

CULTURE AND ECONOMICS

E000282

Modern neoclassical economics has, until recently, ignored the potential role of culture in explaining variation in economic outcomes, largely because of the difficulty in rigorously separating the effects of culture from those of institutions and traditional economic variables. This article selectively reviews some recent attempts to empirically identify the effects of culture on economic outcomes and to answer the question, 'does culture matter and, if so, how much?' Open theoretical and empirical questions are discussed, including the relationship between culture and institutions.

Economic decisions are made within a social context; as Aristotle reminds us, man is a social animal. The relevance of this statement to economics, however, is far from clear. In what ways, if any, do we need to consider the social nature of man in order to study economic questions? This article attempts to provide a partial answer to this question.

Traditionally, economists seek to explain differences in economic outcomes by studying how agents, with given preferences and beliefs, react to changes in the policy environment, institutions and technology. At a deeper level than the taste for apples versus oranges, however, few would deny that preferences and beliefs must be, to some extent, endogenous. Our level of trust in others, the determinants of status in society, our beliefs about the correct trade-off between efficiency and equity, or the 'proper' roles for men and women, are all examples of beliefs or preferences that have differed across societies and over time. These beliefs and preferences impact on individual behaviour and how society allocates scarce resources. At the individual level they help determine whether a woman participates in the formal labour market and the career she follows, the extent to which racism is tolerated, or the degree of assortative matching on wealth in marriages. At a collective level, they help determine, for example, the range and depth of the welfare state, the legality of slavery, or the proportion of the budget that is dedicated to foreign aid.

Although at some general level few may disagree that preferences, beliefs, or values of the type discussed above are endogenous (and may therefore differ across societies), whether they have a quantitatively significant impact on economic outcomes is another matter. Do differences in beliefs and preferences that vary systematically across groups of individuals separated by space (either geographic or social) or time – what I shall henceforth term *culture* – play an important role in explaining differences in outcomes? (For the purposes of this article, I will not give a more rigorous definition of culture than the abbreviated one here. See Elster, 1989, for a discussion of social norms and culture and Manski, 2000, for a discussion of peer effects and social interactions.) Modern economics (as opposed to sociology or anthropology) has largely been, until recently, reluctant to investigate this question. Although in principle there is nothing non-standard about positing preference/belief heterogeneity among individuals to explain differences in outcomes, the Stigler–Becker dictum *de gustibus non est disputandum* (Stigler and Becker, 1977) and its assertion that ‘no scientific behavior has been illuminated by assumptions of differences in taste’ has cast a long shadow in economics. Thus, the main challenge faced by those who believe that culture might matter has been to find a convincing way to show that culture can be studied rigorously and, in particular, that it is possible to separate the influence of culture from institutions and standard economic variables. In this sense, running, say, cross-country regressions on variables that one suspects reflect cultural attitudes (for example, different savings patterns may reflect attitudes towards thrift) to study the effect of culture has long (and correctly) been considered unsatisfactory. Despite one’s best efforts to control for differences in countries’ economic environments, identifying the residual with culture is ultimately unconvincing. It is difficult, if not impossible, to summarize the economic environment faced by agents with a few aggregate variables. Thus, there are bound to be omitted variables and problems of endogeneity, which are all further confounded by mismeasurement.

Hence, despite a long history of writers on the relationship between culture and economics (which includes Marx, Weber, Gramsci, Polanyi, Banfield and, more recently, Putnam and Landes, among others), modern neoclassical economics has been by and large silent on the topic of culture and only in recent years have economists started to

think seriously again about how culture may help explain economic phenomena. In this article I will selectively review some recent attempts to empirically identify the effects of culture on important economic outcomes and to answer the question, ‘does culture matter?’ Answering this question affirmatively naturally leads one to explore the propagation mechanisms of culture, to theorize about the relationship between institutions and culture, and to investigate the dynamic of culture – all topics that I will briefly touch upon at the end.

Empirical evidence on culture

In this section I examine some of the recent evidence on the importance of culture for economic outcomes. For expository ease, I have divided the empirical evidence into that which uses survey data, evidence based on immigrants or their descendants (what I call the ‘epidemiological approach’), and historical case studies. There is also a small body of experimental work that, by showing that across societies there exist marked differences in how individuals play games such as the ultimatum, public good or dictator game, has also shed light on the relationship between culture and economics (see, for example, Henrich et al., 2001).

Survey-based evidence

Perhaps the most natural approach to doing empirical work on culture consists in using the beliefs expressed by individuals in surveys (for instance, the World Value Surveys) on a variety of issues as expressions of culture and correlating them with economic outcomes. This approach, however, must overcome the problem of reverse causality. That is, differences in beliefs may be solely a consequence of different economic and institutional environments. Hence, the use of instrumental variables is required in order to identify causality. Overall, this has been difficult to achieve.

As shown by Guiso, Sapienza and Zingales (2003), the intensity of religious beliefs and religious denomination are correlated with a variety of individual attitudes such as trust in others, government’s role, views of working women and the importance of thrift. Guiso, Sapienza and Zingales (2006) show that these attitudes, aggregated at the country level, are correlated with cross-country aggregate outcomes (for example,

savings, redistributive versus regressive taxation, and trade). In order to ensure that the reverse causality is not at play, the attitudes are instrumented, usually by the religious composition in the country. This work is suggestive but there are several concerns associated with it. In addition to questions about omitted variables, it is not clear that religious composition is a valid instrument since it may also help explain the aggregate outcome through other channels (indeed, the coefficients on the instrumental variable results tend to look very high relative to the ones obtained by ordinary least squares. Running regressions at the individual outcome level would be more convincing, but opinion surveys unfortunately tend not to have high-quality economic data (the World Value Survey, for example, classifies income levels into ten categories). Recent work by Guiso, Sapienza and Zingales (2005) on the relationship between trust and trade, instead instruments trust with the genetic distance between indigenous populations. This seems a promising avenue of research.

Tabellini (2005) takes a significant step towards overcoming some of the weaknesses discussed above. To study whether culture affects economic development across European regions, he also aggregates (at the regional level) individual responses from the World Value Surveys to questions about trust, respect and the link between individual effort and economic success. The scope for omitted variables is reduced by focusing on within-country variation in Europe (by including country fixed effects). The attitudes are then instrumented with historical variables, such as regional literacy rates at the end of the 19th century and indicators of political institutions in the period from 1600 to 1850. The author finds that the proxies for culture are quantitatively significant determinants of per capita GDP levels and growth rates across regions. It is possible of course that the instruments are not valid. For example, they could affect output directly via sectoral composition or public investment. The paper contains a good discussion of these and other alternative hypotheses.

The epidemiological approach

A very different approach to relying on opinion data is to examine the economic outcomes of immigrants or their descendants. This is reminiscent of the epidemiology literature that, in order to attempt to identify the contribution of the environment broadly

defined (namely, physical and cultural) relative to genes in disease, studies various health outcomes for immigrants and compares them to outcomes for natives (see, for example, the classic study by Marmot et al., 1975).

To understand the strengths and weaknesses of such an approach, suppose that the level of, say, heart disease differs markedly between two countries (the source and host countries). If heart disease in immigrants converges to that of 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. There are many reasons why the environment may be solely responsible and still 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.

The epidemiological strategy in economics has its own set of problems. In particular, it is important to recognize that immigrants may be subject to many shocks (language difficulties, worse employment opportunities, greater uncertainty and so forth) which cause them to deviate from their traditional behaviour. Culture, furthermore, is socially constructed: to be replicated, the behaviour may require the incentives—rewards and punishments—provided by a larger social body such as a neighbourhood, school, or ethnic network. Furthermore, immigrants are unlikely to be a representative sample of their home-country's population. Their beliefs, preferences, and unobserved differences in their economic circumstances may differ significantly from the country average. Lastly, the exposure of immigrants (or their descendants) to a different culture from the one prevalent in their country of heritage presumably weakens the latter's impact on their behaviour. Note that all the factors mentioned above introduce a bias towards finding culture to be insignificant. Thus, on the whole, comparisons of behaviour or outcomes across different immigrant groups are a very demanding test of the importance of culture. In epidemiology, when differences across groups remain, one must be careful not to conclude that genetics is determinative when the underlying cause may be cultural; in

economics, when significant differences are not observed, one must be careful not to rule out cultural forces.

In economics, the paper by Carroll, Rhee and Rhee (1994) is the first that, to my knowledge, follows an approach similar to the one described above. The authors are interested in exploring whether cross-country differences in savings rates may be culturally driven. Using individual-level data on immigrants to Canada, they estimate individual consumption levels as a function of permanent income (as captured by labour and asset income), the interaction of this variable with demographic variables, some measures of wealth, and finally the interaction of a region of origin dummy (and years since arrival to Canada) with their measure of permanent income. If there exist different cultural attitudes towards savings, and if this attitude is maintained in immigrants, then one should observe different propensities across immigrants, by region of origin, to consume out of permanent income (that is, the regional dummies should be significantly different from one another). The authors find that the saving patterns of immigrants do not vary significantly by region of origin. Recent immigrants as a whole save less than native-born Canadians, but there is no statistically significant difference in behaviour across immigrant groups.

There are several weaknesses in the data-set used in the study above that may bias it against finding results that show a significant impact of culture. Wealth, for example, is not well measured. In particular, as only South East Asia's saving rate differed markedly from those of other regions in the immigrant population (31 per cent relative to 18–20 per cent across the remaining regions), the small number of immigrants from this group in the sample limits the power of the test. Note also that, if the motivation to save more stems from the desire to provide one's child with greater status via a larger bequest, the incentive to do this may be much less marked in a society in which savings are generally low or in which status stems from consumption behaviour.

The epidemiological approach has recently been used to study the influence of culture for a variety of issues. For example, Antecol (2000) examines the gender gap in labor force participation, Ichino and Maggi (2000) study cultural differences in the propensity to shirk, Guiso, Sapienza and Zingales (2004) examine differences in social

capital, and Giuliano (2007) studies the cultural transmission of family living arrangements.

Fernández and Fogli (2005; 2006) use a similar, but arguably less problematic, methodology than Carroll, Rhee and Rhee by studying second-generation Americans in order to investigate the quantitative importance of culture. Their research focuses on the fertility and work behaviour of married second-generation American women (that is, women who were born in the United States but whose parents were born elsewhere). The use of second-generation immigrants attenuates the problems associated with the first generation's adjustment to a foreign setting (for example, language difficulties) and even some selection problems are less likely to play a role for the second generation. On the other hand, second-generation individuals have been more exposed to the new culture, and that will tend to diminish the role of culture from the country of heritage. Our hypothesis is that attitudes towards woman's 'proper' role in society and towards ideal family size are culturally different across countries and that this culture is likely to be transmitted intergenerationally and show up in systematic differences in female labour force participation (LFP) and fertility, even if individuals were raised in the United States.

In our 2005 paper, the challenge was how to best capture the attitudes towards women and family size in the parents' country of origin. We chose not to use country dummies (as in Carroll, Rhee and Rhee, 1994) but to instead examine whether past values of economic variables in the country of origin that should reflect this culture – in particular, past values of female LFP and total fertility rates (TFR) – are able to play a quantitatively significant role in explaining differences in outcomes across second-generation women in the United States. Our argument is that these economic variables reflect the institutions (for example, markets, legal framework, minimum wages and so on), the strictly economic environment (demand and supply, transportation costs, access to day care, for example), as well as the preferences and beliefs (that is, the culture) of individuals in the country making decisions at that time. If these variables are able to explain the behaviour of women who, by virtue of living in the USA and in a different time period, face different institutions and economic variables, then solely the cultural component of these variables should affect their choices. This is a more demanding test

that is superior to the ‘black box’ approach of using country dummies which leaves open the question of what it is about the country that matters to outcomes.

In individual level regressions, we find that our cultural proxies – past values of female LFP and TFR – help explain both how much second-generation American women work and their fertility. As our data-set – the 1970 US Census – does not allow us to control for family factors such as parental wealth, income, and education, we include the woman’s education, her spouse’s education, and total personal income (as well as location, age, and so on) in our regressions. By including these variables, the coefficient on the cultural proxy only captures the direct effect of culture rather than its full direct and indirect effects (for example, a woman who wants engage in market work is more likely to invest in education and hence, by controlling for education, we are eliminating the effect of culture on this variable), but this is preferable to not controlling for differences in parental background, other than culture, that may affect women’s work and fertility outcomes. We find that the cultural proxies still matter even after including these additional variables. Furthermore, the cultural proxies are quantitatively significant: a one standard-deviation increase in the corresponding cultural proxy is associated with approximately an eight per cent increase in hours worked per week and about a 14 per cent increase in the number of children. The forces of assimilation means that these numbers should be taken, if anything, as a downward biased estimate of the true power of culture in the original setting (that is, in the country of ancestry).

We also examine the most compelling alternative economic explanation for our results, namely, the hypothesis that these are driven by unobserved human capital. We do this by showing that the results are robust to the inclusion the country of ancestry’s level of per capita GDP in various years and to the years of education of immigrants (by country of ancestry) in 1940 (this remains the case when Hanushek and Kimko’s (2000) measures of education quality in the parents’ country of origin are included). We also demonstrate that the work cultural proxy does not have explanatory power in a Mincer wage regression which it would be expected to have if it captured unobserved human capital. Lastly, we show that the work cultural proxy is insignificant in explaining how much married second-generation American men work whereas the fertility cultural proxy retains its explanatory power. (If the work cultural proxy had a negative effect on how

much these men work, that might indicate a substitution effect. In our regressions, the coefficient is basically zero and insignificant.) This is important because it implies that there does not exist some omitted economic variable at the parental country-of-origin level that affects the productivity of both men and women and that helps explain how much they work.

The methods used in Fernández and Fogli (2005) could be profitably extended to examine other issues, such as entrepreneurship or savings behaviour. It might also be interesting to elaborate upon the recent approach by Algan and Cahuc (2006) that attempts to combine survey evidence with the epidemiological approach in order to study the effects of culture on cross-country labour market outcomes. Although this work is too preliminary to discuss in depth, using the attitudes of, say, second-generation Americans to instrument for the attitudes of individuals in the home country seems cleaner than relying on variation in religious denominations. As usual, the question will be whether there is some omitted background economic variable correlated with the country of origin (particularly given the quality of the survey data-sets) that could be driving the results, but it seems a promising avenue of research. As shown recently in Fernández (2007a) using the World Value Survey, the attitudes of individuals in the country of ancestry towards women's market work and housework has explanatory power for the work outcomes of second-generation American women in 1970.

Historical case studies

The analysis of historical episodes in which changes in either culture or environment yield 'natural experiments' is likely to add richness and depth to our understanding of culture and the economy. Greif's 1994 paper is probably the best-known work in economics that makes the link between culture and institutional development. In brief, Greif argues that cultural beliefs (collectivist versus individualist) are reflected in the different ways in which in the 11th-century Genoese traders and Maghrebi traders set up their trading institutions. Both groups of merchants required agents to conduct their business overseas, and in both cases there was an agency problem as the overseas agent might be tempted to cheat the merchant. Maghrebi traders set up 'horizontal' relations in which merchants served as agents for traders and vice versa.

Information was shared among merchants/traders and an agent who was dishonest with one merchant could expect to be shunned by other merchants. The Genoese, on the other hand, set up ‘vertical’ relationships in which individuals specialized as merchants or agents. Information was not shared among merchants. This led the Genoese to set up more formal enforcement institutions. The two different responses, argues Greif, then had important consequences once trading opportunities were expanded in previously inaccessible areas. The Maghrebi expanded trade using other Maghrebi agents whereas the Genoese were able to establish agency relations with non-Genoese, leading to very different economic development paths thereafter (see also Greif (2005)).

Another compelling example is provided by Botticini and Eckstein (2005) who present the thesis that an ‘exogenous’ cultural change gave rise to the pattern of Jewish occupational selection that we see to this day. They argue that with the destruction of the Temple in Jerusalem in 70 CE, the Pharisees became the dominant religious group and transformed Judaism from a religion based on sacrifices to one whose main rule required each male to read and to teach his sons the Torah. This reform was implemented in places where most Jews were farmers who would not gain anything from investing in education. When urbanization expanded many centuries later, Jews had a comparative advantage in the skilled occupations demanded in the new urban centres. Thus, culture – the religious requirement of reading skills for other than human capital reasons – gave rise to the pattern of Jewish occupational selection seen since the ninth century.

Theories of culture

Is it necessary to modify the standard economic model in order to incorporate culture? The answer definitely is ‘no’. What appear to be societal differences in preferences may only be choice of equilibrium strategies in a game with multiple equilibria and standard preferences. This is in fact the most common way to think about the role of culture in economics, and is fully in keeping with our working definition of culture as systematic differences (across groups) in preferences or beliefs. Here the heterogeneity lies in the expectations (beliefs) over the strategies that will be played in equilibrium. Hence differences in culture can be identified with, for example, which equilibrium we play in a static game (for example, do we drive on the right- or left-hand

side of the road) or the degree of cooperation ('trust') sustained in a repeated Prisoner's Dilemma game.

Within the 'culture as multiple equilibria' literature, I find particularly interesting the research that attempts to generate behaviour that looks like social norms (such as determinants of status). Take, for example, a dynamic matching model in which individuals who differ in wealth choose a partner with whom to match and obtain utility from joint consumption and the utility of their child. As shown in Mailath and Postlewaite (2003), in addition to an equilibrium in which there is assortative matching on wealth, there may also be an equilibrium with imperfectly assortative matching that depends also on non-economic characteristics such as whether one has blue eyes. In this equilibrium, blue eyes matter not because of their intrinsic value, but simply because the matching rule allocates, for the same wealth level, a wealthier partner to individuals with blue eyes. Thus, a woman would be willing to match with a man with blue eyes and slightly lower wealth than another man without blue eyes, because although she obtains lower joint consumption, there is a 50 per cent chance that her child would inherit blue eyes and hence a better match and higher consumption in the future. To an outside observer, it might therefore appear that in this society people had an intrinsic preference for blue eyes, although this inference would be incorrect.

Although the example above is interesting, its explanation for a particular social norm seems incomplete and intuitively less than compelling. The preference for blue eyes or light skin may perhaps initially come about as a choice among many equilibria and involve solely a calculation about the trade-off between one's own consumption and that of one's child (though that too seems doubtful and is more likely the result of a history in which these traits are correlated with higher status). Over the longer run, however, one may conjecture that what sustains these equilibria – what makes these cultural traits less fragile to perturbations – is that these calculations are embodied in the individual and in society as preferences and beliefs about the inherent superiority/desirability of such features. People come to prefer blue eyes; people become racist. Thus, what is missing more generally in the theory of culture is an analysis of how preferences and beliefs (about things other than equilibrium strategies) themselves evolve.

The hypothesis that certain features of culture (those that have greater depth than driving on the left or the right side of the street) become part of preferences and beliefs implies that they cannot be discarded easily simply because they are no longer useful or beneficial, though over time this will certainly lessen their appeal. In this way, the operation of culture may be clearest to perceive when it no longer serves any useful societal purpose or particular group interest but nonetheless, at least for some time, persists – for example, religious prohibition on eating pork. (One reason speculated for this prohibition is that consumption of undercooked pork is linked to trichinosis. It is now known that this problem can be eliminated, however, by thoroughly cooking the meat.) In the context of the matching example above, individuals may eventually be willing to match with lower wealth people with blue eyes because this matching rule is incorporated into preferences/beliefs over what type of mate is intrinsically better even if the benefit derived by passing this trait on to their offspring is no longer substantial (say, because family size falls and decreases the payoff from the inheritable trait relative to the decrease in immediate joint consumption).

So far, we have discussed differences in culture as systematic differences in preferences and beliefs without distinguishing much between the two. This is not accidental, since, in general, the distinction between preferences and beliefs for our purposes is rather fuzzy. Even for simple preferences such as the trade-off between apples and oranges, what one knows (or believes) about the nutritional contents of the two may affect how one ‘feels’ about them, as may any other mental associations (for example, whether one is considered more exotic, how they were grown and so forth). In general, there are few pure (or naive) preferences – what one thinks or believes influences how one feels (and the same may be true vice versa. See Damasio, 1995, for an interesting exposition of evidence in favour of the hypothesis that emotions affect – and in fact are necessary for – the ability to think well.). This is not to deny that people have some inherent tastes (for example, it is believed that human beings have a taste for fat, probably because of the evolutionary advantage associated with an inclination to eat meat in an environment in which protein and iron were not easily obtained).

For more complex questions the above is even more likely to be true. Consider, for example, the large increase in female labour force participation in the 20th century. Is

it that woman's disutility from market work decreased or that her beliefs about the meaning or consequences of her working that changed over time? The dichotomy between the two alternatives does not seem very useful in this case. If the focus is on understanding why actions change over time, then using standard preferences and modelling the evolution of beliefs as giving rise to changes in expected payoffs may be the more useful strategy (the latter is the approach taken by Fernández, 2007b, who shows that a model of the evolution of female LFP as an intergenerational learning process does a good job of replicating a century of US female LFP data). If instead one wished to understand the utility from a given action, particularly one in which identity is concerned, then incorporating cultural beliefs into preferences may be a better route (see, for instance, Akerlof and Kranton, 2000). For example, wearing a dress or having a woman as a boss may decrease a man's utility, independently of any expectations of future consequences, simply because it makes him feel (culturally) less masculine.

Culture and institutions

As seen previously, the main challenge faced by most empirical work on culture is to convincingly isolate its effects from the incentives provided by traditional economic variables and institutions. This should not be taken to mean that culture and institutions are independent variables. Indeed, one way to think about institutions is as congealed culture: that is, which institutions are set up and how these evolve depends not only on the problems faced by society (or by a particular group in society) at a particular moment in time but also the beliefs/preferences – the culture – that are prevalent. As elaborated on in our earlier discussion of Greif (1994), cultural beliefs (collectivist versus individualist), for example, were reflected in the different ways in which in the 11th-century Genoese traders and Maghrebi traders set up their trading institutions, leading to very different economic development paths thereafter. My hypothesis is that the reverse causality is also likely to hold: that is, not only does culture affect institutions but also institutions affect the dynamic evolution of culture. In this sense, work that attempts to establish whether institutions or culture are the most important determinants of economic development seems misconceived (see Fernández, 2007c, for a theoretical analysis of the

dynamic dependency of culture and institutions; also Bowles, 1998, for a review of some of the theoretical and empirical evidence on the effect of markets on culture).

Concluding remarks

The rigorous study of culture and economics is in its infancy. We would like to understand, for example, how culture propagates and evolves. In particular, what is the relative importance of family versus other institutions as cultural transmission mechanisms for different beliefs or in different environments? To what extent is cultural transmission purposeful, that is, optimizing on the part of an individual or her parents (as in Bisin and Verdier, 2000) or for a social group and to what extent is it involuntary? (Fernandez, Fogli and Olivetti, 2004, show that whether a man's mother worked while he was growing up is 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.) When and why does culture change abruptly whereas at other times it proceeds glacially?

The relationship between technology and culture also needs to be investigated. How does technology influence culture and how does culture shape technological change? Some papers (for instance, Greenwood and Guner, 2005; Greenwood, Seshadri and Yorukoglu, 2002) argue that sexual norms and female LFP changed because of changes in technology. These papers ignore, among other things, the endogeneity of demand for new technology. Despite the convenient simplification of treating technology as a primitive, it too is endogenous. The extent to which societies put resources into developing technology that 'liberates' individuals from household work, for example, depends on things such as whether slavery is available or whether women expect to work in the market or at home. Put differently, both the relative price of market versus household labour and the elasticity of labour supply depend on the institutions (for example, slavery) and expected division of labour (for example, clearly differentiated gender roles) that are in place. The opposite is also true – the extent to which one can substitute capital for labour, whether at work or at home, helps determine which institutions are viable and may determine the pace and ease with which beliefs or preferences change.

From a theoretical perspective, the endogeneity of preferences and beliefs raises difficult questions for welfare. How should we evaluate policies once we recognize that preferences can change? While this is indeed a vexing and problematic question for welfare economics, recognizing that man is a social animal that is capable of reflecting upon, and hence changing, his preferences and beliefs greatly enriches our view of ourselves and the world and within it the potential role of economic discourse.

Raquel Fernández

See also cultural transmission; social norms

Bibliography

- Akerlof, G. and Kranton, R. E. 2000. Economics and identity. *Quarterly Journal of Economics* 115, 715–33.
- Algan, Y. and Cahuc, P. 2006. Minimum wage: the price of distrust. Mimeo, CREST-INSEE.
- Antecol, H. (2000), “An Examination of Cross-Country Differences in the Gender Gap in Labor Force Participation Rates,” *Labour Economics* 7, 409-426.
- Bisin, A. and Verdier, T. 2000. Beyond the melting pot: cultural transmission, marriage, and the evolution of ethnic and religious traits. *Quarterly Journal of Economics* 115, 955–88.
- Botticini, M. and Eckstein, Z. 2005. Jewish occupational selection: education, restrictions, or minorities? *Journal of Economic History* 65, 922–48.
- Bowles, S. 1998. Endogenous preferences: the cultural consequences of markets and other economic institutions. *Journal of Economic Literature* 36, 75–111
- Carroll, C., Rhee, B. and Rhee, C. 1994. Are there cultural effects on saving? Some cross-sectional evidence. *Quarterly Journal of Economics* 109, 685–99.
- Damasio, A. 1995. *Descartes’ Error: Emotion, Reason, and the Human Brain*. New York: Harper Perennial.
- Elster, J. 1989. Social norms and economic theory. *Journal of Economic Perspectives* 3(4), 99–117.

- Fernández, R. 2007a. Women, work, and culture. *Journal of the European Economic Association* (forthcoming).
- Fernández, R. 2007b. Culture as learning: the evolution of female labor force participation over a century. Mimeo, New York University.
- Fernández, R. 2007c. The co-evolution of culture and institutions. Mimeo, New York University.
- Fernández, R. and Fogli, A. 2005. Culture: an empirical investigation of beliefs, work, and fertility. Working Paper No. 11268. Cambridge, MA: NBER.
- Fernández, R. and Fogli, A. 2006. Fertility: the role of culture and family experience. *Journal of the European Economic Association* 4, 552–61.
- Fernández, R., Fogli, A. and Olivetti, C. 2004. Mothers and sons: preference formation and female labor force dynamics. *Quarterly Journal of Economics* 119, 1249–99.
- Giuliano, P. (2007), “On the Determinants of Living Arrangements in Western Europe: Does Cultural Origin Matter?” forthcoming, *Journal of the European Economic Association*.
- Greenwood, J. and Guner, N. 2005. *Social Change*. Economie d’Avant Garde Research Reports No. 9, Economie d’Avant Garde.
- Greenwood, J., A. Seshadri, and M. Yorukoglu (2005), "Engines of Liberation," *Review of Economic Studies* 72(1) 109-133.
- 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, 912–50.
- Greif, A. 2005. *Institutions: Theory and History. Comparative and Historical Institutional Analysis*. Cambridge: Cambridge University Press.
- Guiso, L., Sapienza, P. and Zingales, L. 2003. People’s opium? Religion and Economic Attitudes. *Journal of Monetary Economics* 50, 225–82.
- Guiso, L., Sapienza, P. and Zingales, L. 2004. “The Role of Social Capital in Financial Development,” *American Economic Review* 94 (3), 526-56.
- Guiso, L., P. Sapienza and Zingales, L. 2005. Cultural biases in economic exchange. Working Paper No. 11005. Cambridge, MA: NBER.

- Guiso, L., Sapienza, P. and Zingales, L. 2006. Does culture affect economic outcomes? *Journal of Economic Perspectives* 20(2), 23–48.
- Hanushek, E. and Kimko, D. 2000. Schooling, labor-force quality, and the growth of nations. *American Economic Review* 90, 1184–208.
- Henrich, J., Boyd, R., Bowles, S., Camerer, C., Fehr, E., Gintis, H. and McElreath, R. 2001. In search of homo economicus: behavioral experiments in 15 small-scale societies. *American Economic Review* 91(2), 73–8.
- Ichino, A. and Maggi, G. 2000. Work environment and individual background: explaining regional shirking differentials in a large Italian firm. *Quarterly Journal of Economics* 115, 1057–90.
- Mailath, G. and Postlewaite, A. 2003. The social context of economic decisions. *Journal of the European Economic Association* 1, 354–62.
- Manski, C. 2000. Economic analysis of social interactions. *Journal of Economic Perspectives* 14(3), 115–36.
- Marmot, M. G., Syme, S. L., Kagan, A., Kato, H., Cohen, J. B. and Belsky, J. 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, 514–25.
- Stigler, G. and Becker, G. 1977. De gustibus non est disputandum. *American Economic Review* 67, 76–90.
- Tabellini, G. 2005. Culture and institutions: economic development in the regions of Europe. Working Paper No. 1492, CESifo.