zingales (2006) use the answer to the binary question in WVS "Generally speaking, would you say that most people can be trusted or that you have to be very careful in dealing with people?" to construct a dummy which takes a value 1 if the person answered that people could be trusted. Luttmer and Singhal (2010) measure preferences for redistribution using the European Social Survey (ESS). Individuals are given the statement "the government should take measures to reduce differences in income levels" and the responses are measured in a scale of 1 to 5, ranging from strong disagreement to strong agreement.




Figure 5: Preferences for redistribution by immigrant group and country of birth. (Data Source: ESS.) Picture from Luttmer and Singhal (2010).

3.1 An Example

In order to understand the strengths and weaknesses of the epidemiological approach to culture, I will develop it in greater detail in the context of a simple model of a married woman's decision to work in the formal labor market.

3.1.1 A Simple Model

Let's start with the work decision of a (married) woman i in country k . For simplicity, I model solely the extensive margin and treat the utility from consumption and disutility from working as separable. Thus, a woman's work decision can be thought of as the solution to the maximization of the following utility function:

$$U(c, v_i) = u(c) - \mathbf{1}v_i$$

where u is a strictly increasing, concave utility function, $\mathbf{1}$ is an indicator function that takes the value one if she works and zero otherwise, c is household consump-

tion, and v_i is woman i 's disutility from working.

A woman's consumption is the sum of her labor income (if she works), w_f , and her husband's wages, w_h (i.e., consumption is a public good at the household level). Husbands are assumed to always work. For simplicity, the level of wages is taken as exogenous and identical for individuals within the same country but is potentially different across countries. Thus,

$$c = w_{hk} + \mathbf{1}w_{fk}$$

The disutility of work, v_i , varies across women and is assumed to be a random draw from a country-specific distribution with mean m_k and variance σ^2 , with a cdf denoted by $G_k(m_k, \sigma)$. Thus, for simplicity, differences in culture across countries with respect to women's work are modeled as (exogenous) differences in the mean of the distribution that characterizes the disutility of working. Note that there are two sources of heterogeneity in the model: a within-country heterogeneity across women in the same country reflected in the fact that they obtain different preference draws from a given distribution, and cross-country heterogeneity reflected in the mean of the preference distribution from which women obtain individual-level preference draws.

Given wages in country k , w_{hk} and w_{fk} , and the distribution of preferences in the country, G_k , we can solve for the level of female labor force participation in that country, L_k . It is given by the cdf evaluated at v_k^* , i.e.,

$$L_k = G_k(v_k^*)$$

where

$$v_k^* \equiv v^*(w_{hk}, w_{fk}) = u(w_{hk} + w_{fk}) - u(w_{hk})$$

is the level of disutility from work in country k which makes a woman indifferent between participating in the labor market or not.

If, for concreteness, we assume that G is a normal distribution and $\Phi(x)$ is the

standard normal cumulative distribution evaluated at x , then $L_k = \Phi\left(\frac{v_k^* - m_k}{\sigma}\right)$. To summarize, the women that choose to work are those whose disutility of working lies below the critical level v_k^* , which depends only on the wages in that country, w_{hk}, w_{fk} .

Note that both culture and the economic/institutional environment play a role in determining the level of women's labor force participation. Culture matters since it shifts the distribution from which preferences are drawn, by changing m_k (without affecting v_k^*).¹⁸ Cultures that have more negative views about women working, i.e., those with higher values of m_k , will have, ceteris paribus, lower female LFP, i.e.,

$$\frac{\partial L_k}{\partial m_k} = -\phi\left(\frac{v_k^* - m_k}{\sigma}\right) \frac{1}{\sigma} < 0 \quad (2)$$

where $\phi(x)$ is the pdf associated with standard normal distribution $\Phi(x)$.

Economic/institutional differences across countries also matter. In this simple framework these differences are reflected in wages and thus affect v_k^* . In particular, $\frac{\partial L_k}{\partial w_{fk}} = \frac{1}{\sigma} \phi\left(\frac{v_k^* - m_k}{\sigma}\right) u'(w_{hk} + w_{fk}) > 0$ and $\frac{\partial L_k}{\partial w_{hk}} = \frac{1}{\sigma} \phi\left(\frac{v_k^* - m_k}{\sigma}\right) [u'(w_{hk} + w_{fk}) - u'(w_{hk})] < 0$. Thus an increase in female wages will lead to an increase in a country's level of female LFP whereas an increase in the male wage leads to a decrease.

Next, consider a random sample of women from different countries of ancestry k , ($k = 1, 2, \dots, n$), all living in the same country j . Suppose that these women are identical in all but their cultural beliefs. In particular, suppose that they are endowed with identical husbands (i.e., they do not differ in their earnings) and that they face the same formal institutional environment so that their market wages, w_{fj}, w_{hj} , are the same. Hence v^* will be the same across all women, even though they have different countries of ancestry. We will also assume that culture is transmitted perfectly by parents, i.e., these women inherit the same v_i draws as their (foreign-born) mothers. The proportion of women who will work, however, will differ across countries of ancestry k since each of their v_i are drawn from different distributions. In particular, assuming a normal distribution for

¹⁸In more general models, culture will also affect v_k^* by affecting, for example, the supply of labor and through it, wages.

G , the proportion of women from ancestry k who work in country j is given by $L_{kj} = \Phi\left(\frac{v_j^* - m_k}{\sigma}\right)$. Note that, as shown in equation (2), women whose mothers were born in countries that have less favorable views of working women will be less inclined to work, as their disutilities v_i will be drawn from distributions with higher values of m . Thus, from a theoretical perspective, culture not mattering requires the distribution of distribution of preferences/beliefs to be identical across countries ($m_k = m, \forall k = 1, \dots, n$).

3.1.2 Empirical Issues

Because cultural differences (m_k) are not observable, in order to conduct an empirical analysis akin to equation (1) one needs to find variables which can function as proxies for cultural attitudes. In the context of the model above, a good proxy for a woman's culture could be the level of married women's LFP in her country of ancestry or a measure of attitudes in that country towards married women who work.¹⁹

There are several important issues that must be addressed before one can conclude that a statistically significant coefficient on the cultural proxy in equation (1) constitutes even moderately persuasive evidence that culture matters. First, there are many sources of heterogeneity across women other than their cultural beliefs. To the extent that these sources of heterogeneity are orthogonal to culture, it is simple to include them in the vector of individual characteristics. The comparative statics of equation (2) will then still be valid. Many of these characteristics, however, are endogenous outcomes that may well be influenced by culture. In the context of a woman's work decisions, her desire to acquire higher levels of education, the state/neighborhood where she lives, and the characteristics of her husband are a few of the more salient variables that may be influenced by culture. Thus, by including them in a regression one is effectively testing whether culture has an influence on work outcomes beyond the ways in which it is already reflected in these choices. It is important to note that in this case the failure to find a

¹⁹See Fernández and Fogli (2009) for the former and Fernández (2007a) for the latter.

significant coefficient on the cultural proxy is not an indication that culture does not matter.

Second, these women may not be randomly selected in the sense that their immigrant parents may be a selected sample from the distribution of beliefs in the country of ancestry. Thus, the cultural attitudes which were transmitted to their descendants may not be representative of the country's culture. Once again, a finding of an insignificant coefficient on the cultural proxy cannot lead one to rule out the possibility that culture matters. The interpretation of a significant coefficient on the cultural proxy, on other hand, depends on the issue being studied. In the case of women's work outcomes, selection would invalidate this finding only if parents from different countries came from systematically different parts of an otherwise identical distributions of beliefs and preferences across countries. For a positive coefficient on the cultural proxy of female LFP to be driven by selection would require countries to have identical distributions of preferences but that those parents (immigrants) from high female LFP countries be drawn from a different part of the distribution than the ones from low female LFP countries. In particular, it would require the former to be drawn disproportionately from the low disutility-of-labor portion of the distribution and the opposite for latter. How reasonable this possibility is depends upon the issue being studied. It is worth noting, in any case, that selection is a problem for all empirical methodologies, as even random experiments can suffer from the possibility of attrition with selection on unobservable variables.

Third, and perhaps the most critical issue, is whether there exists an omitted variable that varies in a systematic fashion across country of origin for purely economic reasons. In the context of our example of married women's LFP, variables such as her education, her husband's income, and her geographic location are all likely candidates that might reflect underlying economic differences across individuals rather than culture. Thus, controlling for these characteristics is important in this regard as well, despite their potential endogeneity.

While controlling for observed individual characteristics is straightforward (as-

suming that the data is available), the issue of unobserved heterogeneity remains; there may well be an omitted variable that is correlated with the country of ancestry. How to deal with this possibility depends on the question that is being studied. Tackling this issue is fundamental, however, as the ability of the new epidemiological literature to be persuasive depends on how well this issue is addressed.

In the case of culture and female LFP, the most likely candidate for an omitted variable is unobserved human capital. For example, it is possible that women with higher levels of (unobserved) human capital would choose to supply more labor to the market since their wages are higher. The most direct way to test for the presence of different human capital levels is via a Mincer regression on wages. After controlling for the usual variables in a Mincer regression (schooling, experience, experience squared and location), the cultural proxy should not have additional explanatory power for women’s wages. If it does, then it is more likely that unobserved human capital is responsible for the correlation between culture and women’s labor supply. Other possible variables that one can use to control for an individual’s unobserved human capital (assuming that individual IQ or individual test scores are unavailable) include proxies for the quality of education of the parents or parental human capital. In order to do this, one can employ either direct measures of the latter’s education or average test scores on standardized international tests (a la Hanushek and Kimko (2000)) in the country of ancestry as a measure of the parents’ quality of education.²⁰

Fourth, it should be noted that although using variables related to the economic outcome of interest is in many ways a superior approach to the “black box” of a country dummy, there may be issues with this alternative. On the one hand, making use of an economic variable is preferable since it facilitates formulating alternative hypotheses regarding the critical issue of a potentially omitted variable. On the other hand, the variable used as a cultural proxy may itself not reflect

²⁰Fernández and Fogli (2009) conduct a large battery of tests, including the ones mentioned above, to persuade the reader that an omitted variable is not responsible for their results.

cultural differences across countries since this source of variation may be swamped by their economic and institutional differences. For example, one could imagine a country which despite having more conservative attitudes towards women working, could also have higher female wages or a better child-support mechanism, resulting in an overall higher female-LFP-rate than that in a less conservative country. Using alternative proxies (or more directly a country dummy) will eliminate this concern.

It should be noted explicitly that the epidemiological approach is biased towards finding that culture does not matter. As mentioned previously, the fact that parents are only one source of cultural transmission among many and that they may have cultural attitudes that differ from the average ones in the country of ancestry, implies that one is more likely to rule the cultural proxy insignificant. Thus, just like the absence of convergence in disease does not provide definitive evidence in favor of genetics, the absence of a significant coefficient on the cultural proxy does not imply that only the economic and institutional setting matters.

4 The Epidemiological Literature in Economics

The first paper to use the epidemiological approach was Carroll, Rhee, and Rhee (1994). They used individual-level data on immigrants to Canada to investigate whether cross-country differences in savings rates were culturally driven. They estimated individual consumption levels as a function of permanent income, demographics, and region of origin. The authors found that although recent immigrants tended, on the whole, to save less than native-born Canadians, their saving patterns did not vary significantly by region of origin. Given this negative finding, it is perhaps not surprising that for some time no further attempts were made using this methodology. In more recent years, however, the epidemiological approach has been used to study the impact of culture on various economic outcomes such as women's work, fertility, labor market regulation, corruption, redistribution, and financial participation to name a few topics. Below I review some of this literature.

4.1 Women’s Work, Fertility, and Gender Preferences

Not surprisingly, issues which concern women (e.g. female labor force participation and fertility) have been a popular focus for work in culture and economics since attitudes towards women have evolved significantly over the last century across most of the developed countries.

Reimers (1985) is an early attempt to examine the role of ethnicity in married women’s labor force participation in the United States. Using a standard regression approach with ethnicity dummies, she finds mixed results regarding the importance of ethnicity for female LFP. The women in the sample she studies have been in the US for varying time periods, however, which is perhaps partially responsible for these results.

Antecol (2000) uses male and female LFP in the country of ancestry to examine whether culture plays a role in determining the inter-ethnic gender gap in labor force participation rates in the United States. She studies both first-generation immigrants and second and higher-generation individuals, which she groups into the same category. The results are suggestive that ethnicity matters although the absence of key individual-level variables leaves open the possibility that the results could be driven by omitted factors such as education or differences in parental background that lead to systematic variations in unobserved human capital.

The concern above is mitigated in Fernández and Fogli (2009) who use various measures of parental education and unobserved human capital (including average test scores in the country of ancestry and wages as described in the prior section) to rule out this alternative transmission channel. They show that culture plays a quantitatively significant role in explaining variation in women’s work and fertility outcomes. The authors also examine whether it is her or her husband’s country-of-ancestry that drives their results. Interestingly, they find that both matter but, if anything, the husband’s culture has a larger impact on his wife’s labor supply than her own cultural background.

An alternative to proxying culture with aggregate economic variables from

the country-of-ancestry (such as female LFP above) is to make use of indicators of social attitudes prevalent in those countries. The important issue of reverse causality discussed previously is avoided by using the epidemiological approach. The first paper to do this is Fernández (2007a). She uses the attitudes towards women’s work expressed by individuals in the woman’s country of ancestry as a cultural proxy to study the work outcomes of second-generation American women. She finds that cross-European variation in answers to questions about women’s role in the 1990 WVS has explanatory power for the 1970 work outcomes of second-generation American women from these countries of ancestry, even after controlling for individual differences such as those in education, location, and husband’s characteristics.²¹ Figure (6) from Fernández (2007b) shows the raw correlation between the cultural proxy (in this case, female labor force participation in 1990 in the country of ancestry) and the 1970 work outcome for second-generation American women from that country of ancestry, measured in hours worked per week.²²

Another strand of literature focuses on the effect of culture over another important outcome for women, fertility. This literature includes, for example, Guinane, Moehling, and Ó’Gráda (2006) who study Irish fertility in the United States in 1910, and Blau (1992) who examines the fertility behavior of first-generation immigrant women in the United States. These investigations are based on immigrants directly and therefore face the usual issues associated with immigration such as selection and the possible disrupted and delayed fertility behavior. The analysis of Fernández and Fogli (2006) and Fernández and Fogli (2009) mitigates these concerns by studying second-generation American women. Using past values of the total fertility rates from the woman’s country of ancestry as a cultural proxy,

²¹The WVS statements with which individuals are asked to agree or disagree (with various degrees of intensity) are: 1. Being a housewife is just as fulfilling as working for pay; 2. Having a job is the best way for a woman to be an independent person.

²²Note that if the way in which culture evolves is relatively stable across countries, it is possible to use future levels of the outcome as a cultural proxy, e.g., female LFP in 1990 is used to explain work behavior in 1970.




Figure 6: Hours worked and culture (Female LFP in 1990). (Data Source: ILO and US Census.) Picture from Fernández (2007).

they find that the latter has explanatory power for fertility outcomes, leading them to conclude that culture plays an important role.

On a related issue, Almond, Edlund, and Milligan (2009) investigate the role of culture in son preference. The authors note that sex ratios at birth are above the biologically normal level in a number of Asian countries. They investigate whether this is a result due to traditional economic reasons associated with poverty and/or to the existence of a rural or discriminatory environment that renders sons more valuable, or whether it is instead due to culture. To do this, they use an epidemiological approach and study Asian immigrants to Canada. They find that sex ratios for these immigrants rise with parity (i.e. with the number of children) if there was no previous son. In particular, those families whose first two children were girls are significantly more likely both to have a third child and for that child to be a boy, if they originally emigrated from India, China, Korea, or Vietnam. Since these immigrants no longer live in rural environments and poverty is presumably no longer an issue, culture is likely to be responsible for these results.²³

²³Unfortunately, they do not control for household income explicitly.

4.2 Family Ties, Political Engagement, and Labor Market Regulation

The type of relationships people possess may have a cultural component which can affect economic outcomes. For example, the degree of attachment to one's family may influence one's political attitudes or lead to economic concessions as individuals may have greater stakes in remaining in the same location. Relatedly, the perceived quality of the relationship between workers and management may also have important economic consequences.

In a series of papers, Giuliano and various coauthors establish that there is cross country variation in how families are viewed and that these views are correlated with a series of political and economic outcomes. First, as shown in Giuliano (2007), the living arrangements of second-generation immigrants in the US tend to follow the cross-European cultural patterns from their country of ancestry. In particular, individuals of Southern European descent in the US are more likely to live with their parents during the ages of 18 to 33 than the descendants of immigrants from other European countries.

Second, Alesina and Giuliano (2009) use questions in the WVS to construct a measure of the average "strength" of family ties across European countries.²⁴ They first show that, within a country, individual answers to these questions have predictive value for an individual's political participation and general interest in politics.²⁵ Next, they follow the epidemiological approach by using the country-level measure of the strength of family ties as a cultural proxy for second-generation nationals. That is, within a given host country, say Germany, second-generation Germans are associated with the Turkish value of the strength of family ties if their parents came from Turkey and with the Italian value if their parents came

²⁴The authors use answers to a series of questions in the World Value Survey that attempt to assess how important the family is in a person's life, the degree to which one should "love and respect" one's parents regardless of their characteristics, and whether parents have a duty to do their best for their children even at the expense of their own well-being.

²⁵ To assess the latter, the authors use questions which ask respondents about their general interest in politics and their interest in engaging in political conversations with friends. Political action is measured using a list of political activities that the respondent has engaged in.

from Italy. Using host country dummies and data from European Social Survey, the authors show that the cultural proxy has explanatory power for within-country variation in political attitudes of individuals from different countries of ancestry.²⁶ In particular, they find that second-generation immigrants are themselves less likely to be interested in politics if their father’s country of origin had a high average level of family ties. Their analysis includes a series of individual-level characteristics such as education categories, employment status, and a measure of family income. The fact that they consider second-generation immigrants across 32 destination countries strengthens the analysis as it is less likely that the results are driven by some special feature of a destination country. They interpret their finding as evidence of the importance of “amoral familism” in which strong family ties have a negative influence on social capital.

The strength of family ties also matters for labor market outcomes. Using the CPS from various years, Alesina, Algan, Cahuc, and Giuliano (2010) show that second-generation Americans whose parents come from countries with stronger family ties tend to have lower geographic mobility, a higher probability of unemployment, and lower hourly wages even after controlling for individual characteristics such as age, education, marital status, gender, and number of children as well as state fixed effects. They interpret this result as evidence that individuals who have a more family-focused culture are less able to take advantage of labor market opportunities due to their lower willingness to move away from their family in response to adverse local conditions.

Culture can also impact the labor market at the institutional level. Aghion, Algan, and Cahuc (2008) argue that bad labor relations and low unionization rates lead governments to set more stringent minimum wage policies in order to better protect workers. In a series of cross-country comparisons, they first show that the stringency of the state’s regulation of the minimum wage in OECD countries is negatively correlated with both executives’ and workers’ beliefs in the quality

²⁶The authors use questions concerning political attitudes in the ESS, which are very similar to the ones mentioned in footnote (25). For further details, see Alesina and Giuliano (2009).

of labor relations, whereas the unionization rate is positively correlated with these beliefs. They next use an epidemiological approach to show that there exists a cultural component to an individual's attitudes towards unions and her/his likelihood of belonging to a union. In particular, they examine the relationship between two cultural proxies for the country of ancestry - union density and a composite measure of state regulation of minimum wage - and two outcomes for second generation immigrants in the US: the degree of confidence an individual expresses about labor unions as well as the probability that the respondent belongs to a union.²⁷ They control for various individual-level characteristics but their use of the General Social Survey significantly restricts the number of countries of ancestry (twelve only) and provides only rough categories for critical variables such as income. Thus, although their finding that union density and minimum wage legislation in the country of ancestry has a significant impact on both an individual's confidence in unions and her/his probability of participating in one (in the US) is suggestive, it is also open to other interpretations. For example, it may be that an individual's occupation may be more or less prone to being unionized in a way that is correlated with her/his country of ancestry.

4.3 Corruption, Redistribution, and Violence

Is there a link between culture and the extent to which countries engage in redistribution?²⁸ Luttmer and Singhal (2010) take a step towards establishing this link by using an epidemiological approach to show that individual preferences for redistribution exhibit a cultural component. They study (mostly European) immigrants to 32 European host countries and show that preferences for redistribution in the country of origin can help explain the variation in the immigrants' preferences for

²⁷The authors construct a composite index to measure the degree of state regulation of the minimum wage. It is a combination of stringency measures, such as the existence of minimum wage legislation, and the "level" of the minimum wage, which the authors measure as the ratio of the minimum wage over the median wage in the economy. For further details, see Aghion, Algan, and Cahuc (2008).

²⁸For of a review of this literature, see Alesina and Giuliano (this volume).

redistribution in the host country.²⁹ This result holds even after controlling for several individual characteristics such as income, education, employment status, and host country fixed effects.³⁰ The fact that the authors consider immigrants across 32 destination countries strengthens the analysis as it makes it less likely that the results are driven by some special feature of a destination country. The cultural effects are large in the sense that a one-standard deviation increase in the average preference for redistribution across birth countries is associated with a greater than one-standard-deviation decrease in the log of household income.

A rather different take on the epidemiological approach is Fisman and Miguel (2007). The authors investigate the parking behavior of United Nations officials in Manhattan. As in studies based on immigrants or their descendants, this work follows an epidemiological approach by studying a select group of individuals (UN officials) in the same geographical environment (Manhattan). Until 2002, diplomatic immunity protected U.N. diplomats from parking enforcement prosecution, so their actions were presumably constrained by cultural norms alone. The authors find that diplomats from countries with high levels of corruption (based on existing survey-based indices) accumulated significantly more unpaid parking violations.

Fisman and Miguel's finding is intriguing as it seems to indicate that countries with high levels of corruption also have cultures which facilitate corrupt behavior. Does the failure to pay parking tickets when one is not legally required to do so, however, indicate corruption? An alternative explanation may be that highly corrupt countries face a different set of social problems that are far more serious than parking violations, leading to a culture in which these comparatively trivial issues are ignored. Moreover, even if one accepts the authors' interpretation of their results, an important remaining issue is whether the UN officials from countries with different levels of corruption face different likelihoods of punishment at home. If they do, then it would be unclear whether culture or economic

²⁹Preferences for redistribution are measured by the average answer to the ESS question which asks respondents how strongly they agreed/disagreed with the statement that "the government should take measures to reduce differences in income levels".

³⁰A similar analysis, but for second-generation immigrants to the US rather than Europe, is performed by Alesina and Giuliano (this volume).

rewards/punishments underlie their findings since this would imply that the institutional setting in which these officials operate may not truly be one and the same (the UN and Manhattan) but may also involve the institutions from their country of origin.³¹

Violence may also have a cultural component. By studying individuals from different nationalities who are all involved in the same activity – soccer – Miguel, Saiegh, and Satyanath (2008) find an ingenious way to keep the environment constant. The authors examine the relationship between a country’s history of civil war and a soccer player’s propensity to engage in violence on the soccer field as evidenced in his incidence of yellow and red cards (indicating a violent foul) when playing in one of six major European leagues. These leagues include players from 70 countries and all continents.

Controlling for a variety of important characteristics such as the position and league played in, the number of games, the quality of play (goals scored), etc., Miguel et al. find that players from countries with higher civil war incidence accumulate a greater number of yellow and red cards. The inclusion of continent dummies to some extent helps rule out alternative explanations such as racial discrimination by the referees. While it may be that, as in the study of parking violations and corruption, different home institutions are responsible for this behavior (e.g., perhaps future coaching opportunities on a home team depend on the degree to which violence is punished domestically), this concern seems less pressing in this arena than in the former study.

4.4 Within-Country Migration: Shirking and Financial Participation

Different cultures can coexist within the same country, particularly across different geographical regions. Ichino and Maggi (2000) use movers from and to different regions of Italy in an attempt to investigate the role of culture in the higher inci-

³¹The authors only deal with this issue partially by ascertaining that the length of a diplomat’s tenure is uncorrelated with the number of parking violations early in her/his career.

dence of shirking found in Southern versus Northern Italian employees. As shown by the authors, the rate of absenteeism in the South is almost double that in the North of the country. The authors' results are suggestive of a role for culture in this phenomenon since, when faced with a common environment, and after controlling for several individual and local characteristics, individuals born in the South but working in the North continue to have greater shirking rates than comparable Northern workers.

Guiso, Sapienza, and Zingales (2004) also study movers within Italy to attempt to identify the effect of civic capital on financial development. They use indicators of how much people rely on financial markets, such as the use of checks, reliance on cash, stock holdings and access to credit markets, as these are presumably correlated with financial development. They measure civic capital in an ingenious fashion, using not only the degree of electoral participation but also the quantity of voluntary blood donations in each province.³² As in the prior study, the use of movers allows the authors to control for cross-regional variations in the efficiency of institutions.³³

Guiso et al. use a dummy variable for the individual's place of residence and another one for the individual's origin to identify the effect of her/his culture. They find that people who were originally from provinces with higher civic capital make larger investments in stocks, rely more on checks to settle transactions, and have easier access to loans. It should be noted that while the authors interpret the latter finding as resulting from trustworthiness, it is also consistent with discrimination.

4.5 Cultural Change and Changes in Economic Outcomes

As discussed previously, there is no reason to believe, a priori, that culture changes only slowly. Algan and Cahuc (2010) exploit time variation in measures of individual trust to show that trust can impact economic growth. Suppose that income

³²This is the number of 16 oz blood bags per individual in each province in 1995.

³³As an example of institutional cross-variation, the completion of similar courtroom trials can range from 1.4 to 8.3 years across different regions of Italy.

per capita Y in country c at time t can be written, in a cross-country regression form, as:

$$Y_{ct} = \alpha_0 + \alpha_1 S_{ct} + \alpha_2 X_{ct} + F_c + F_t + \epsilon_{ct}$$

where S_{ct} measures the country average of social attitudes of individuals who live in country c in period t ; X_{ct} denotes a vector of average characteristics of the population and past economic development of the economy; F_c stands for country fixed effects and captures all other time invariant specific features in the country such as legal origins, endowments, or past institutions with long-lasting effects; F_t stands for period fixed-effects common to all countries and ϵ_{ct} denotes an error term.

The problem with the specification above is that contemporaneous social attitudes, S_{ct} , are likely to be correlated with the unobserved error term (if, for example, higher per-capita income increases trust). Here is where the authors employ the epidemiological approach.³⁴ Assuming that contemporaneous social attitudes are formed both by attitudes inherited from previous generations as well as by the contemporaneous environment allows the authors to write:

$$S_{ct} = \gamma_0 + \gamma_1 S_{c,t-1} + \gamma_2 X_{ct} + \Phi_c + \Phi_t + \nu_{ct}$$

where Φ_c and Φ_t stand for country and time dummies respectively; $S_{c,t-1}$ denotes the social attitudes of the prior generation; and ν_{ct} is an error term. The assumption that social attitudes from period $t - 1$ do not directly affect Y_{ct} , along with the assumption that $\nu_{ct} \perp S_{c,t-1}$, allows the authors to identify the parameters of the system of equations.

Given that standardized cross-country databases on social attitudes of earlier generations are not available, the authors proxy the inherited attitudes of people living in country c at time t by the social attitudes that Americans born in the US inherited from forebears coming from country c . As shown by Guiso, Sapienza,

³⁴In fact, a prior version of this paper was titled "Social Attitudes and Economic Development: An Epidemiological Approach".

and Zingales (2006) (see Figure (4)), there is a positive correlation between the trust levels of immigrants and their descendants in the US and trust levels in the country of ancestry.³⁵ Using the fact that the GSS identifies whether one's parents or grandparents were born outside the US, the authors use variation in the arrival times of the individual's ancestors to the US to proxy for attitudes in two different time periods: 1935-1938 and 2000-2003. Note that this strategy deals not only with the lack of historical data on trust attitudes but also ensures that contemporaneous events that might affect attitudes in the country of ancestry do not affect the cultural proxy, which is the inherited portion of culture for second, third, and fourth generation Americans.³⁶

The authors first show that the level of trust transmitted from the source countries has changed over the two time periods. They then demonstrate that the change in trust explains a significant portion of the variation in change in per capita income for the 24 countries in their sample. This is an intriguing finding. The causal interpretation relies on inherited attitudes and contemporaneous economic outcomes not being codetermined by some common factors, however. The authors attempt to mitigate this concern by using longer time lags between the outcomes and the inherited attitudes. A theory that would allow us to understand why trust changed over time and to identify the sources of change in the data would further strengthen their finding.

³⁵The authors use the answers to simple binary question on trust. See footnote (17).

³⁶Assuming generations of 25 years, inherited trust in 1935-1938 would be selected in the beliefs of second-generation Americans born before 1910 (i.e. whose parents arrived for sure one generation before 1935), of third-generation Americans born before 1935 and of fourth-generation Americans born before 1960. In the same way, inherited attitudes in 2000-2003 are those inherited by: second-generation Americans born between 1910 and 1975, by third-generation born after 1935 and by fourth-generation Americans born after 1960. For the authors, second-generation Americans are those whose both parents were born in the US; third-generation Americans at least two grand-parents but not all immigrated to the US; and fourth-generation Americans had all grand-parents born in the US.



Figure 7: Correlation between change in income per capita and change in inherited trust between 2000 and 1935. (Data Sources: Maddison database and GSS 1977-2004.) Picture from Algan and Cahuc (2010)

5 Concluding Questions and Remarks

The empirical work on culture has evolved considerably over time. It is my belief that the evidence that culture matters for a large variety of economic outcomes is by now sufficiently strong that most readers would find it convincing. There are many exciting questions left open, however. We would like to understand, for example, how culture propagates and evolves. The evidence presented in this paper shows that cultural preferences and beliefs have a life of their own in the sense that, even when removed from the environment in which they originated, they continue to exercise influence over individual outcomes. The evidence also shows, however, that there is some convergence over time both in economic outcomes and in attitudes. This indicates, not surprisingly, that culture changes in response to a new environment. Culture and the economic environment are, moreover, unlikely to be independent variables. Take, for example, attitudes towards premarital sex; these are likely to depend on contraceptive technology, the availability of abortion,

and a woman's ability to support a family on her own.³⁷ Culture, however, also influences the economic and institutional environment. A culture that considers sex to be shameful is less likely to make contraception or abortion easily available.³⁸ Thus culture and the economic and institutional environment interact and influence one another. Studies of this interaction would be an important addition to the literature.

Related to the topic discussed above is the question of why culture sometimes changes quickly and at other times glacially. This may be, at least in part, a response to the pace of change in the technological environment. This is not the only possibility, however. As shown in Fernández (2007a), cultural change can also arise from people endogenously learning about their environment. The author develops a dynamic model of culture in which individuals hold heterogeneous beliefs regarding the relative long-run payoffs for women who work in the market versus the home. These women do not know the long-term consequences of market work for their marriages and their children's welfare. Their beliefs, however, and those of their descendants, evolve rationally via an intergenerational learning process in which they learn about the long-term payoffs from working by observing (noisy) private and public signals.

The process described above generically generates an S-shaped figure for female labor force participation, which is what is found in the data. The S shape results from the dynamics of learning. When either small or large proportions of women work, learning is very slow and the changes in female labor force participation are also small. When the proportion of women working is close to 50%, rapid learning and rapid changes in female LFP take place. Thus, a learning model is also able to explain why culture changes at times slowly and at other times quickly, giving rise to an evolution in social attitudes similar to that shown in Figure (1).

It is also important to gain a deeper understanding of when cultural differences

³⁷See Fernández-Villaverde, Greenwood, and Guner (2010) for an interesting study of this issue.

³⁸For example, although the FDA approved the first oral contraceptive in 1960, it was not until the Supreme Court's decision in 1972 that it became available to unmarried women in all states.

are simply manifestations of multiple equilibria versus when they reflect a deeper disagreement. As discussed in Postlewaite (this volume), for example, the concern with rank, which varies across societies, may not indicate fundamental differences in preferences but arise instead from selecting different equilibria in a model of multiple equilibria.

In the model developed by Cole, Mailath, and Postlewaite (1992), individuals have standard preferences over consumption and their children's utility. Individuals are assigned some initial distribution of wealth, and men in the first generation are arbitrarily assigned a social rank which has no assumed correlation with wealth. The social arrangement (i.e., the equilibrium behavior) prescribes assortative matching between men's rank and women's wealth, with the highest ranked man matching with the wealthiest woman, etc.. The punishment for violating this prescription (which only women would be tempted to do) is that the rank of the male offspring from such a union would be reduced to zero. This implies that these sons will be matched with relatively poor women. Thus, a woman will rationally choose to match with a less wealthy but higher-ranked man if the decrease in her son's future consumption due to her deviation is sufficiently large.

The authors show that the behavior described above in which rank matters (called "aristocratic matching") is an equilibrium for some parameters of the infinite-horizon model. There is also always an equilibrium, however, in which aristocratic rank plays no role and individuals sort simply on the basis of wealth, with the wealthiest man marrying the wealthiest woman etc. It is important to note that these two societies will look very different not only because they give rise to different marital patterns, but also because they will give rise to different savings and bequest levels. In particular, in the aristocratic matching equilibrium, parents have less of an incentive to leave a large bequest to their male offspring since the bequest itself does not change their son's match in the marriage market. This is not so in the equilibrium in which both sexes match on wealth alone. Thus, one would expect families to save more in the latter equilibrium, changing fundamental economic outcomes.

As noted in Fernández (2008) in a slightly different context, behavioral differences arising from true differences in preferences versus those which are simply manifestations of multiple equilibria may be more difficult to distinguish in reality. It is likely that over time, the concerns for aristocratic rank (like the preferences for blue eyes in Mailath and Postlewaite (2003)) evolve in such a way that they become incorporated in deep preferences or beliefs. This would make them more robust in the sense that small changes in the environment would not necessarily eliminate this equilibrium even if it were no longer tenable as equilibrium behavior with standard preferences. This raises the important question of where preferences and beliefs come from and the extent to which cultural transmission is purposeful, that is, optimizing on the part of an individual or her parents.^{39,40} Lastly, it should be noted that if cultural variations are manifestations of endogenous differences in preferences rather than reflections of either different priors (as in Fernández (2007a)) or multiple equilibria arising from standard preferences (as in the Cole, Mailath, and Postlewaite (1992) paper discussed above), this raises difficulties for welfare analysis. Once preferences are endogenous, the standard welfare theorems no longer apply leaving open the question of how policies should be evaluated. This is an important question that requires further investigation.

³⁹Cultural transmission may well be involuntary. Fernández, Fogli, and Olivetti (2004) show that whether a man's mother worked while he was growing up is positively correlated with whether his wife works, even after controlling for a whole series of socioeconomic variables. They interpret this as preference transmission, but whether it is voluntary – optimizing – or simply by example is an open question.

⁴⁰See Bisin and Verdier (2000) for a model in which parental effort affects the probability with which children inherit their parents' preferences. See Bisin and Verdier (this volume) and Postlewaite (this volume) for excellent reviews of this literature.

References

- AGHION, P., Y. ALGAN, AND P. CAHUC (2008): “Can Policy Interact with Culture? Minimum Wage and the Quality of Labor Relations,” mimeo.
- ALESINA, A., Y. ALGAN, P. CAHUC, AND P. GIULIANO (2010): “Family Values and the Regulation of Labor,” mimeo.
- ALESINA, A., AND G.-M. ANGELETOS (2005): “Fairness and Redistribution,” American Economic Review, 95(4), 960–980.
- ALESINA, A., AND N. FUCHS-SCHUNDELN (2007): “Goodbye Lenin (or Not?): The Effect of Communism on People’s Preferences,” The American Economic Review, 97, 1507–1528.
- ALESINA, A., AND P. GIULIANO (2007): “The Power of the Family,” Working Paper 13051, NBER.
- (2009): “Family Ties and Political Participation,” mimeo.
- (this volume): “Preferences for Redistribution,” in this volume.
- ALGAN, Y., AND P. CAHUC (2010): “Inherited Trust and Growth,” American Economic Review, forthcoming.
- ALMOND, D., L. EDLUND, AND K. MILLIGAN (2009): “Son Preference and the Persistence of Culture: Evidence from Asian Immigrants to Canada,” Working Paper 15391, NBER.
- ANTECOL, H. (2000): “An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates,” Labour Economics, 7(4), 409–426.
- BISIN, A., AND T. A. VERDIER (2000): “Beyond The Melting Pot: Cultural Transmission, Marriage, And The Evolution Of Ethnic And Religious Traits,” Quarterly Journal of Economics, 115(3), 955–988.
- (this volume): “Cultural Transmission and Socialization,” in this volume.

- BLAU, F. D. (1992): “The Fertility of Immigrant Women: Evidence from High Fertility Source Countries,” in Immigration and the Workforce: Economic Consequences for the United States and Source Areas, ed. by G. Borjas, and R. Freeman. Chicago: University of Chicago Press.
- BOTTICINI, M., AND Z. ECKSTEIN (2005): “Jewish occupational selection: education, restrictions, or minorities?,” Journal of Economic History, 65(922-948).
- CARROLL, C. D., B.-K. RHEE, AND C. RHEE (1994): “Are there cultural effects on saving? Some cross-sectional evidence,” Quarterly Journal of Economics, 109(685–699).
- CHUAH, S.-H., R. HOFFMANN, M. JONES, AND G. WILLIAMS (2007): “Do Cultures Clash? Evidence from Cross-National Ultimatum Game Experiments,” Journal of Economic Behavior and Organization, 64, 35–48.
- (2009): “An economic anatomy of culture: Attitudes and behaviour in inter- and intra-national ultimatum game experiments,” Journal of Economic Psychology, 30, 732–744.
- CIPRIANI, M., P. GIULIANO, AND O. JEANNE (2007): “Like Mother Like Son? Experimental Evidence on the Transmission of Values from Parents to Children,” IZA Discussion Papers 2768.
- COLE, H. L., G. J. MAILATH, AND A. POSTLEWAITE (1992): “Social norms, savings behavior, and growth,” Journal of Political Economy, 100(1092-1125).
- DiTELLA, R., S. GALIANI, AND E. SCHARGRODSKY (2006): “The Formation of Beliefs: Evidence from the Allocation of Land Titles to Squatters,” mimeo.
- DOHMEN, T. J., A. FALK, D. HUFFMAN, AND U. SUNDE (2008): “The Intergenerational Transmission of Risk and Trust Attitudes,” Working Paper 2307, CESifo.

- DOHMEN, T. J., A. FALK, D. HUFFMAN, U. SUNDE, J. SCHUPP, AND G. G. WAGNER (2005): “Individual Risk Attitudes: New Evidence from a Large, Representative, Experimentally-Validated Survey,” IZA Discussion Papers 1730.
- FARRÉ, L., AND F. VELLA (2007): “The Intergenerational Transmission of Gender Role Attitudes and its Implications for Female Labor Force Participation,” IZA Discussion Papers 2802, IZA.
- FERNÁNDEZ, R. (2007a): “Culture as Learning: The Evolution of Female Labor Force Participation over a Century,” mimeo.
- (2007b): “Women, Work and Culture,” Journal of the European Economic Association, 5(2-3), 305–332.
- (2008): “Culture and Economics,” in New Palgrave Dictionary of Economics, ed. by S. Durlauf, and L. Blume. Palgrave Macmillan, Basingstoke and New York, 2nd edn.
- FERNÁNDEZ, R., AND A. FOGLI (2006): “Fertility: The Role of Culture and Family Experience,” Journal of the European Economic Association, 4(2-3), 552–561.
- (2009): “Culture: An Empirical Investigation of Beliefs, Work and Fertility,” American Economic Journal: Macroeconomics, 1(1), 146–177.
- FERNÁNDEZ, R., A. FOGLI, AND C. OLIVETTI (2004): “Mothers and sons: preference formation and female labor force dynamics,” Quarterly Journal of Economics, 119(1249-1299).
- FERNÁNDEZ-VILLAVERDE, J., J. GREENWOOD, AND N. GUNER (2010): “From Shame to Game in One Hundred Years: An Economic Model of the Rise in Premarital Sex and its De-Stigmatization,” mimeo.
- FISMAN, R., AND E. MIGUEL (2007): “Corruption, Norms, and Legal Enforcement: Evidence from Diplomatic Parking Tickets,” Journal of Political Economy, 115(6), 1020–1048.

- GIULIANO, P. (2007): “Living Arrangements in Western Europe: Does Cultural Origin Matter?,” Journal of the European Economic Association, 5(5), 927–952.
- GIULIANO, P., AND A. SPILIMBERGO (2009): “Growing Up in a Recession: Beliefs and the Macroeconomy,” NBER Working Paper 15321.
- GREIF, A. (1994): “Cultural beliefs and the organization of society: a historical and theoretical reflection on collectivist and individualist societies,” Journal of Political Economy, 102, 912–950.
- GUINNANE, T. W., C. M. MOEHLING, AND C. Ó’GRÁDA (2006): “The Fertility of the Irish in the United States in 1910,” Explorations in Economic History, 43(3).
- GUIISO, L., P. SAPIENZA, AND L. ZINGALES (2004): “The Role of Social Capital in Financial Development,” The American Economic Review, 94, 526–556.
- (2006): “Does Culture Affect Economic Outcomes?,” Journal of Economic Perspectives, 20, 23–48.
- HANUSHEK, E. A., AND D. D. KIMKO (2000): “Schooling, Labor-Force Quality, and the Growth of Nations,” The American Economic Review, 90(5), 1184–1208.
- HENRICH, J. (2000): “Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining Among the Machiguenga of the Peruvian Amazon,” The American Economic Review, 90(4), 973–979.
- HENRICH, J., R. BOYD, S. BOWLES, C. F. CAMERER, E. FEHR, H. GINTIS, AND R. MCELREATH (2001): “In search of Homo economicus: Behavioral experiments in 15 small-scale societies,” American Economic Review, 91(2), 73–78.
- ICHINO, A., AND G. MAGGI (2000): “Work Environment and Individual Background: Explaining Regional Shirking Differentials in a Large Italian Firm,” Quarterly Journal of Economics, 115, 1057–1090.
- KROEBER, A., AND C. KLUCKHOHN (1952): Culture. Meridian Books, New York.

- LUTTMER, E. F., AND M. SINGHAL (2010): “Culture, Context, and the Taste for Redistribution,” mimeo.
- MAILATH, G., AND A. POSTLEWAITE (2003): “The social context of economic decisions,” Journal of the European Economic Association, 1(354-362).
- MARMOT, M. G., S. L. SYME, A. KAGAN, H. KATO, J. B. COHEN, AND J. BELSKY (1975): “Epidemiologic studies of coronary heart disease and stroke in Japanese men living in Japan, Hawaii and California: prevalence of coronary and hypertensive heart disease and associated risk factors.,” American Journal of Epidemiology, 102, 514–525.
- MIGUEL, E., S. M. SAIEGH, AND S. SATYANATH (2008): “National Cultures and Soccer Violence,” mimeo.
- OOSTERBEEK, H., R. SLOOF, AND G. VAN DE KUILEN (2004): “Cultural Differences in Ultimatum Game Experiments: Evidence from a Meta-Analysis,” Experimental Economics, 7, 171–188.
- POSTLEWAITE, A. (this volume): “Social Norms and Preferences,” in this volume.
- REIMERS, C. W. (1985): “Cultural Differences in Labor Force Participation Among Married Women,” American Economic Association Papers and Proceedings, 75(2), 251–255.
- ROTH, A. E., V. PRASNIKAR, M. OKUNO-FUJIWARA, AND S. ZAMIR (1991): “Bargaining and Market Behavior in Jerusalem, Ljubljana, Pittsburgh, and Tokyo: An Experimental Study,” The American Economic Review, 81(5), 1068–1095.
- TABELLINI, G. (2010): “Culture and Institutions: Economic Development in the Regions of Europe,” Journal of the European Economic Association, forthcoming.

VELLA, F. (1994): “Gender Roles and Human Capital Investment; The Relationship between Traditional Attitudes and Female Labor Market Performance,” Economica, 61, 191– 211.