



Cultural Change and Business History

Naomi R. Lamoreaux

This paper draws on the literature of experimental economics to suggest a model of cultural change with applications to business history. The model is based on experiments involving the public goods game, in which players are given an initial endowment of money and told that they can keep it or contribute some or all of it to a common fund. The fund earns interest, and, at the conclusion of the game, the total is divided among all the players, regardless of the magnitude of their contributions. In most settings, players initially contribute a significant fraction of their endowment to the fund, but some choose to free ride on others' investments. If the game is repeated for multiple periods, players observe this free riding and stop putting their money in the fund. If the rules are changed, however, so that free riders can be punished, players will start contributing again and the common fund will grow and provide general benefits. Although this game is typically used to study topics such as tax avoidance and the provision of schools, roads, hospitals, and other similar investments, I argue that cultural practices have many of the features of public goods and that insights from these experiments can be used to explore systematically the dynamics of cultural change.

Cultural Change and Business History

In recent years economists have displayed an increasing interest in culture, so much so that the National Bureau of Economic Research now runs an annual conference on the "Economics of Culture and Institutions."¹ Unfortunately, this new literature has not done much to spur

I have benefited greatly from the comments of Louis Galambos, Taylor Jaworski, Sophia Rosenfeld, Francesca Trivellato, and questioners at the Thursday evening plenary session of the 2017 Business History Conference.

Naomi R. Lamoreaux <naomi.lamoreaux@yale.edu> is the Stanley B. Resor Professor of Economics and History at Yale University and a Research Associate at the National Bureau of Economic Research. She is also Chair of the Department of History at Yale University.

© Business History Conference, 2017. All rights reserved.

URL: http://thebhc.org/sites/default/files/Lamoreaux%20Cultural%20Change%20for%20BEH_BEH%20Formatted%20Draft%20Final.pdf

interdisciplinary conversation, especially with historians. Whereas historians are interested in how cultures change over time, most economists take culture as a given—as a set of norms, values, and beliefs that developed at some point in the distant past and continue to shape behavior, particularly economic behavior, in the present. For economists, culture is useful precisely because they see it as persistent over the long term and hence can treat it as a source of exogenous variation for purposes of causal identification.²

Although the best of these so-called “persistence” studies raise provocative questions about the relationship of the present to the past, they rarely examine the mechanisms that maintain particular norms, values, and beliefs over time.³ Yet cultural practices do not simply persist. They have to be transmitted in some way from one generation to the next, and they have to be reinforced so that people who violate them are disciplined or otherwise persuaded to toe the line. A related problem with this literature is that economists who find statistically significant and economically important effects of past cultural practices on current behaviors are much more likely to have their work published than those that find statistically insignificant or small effects. As a body, therefore, this literature is misleading because it obfuscates the extent to which cultural practices have changed along with (and perhaps as a result of) the behaviors economists are interested in using culture to explain. The persistence of particular cultural practices and not others may in fact be endogenous. If so, these practices cannot be used for the purpose of causal identification.

Of course, some economists have grappled with the problem of understanding cultural change. The models they have proposed, however, are mostly functionalist in character. For example, some scholars have postulated that cultural practices die out when they are no longer efficient, that the values parents choose to pass on to their children are those that are suited to their economic environment, and that cultural traits are selected through a Darwinian struggle for survival.⁴ Historians have generally resisted these approaches as too simplistic and have been

¹ For a recent overview of the economics literature on culture, see Boris Gershman, “Long-Run Development and the New Cultural Economics,” in *Demographic Change and Long-Run Development*, eds. Matteo Cervelatti and Uwe Sunde (Cambridge, Mass., 2017), 221-61.

² According to Luigi Guiso, Paola Sapienza, and Luigi Zingales, for example, culture is “those customary beliefs and values that ethnic, religious, and social groups transmit fairly unchanged from generation to generation.” See “Does Culture Affect Economic Outcomes?” *Journal of Economic Perspectives* 20 (Spring 2006): 23-48, quotation at 23. The authors quote approvingly Gary S. Becker’s assertion, “Because of the difficulty of changing culture and its low depreciation rate, culture is largely a ‘given’ to individuals throughout their lifetimes.” Becker, *Accounting for Tastes* (Cambridge, Mass, 1996), 16.

³ See, for examples, Nathan Nunn and Leonard Wantchekon, “The Slave Trade and the Origins of Mistrust in Africa,” *American Economic Review* 101(December 2011): 3221-52; Nico Voigtländer and Hans-Joachim Voth, “Persecution Perpetuated: The Medieval Origins of Anti-Semitic Violence in Nazi Germany,” *Quarterly Journal of Economics* 127 (August 2012): 1339-92; Alberto Alesina, Paola Giuliano, and Nathan Nunn, “On the Origins of Gender Roles: Women and the Plough,” *Quarterly Journal of Economics* 128 (May 2013): 469-530.

⁴ Examples include Pauline Grosjean, “A History of Violence: The Culture of Honor and Homicide in the US South,” *Journal of the European Economic Association* 12 (October 2014): 1285-1316; Matthias Doepke and Fabrizio Zilibotti, “Occupational Choice and the Spirit of Capitalism,” *Quarterly Journal of Economics* 123 (May 2008): 747-93; Harriet Orcutt Duleep, “The Labor/Land Ratio and India’s Caste System,” College of William and Mary Department of

drawn instead to the work of social theorists, such as Anthony Giddens, who emphasize human agency. In business history, for example, Kenneth Lipartito has argued that cultures change when people confront a world that does not fit their inherited mental constructs and so respond by reshaping both the world and the constructs they deploy to understand it.⁵ Lipartito's call for a non-functionalist approach to studying the relationship between cultural and economic change is appealing, but in the end it offers little more than a general injunction to be open to feedback loops and multiple causes. My aim in this paper is to push further toward a model of cultural dynamics that will not be functionalist in spirit but will have nonetheless have predictive utility, particularly for business historians. The goal of this model is to understand the circumstances under which cultural practices are most likely to change and also when change is most likely to be rapid and transformative.

The inspiration for this paper comes from the literature on experimental economics.⁶ Scholars have run simple economic games all over the world during the past couple of decades. In the process, they have amassed considerable evidence showing that the way in which people play one-shot games reflects cultural norms. They have also shown, using experiments in which people play the same game repeatedly, that the extent to which subjects adhere to cultural norms in the game can weaken or strengthen depending on how the play is structured. In the next section, I describe those results. I then go on to suggest how these very simple games provide insight into the dynamics of cultural change.

Games and Culture

In a one-shot game, subjects receive instruction on how to play, but they do not have any experience to guide them. They have not seen how others might play the game or how different ways of playing might work out, and this lack of familiarity limits their ability to think through the various strategic possibilities. When asked to make choices in the game, therefore, most people fall back on general cultural norms and play the game, in a sense, as they think it *should* be played.

A good example is the ultimatum game. This game has two players who are given the chance to share an allotment of money. The first player is told that she must decide how much of the money to share with the second player. The second player is then asked to decide whether to accept or reject the first player's offer. If he accepts, the two players take home the amounts allocated by the first player. If he rejects, neither player gets anything. Basic economic theory predicts that the first player will keep most of the endowment and offer the other player only a small share. The second player will then accept whatever is offered because anything is better than nothing. However, when experimenters actually tried running the game on student subjects in university computer labs, the results did not conform to these theoretical predictions. Instead, the first players tended to offer the second players shares of forty percent or more and often

Economics Working Paper 137 (March 2013). For an overview, see Gershman, "Long-Run Development and the New Cultural Economics."

⁵ Kenneth Lipartito, "Connecting the Cultural and the Material in Business History," *Enterprise & Society* 14 (December 2013): 686-704. See also Philip Scranton, "Determinism and Indeterminacy in the History of Technology," *Technology and Culture* 36 (April 1995): S31-S53; and JoAnne Yates, "Using Giddens' Structuration Theory to Inform Business History," *Business and Economic History* 26 (Fall 1997): 159-83.

⁶ For a comprehensive introduction to experimental economics, see Colin F. Camerer, *Behavioral Game Theory: Experiments in Strategic Interaction* (Princeton, N.J., 2003).

divided the allotment equally. Offers that were more unequal than a 60/40 split were usually rejected, leaving both players with nothing.⁷ To understand these results, experimenters used medical imaging technologies to observe what happened in the brains of the students when they made their choices. Subjects who rejected unequal distributions used different parts of the brain than those who offered and accepted more equal offers. The former reacted emotionally to an offer they regarded as unfair, rather than rationally calculating their self-interest.⁸

That the students' way of playing the ultimatum game was a cultural response is suggested by experiments involving very different populations.⁹ Joseph Henrich was the first scholar to run simple games in the field. As a graduate student in anthropology at UCLA, he tried out the ultimatum game on the Machiguenga, a hunter-gatherer society in the Peruvian Amazon. The results were stunning. Unlike students in the lab, the Machiguenga played the game the way economic theory predicted. That is, the first players generally made very low offers and the second players almost always accepted them. When Henrich interviewed the players afterwards, he found that they thought about the game very differently from his students at UCLA. "Rather than viewing themselves as being 'screwed' by the proposer, they seemed to feel that it was just bad luck that they were responders, and not proposers." They did not punish the proposers because they thought they would have behaved the same way if the roles were reversed.¹⁰

Henrich hypothesized that cultural differences explained the divergent ways in which the Machiguenga and UCLA students responded to the choices posed by the ultimate game.¹¹ To explore this idea further, he joined forces with an interdisciplinary team of scholars to run simple economic games on a set of fifteen small-scale societies around the world. The results confirmed Henrich's hypothesis that culture mattered. At one extreme were the Lamalera, a whale hunting society in Indonesia, where proposers in the ultimatum game tended to make offers of 50 percent or more. At the other extreme were the Machiguenga and the Quichua, a foraging people in Ecuador, where proposers tended to offer less than 30 percent. Henrich and his collaborators offered a functionalist explanation for their results, hypothesizing that the cultural differences they observed

⁷ Colin Camerer and Richard H. Thaler, "Anomalies: Ultimatums, Dictators and Manners," *Journal of Economic Perspectives* 9 (Spring 1995): 209-19.

⁸ Alan G. Sanfey et al., "The Neural Basis of Economic Decision-Making in the Ultimatum Game," *Science* 300 (13 June 2003): 1755-58.

⁹ Scholars have also examined the ways in which other cultural constructs, such as gender, affect the way these simple games are played. See, for example, Sara J. Solnick and Maurice Schweitzer, "The Influence of Physical Attractiveness and Gender on Ultimatum Game Decisions," *Organizational Behavior and Human Decision Processes* 79 (September 1999): 199-215.

¹⁰ Joseph Henrich, "Does Culture Matter in Economic Behavior? Ultimatum Game Bargaining among the Machiguenga of the Peruvian Amazon," *American Economic Review* 90 (September 2000): 973-79, quotation at 977.

¹¹ Henrich defined culture as "socially transmitted rules about how to behave in certain circumstances (economic or otherwise)." He approached his subject in essentially the same way as the persistence scholars. But his findings also included hints of how cultures changed. A few of the Machiguenga offered to share their allotments more equally with the responders. They were "without exception, those who had had greater exposure and dealings with Westerners and especially North American evangelical missionaries—so they may have acquired some Western notions of fairness from these contacts." Henrich, "Does Culture Matter," 973, 977-78.

had evolved over time to fit the social and economic needs of these societies. Killing whales required the cooperation of the entire community, but foraging was a family-based activity.¹²

If the offers that experimental subjects make when they play the ultimatum game are shaped by culture, however, they are also social acts. By varying the circumstances of the game in ways that obscure the actions of the proposer, it is possible to get results that approach more closely the predictions of economic theory. One experiment allocated a number of chips to the players (students) rather than a sum of money and then varied the value of the chips to each of the players, as well as the information that each player had about the valuations. When both players had full information, the results were similar to the standard game, but when the value of the chips was higher for the first player and only that player knew, the proposer tended to offer the responder an equal number of chips but a lower share of the total valuation. That is, the proposer took advantage of asymmetric information to uphold norms so as to appear fair to the responder but actually divided the allocation in a way that the other player, if s/he had full information, would have regarded as unfair.¹³

These results are all for one-shot games. If the games are repeated, players learn from how others play and modify their choices accordingly. By varying the experimental design and observing how players' behavior changes in response to those changes and to others' play, one can see how the cultural norms exhibited in the one-shot games can be strengthened or alternatively eroded. One game strikes me as particularly useful in this respect: the public goods game. This game is also particularly illuminating for the purposes of this paper because, as I will suggest below, cultural norms have many features of public goods.

The public goods game is played by multiple subjects. Each player is given an endowment of money and told that she can keep it or contribute some or all of it to a public fund. Money put in the public fund generates earnings. After all the players have made their decisions, the fund plus the earnings are divided among all the players regardless of their contributions. Total earnings are maximized when everyone in the game contributes, but the economically rational strategy from the standpoint of the individual is to free ride—not contribute anything but get an equal share of the earnings.¹⁴ If the game is played just once, players tend to respond according to their cultural norms. For example, students in the U.S. and many other western countries tend on

¹² See Joseph Henrich et al., "In Search of Homo Economicus: Behavioral Experiments in 15 Small-Scale Societies," *American Economic Review* 91 (May 2001): 73-78; and Henrich et al., eds., *Foundations of Human Sociality: Economic Experiments and Ethnographic Evidence from Fifteen Small-Scale Societies* (Oxford, 2004).

¹³ John H. Kagel, Chung Kim, and Donald Moser, "Fairness in Ultimatum Games with Asymmetric Information and Asymmetric Payoffs," *Games and Economic Behavior* 13 (March 1996): 110-10. For similar results, see also Rachel T. A. Croson, "Information in Ultimatum Games: An Experimental Study," *Journal of Economic Behavior & Organisation* 30 (August 1996): 197-212. Similarly, Elizabeth Hoffman, Kevin McCabe, and Vernon L. Smith found that proposers in the dictator game (where the respondent had no choice but to accept whatever the proposer offers) kept significantly more for themselves when the game was structured in ways that guaranteed that no one (not even the person running the game) would know what they had offered. See "Social Distance and Other-Regarding Behavior in Dictator Games," *American Economic Review* 86 (June 1996): 653-60.

¹⁴ Henrich et al., eds., *Foundations of Human Sociality*, 12-13; Camerer, *Behavioral Game Theory*, 43-46.

average to contribute a substantial share (around 60-70 percent) of their funds to the public pot. By contrast, students in Saudi Arabia, Greece, and Turkey put in much less on average (about 40 percent).¹⁵ Henrich's collaborators found similar variation among the small-scale societies they studied, with the Machiguenga again contributing relatively little on average.¹⁶ Within all the societies studied, however, individuals varied in their willingness to contribute—that is, in their adherence to general norms—with some players contributing much less than the others and in some instances nothing at all.¹⁷

When the public goods game is played multiple times, the amount contributed to the public pot tends to decline as players figure out that some of their counterparts are free riding on others' contributions. In response to this “unfair” behavior, those who were initially generous in their donations give less, and with repeated play the contributions get smaller and the public fund dwindles essentially to nothing. However, if one then introduces the possibility of punishment in the game—that is, if one allows players to devote a small amount of their resources to inflict a financial penalty (typically larger than the cost of the punishment) on free riders—in most settings public contributions will gradually rise again, sometimes above the levels observed in one-shot games.¹⁸ The ability to punish not only reduces the incentive to free ride but also increases players' sense that the allocation of benefits is fair.

Here again culture matters, however. In places where there were substantially more low contributors in one-shot games, punishment has comparatively little effect. Thus, in one experiment, adding punishment to the game led to increased contributions among players in Boston, Copenhagen, and Zurich, where donations had been high to begin with, but had essentially no effect on players in Athens, Riyadh, and Istanbul, where initial contributions had been low. In intermediate cases the effect of adding punishment varied significantly, with contributions rising in Melbourne, Seoul, and Chengdu, but changing little in Minsk, Samara, and Muscat. When researchers regressed behavior in the game on data from the World Values Survey, they found that other features of the cultures from which the players came, such as attitudes toward evading taxes, fraudulently obtaining government benefits, and ducking transit fares, significantly affected the subjects' responsiveness to the addition of punishment.¹⁹

The Public Goods Game as a Model for Cultural Change

The public goods game is particularly useful for the purposes in this paper because cultural norms have many aspects of public goods. Conforming to cultural norms imposes costs on individuals, who might prefer to behave in another way if left to their own devices, but there are also significant benefits to be reaped by the community if the values and behaviors are widely

¹⁵ Benedikt Herrmann, Christian Thöni, and Simon Gächter, “Antisocial Punishment Across Societies,” *Science* 319 (7 March 2008): 1362-67.

¹⁶ Joseph Henrich and Natalie Smith, “Comparative Experimental Evidence from Machiguenga, Mapuche, Huinca, and American Populations,” in *Foundations of Human Sociality*, eds. Henrich et al, 125-67.

¹⁷ For a clever experiment that aims to measure this heterogeneity, see Erik O. Kimbrough and Alexander Vostroknutov, “Norms Make Preferences Social,” *Journal of the European Economic Association* 14 (June 2016): 608-38.

¹⁸ Ernst Fehr and Simon Gächter, “Cooperation and Punishment in Public Good Experiments,” *American Economic Review* 90 (September 2000): 980-94.

¹⁹ Herrmann, Thöni, and Gächter, “Antisocial Punishment Across Societies.”

shared. As in the public goods game, the benefits accrue to everyone regardless of whether they bear the costs. To give a very simple example, in most economics departments, it is expected that faculty members will attend weekly workshops in their areas of specialization. Attending workshops takes time out of the day and requires advance preparation in the form of reading and thinking about the paper, so faculty members might be tempted either to shirk by not showing up or to free ride on others' efforts by not reading the paper in advance. But if people put in the requisite time and effort, the seminars are better—more intellectually stimulating—and hence more worth attending.

Individuals are likely to differ in the value they obtain from widespread adherence to a cultural norm (for example, from attending a stimulating seminar) and also in their willingness to bear the costs of conforming (attending the seminar and reading the paper in advance), so there are usually at least a few people who behave in some kind of deviant fashion.²⁰ Non-conforming faculty members may skip seminar or just sit there silently without contributing to the discussion. Worse still, they may show up and ask stupid questions because they have not read the paper. Societies usually have some way of disciplining deviants in order to maintain adherence to norms. In the academic example, faculty members who do not contribute to seminars or ask stupid questions may lose the respect of their colleagues and may earn lower salaries as a consequence.

Just as free riding in the public goods game leads to lower contributions, non-adherence to cultural norms undermines their power. In the faculty example, if the quality of seminars falls because free riders ask stupid questions or if people feel that they are unfairly burdened because they attend and prepare while others shirk, non-attendance may increase, preparation decrease, and the commitment to the value of attending seminars decline. But just as the addition of punishment to the public goods game raises contributions, enforcement can also increase adherence to norms, such as those that enhance the quality of seminars. Punishment, however, is more or less effective depending on initial conditions—that is, the extent to which there is compliance with the norm to begin with. Moreover, its effectiveness can also be adversely affected by changes that alter the proportion of non-adherents in the population—for example, hiring new faculty from places that do not share the seminar norm.

The example I have been using to illustrate these points is trivial, but it is easy to think of much more profound situations where similar dynamics might induce changes in cultural practices. For example, cultures may change as a result of the in-migration of peoples whose value systems are different from those of the host population, reducing the proportion of people who adhere to particular norms. Cultures may change when there are shifts in leadership that favor alternative value systems—for example, when one people is conquered by another and the conquerors actively discourage or repress existing norms or undermine the social hierarchies that previously had been responsible for enforcement. They may also change as a result of the development of new information systems that make people more aware of the extent of non-conforming behavior or that help to popularize alternative values or make them seem more attractive.

The operative word in all of these scenarios is “may.” Cultures “may” change, but how can we predict whether such shifts in circumstances make it likely that they will? The experiments with public goods games help us answer such questions by suggesting testable hypotheses that involve a limited number of variables. The first is the extent of adherence to a norm in the

²⁰ On this point, see especially Kimbrough and Vostroknutov, “Norms Make Preferences Social.”

existing population. Not only does this variable affect the amount of free riding—that is, the extent to which individuals can benefit from general adherence to a norm without bearing any of the costs—but it also influences the extent to which enforcement is likely to be effective in improving adherence. To give another simple example, drivers who speed benefit from a general adherence to traffic rules, but the more drivers who speed the less likely the police will be able to keep order on the road.

The second variable is the degree of enforcement. Norms that are not enforced are likely to erode. If police do not enforce the speed limit, for example, because of tight government budgets that reduce the number who can be assigned to staff speed traps, drivers are more likely to speed. But the more norms erode, as just indicated, the less likely that any particular degree of enforcement is going to be able to reverse the trend. One might call this point the slippery slope corollary.

The third variable is the cost of adhering to the norm. An important point to take away from the experimental results is that cost does not matter in and of itself so long as there is widespread adherence. As we have seen, players in public goods games are willing to put the lion's share of their endowments into the general fund. It is only when there are free riders that cost matters. Then corrosive worries about unfairness kick in, and people become increasingly unwilling to bear costs for the general good. At this point, the magnitude of the costs becomes important. If you are late to work, you are more tempted to speed, but you are more likely to give way to this temptation when lots of other drivers speed. All things being equal, one would expect that the higher the cost of adhering to a norm, the greater the negative effect of free riding on conformity to the norm.

So far, I have set up this discussion so the erosion of norms is treated as an undesirable development. But it is possible to flip the analysis and think about how to change cultural practices that we would take to be bad—for example, racial discrimination. Such practices will be more entrenched to the extent that they are widely shared and systematically enforced, even if they are costly. For example, discrimination in hiring imposes costs on employers because they could lower their wage bill by increasing the size of the labor pool. But, as Gavin Wright has shown for the American South, if discrimination is socially enforced, employers cannot reap the gains from opening jobs to black workers. For example, when an Atlanta textile factory attempted to hire twenty black women for its folding department, it faced protests, strikes, and property damage until it gave up and discharged the women.²¹ In the 1960s, however, the Civil Rights Movement altered the enforcement environment. New federal rules that made government contracts contingent on non-discriminatory hiring practices resulted in a dramatic increase in black employment in Southern textile mills, which sold a large fraction of their output to the federal government at that time.²² Similarly, blacks' readiness to file antidiscrimination suits helped to put teeth in the Civil Rights Act of 1964.²³

²¹ Gavin Wright, *Old South, New South: Revolutions in the Southern Economy since the Civil War* (New York, 1986), 189-90.

²² James J. Heckman and Brook S. Payner, "Determining the Impact of Federal Antidiscrimination Policy on the Economic Status of Blacks: A Study of South Carolina," *American Economic Review* 79 (March 1989): 138-77. See also John J. Donohue III and James Heckman, "Continuous Versus Episodic Change: The Impact of Civil Rights Policy on the Economic Status of Blacks," *Journal of Economic Literature* 39 (December 1991): 1603-43.

²³ Gavin Wright, *Sharing the Prize: The Economics of the Civil Rights Revolution in the American South* (Cambridge, Mass., 2013), Ch. 4.

Business History and the Study of Cultural Change

I have argued in this paper that cultural norms have many of the attributes of public goods and that changes in their strength are driven by essentially the same set of variables as contributions in public goods games: the relative number of adherents to a norm, the extent of enforcement, the costs of conforming, and the interactions among these three variables. It should be noted that there is nothing inherently functionalist about this approach. Values may change because those at top of the social hierarchy find it in their interest to stop (or start) enforcing them or because individual members of a society increasingly find it in their interest to stop (or start) conforming to particular norms. But values may also change for other reasons—for example, because there is in- or out-migration or because enforcers are distracted by other concerns—and there is no reason to believe that the resulting changes are in any way efficient in the sense that they are a better fit to the economic or social environment. Nonetheless, I would argue that this non-functionalist approach has a great deal of predictive utility for highlighting the kinds of situations that are likely to lead to cultural change. It also has value for us as historians in sharpening our understanding of the sources of cultural change in the past.

Finally, I would suggest by way of conclusion that business history is particularly well suited for exploring the explanatory power of these hypotheses about cultural change. Business enterprises have their own “corporate cultures” and also frequently find themselves in situations where those cultures are subject to challenge. Businesses may merge with other enterprises that have different corporate cultures and may therefore face the imperative of imposing their value systems on their acquisitions or changing their own cultures in turn. Business leaders may seek to change their enterprise’s culture from above in order to make it more competitive, perhaps by imitating a rival. Businesses may find their cultures challenged when they expand their operations and take on new employees, particularly when they expand abroad. Because these challenges occur in relatively compressed timeframes, they can be readily be studied from start to finish. Business history thus promises to offer the study of cultural change what fruit flies provide for the study of biological evolution.