Female genital cutting without a cutting norm

Charles Efferson,[1*] Sonja Vogt,[1*] Amy Elhadi,[2]
Hilal El Fadil Ahmed,[2] and Ernst Fehr[1*]
[1]Department of Economics, University of Zürich,
Blümlisalpstrasse 10, 8006 Zürich, Switzerland
[2]Khartoum, Sudan

*To whom correspondence should be addressed;
E-mail: charles.efferson@econ.uzh.ch, sonja.vogt@econ.uzh.ch, ernst.fehr@econ.uzh.ch;
Efferson and Vogt contributed equally to this project.

The World Health Organization defines female genital cutting as any procedure that removes or otherwise injures any part of a female's external genitalia for non-medical reasons (*1*). Females are typically cut when they are young girls. Although different forms of cutting exist, the most aggressive forms involve extensive tissue removal followed by a procedure known as infibulation. To infibulate a female, the wounds produced by cutting are sewn up so that they heal to leave only a small opening for urine and menstrual blood. Female genital cutting brings no documented health benefits, and it can lead to serious health problems throughout life. These health problems include severe pain, bleeding, infection, psychological trauma, infertility, and a variety of obstetric complications that in the worst cases lead to death. In Sudan, as but one example, obstetric complications related to cutting account for approximately 1% of all government health spending on women between the ages of 15 and 49 (*2*). More broadly, across six African countries, a single cohort of 15-year-old girls is expected to lose altogether nearly 130,000 years of life because of problems during pregnancy and childbirth due to genital cutting (*2*).

The prevailing view of female genital cutting is that the practice is a social norm that evolved culturally within the context of a coordination game (*3–11*). Development agencies spend millions of dollars every year for programs based on this model (*10, 12–14*). Although the details vary considerably, these programs share the assumption that coordination incentives play a key role in perpetuating female genital cutting, and they also play a key role in leading people to abandon the practice. Moreover, the social norms approach and its emphasis on coordination have spread from an initial focus on female genital cutting to other domains where development organizations would like to promote changes in behavior (*7, 9, 10, 14, 15*).

The coordination game approach to female genital cutting can be illustrated with the following simple model related to families and the marriage of their children (*3, 6, 8, 16*). Assume that families have sons and daughters. Some families cut their daughters and demand cut wives for their sons. We will call this the "cutting" strategy. Other families do not cut their daughters, and they demand uncut wives for their sons. This is the "non-cutting" strategy. If a family cuts, family members pay a total cost of $c > 0$. This cost includes the daughter's pain, psychological trauma, and risk of infection, infertility, or death. It also includes any medical expenses associated with treating the daughter's health problems.

As sons and daughters reach a marriageable age, families would like them to find suitable partners. Families are randomly matched for their sons and daughters to consider potential marriages. A successful marriage brings a total benefit of $b > c$ to the family. This benefit includes the prestige from successful reproductive matches in natalist societies. It further includes advantages like increased social influence and risk spreading that result when families use marriage to create or strengthen alliances. If a marriage is not successful or does not happen, the families involved receive no benefit. For a match to result in a successful marriage, the two families must have the same strategy. This is the key to coordination. A successful marriage requires that matched families have the same practices and demands related to cutting, regardless of whether these practices and demands favor cutting or not cutting.

Under this simple model, two cutting families that are matched each get a payoff of $b - c$. Two non-cutting families that are matched each get a payoff of b. If two families with different practices are matched, the cutting family gets $-c$, and the non-cutting family gets 0. Finally, to generate dynamics, families update their strategies over time, and at some positive rate they adopt the strategy with the highest expected payoff. This model has three equilibria (Fig. 1). Namely, all families cut, no families cut, or a specific intermediate proportion of families cut. In the latter case, the equilibrium proportion of families cutting is $(b + c)/(2b)$. Of the three equilibria, only the two pure equilibria are dynamically stable; the mixed equilibrium is unstable (Fig. 1). Consequently, within a group of intermarrying families, the model predicts that either all families cut or none of them do. Within a group, behavior and the demands that families bring to the marriage pool are homogeneous in equilibrium, and for this reason they constitute a social norm that efficiently facilitates coordination. Once a norm is in place, no family can afford to deviate unilaterally, and so the norm can potentially persist for a very long time. This is true even if the norm in question is cutting, which yields a lower payoff for all families than the non-cutting norm. To see this, consider an individual family facing a group in which all other families cut. If the family in question does not cut its daughters, the payoff will be 0. If it does cut, the payoff will be $b - c > 0$. Importantly, the theoretical prediction of homogeneity within groups holds under assumptions much more general and realistic (*17*) than those presented in this simple model.

The coordination game logic has been extremely influential both among researchers with an interest in the cultural evolution of social norms (*3, 6, 18–20*) and among practitioners with an applied interest in cutting as a health risk and a violation of children's rights (*4–8, 10, 11, 14*). In particular, for several years the social norms approach and its emphasis on coordination incentives have been a dominant theoretical framework for programs promoting the abandonment of female genital cutting (*5, 7–10, 14*). The key idea is that no single family can afford to deviate from the cutting norm because the chance would be too high that the children of the family do not grow up to find marriage partners (*5–7, 14*). However, if a sufficient number of families deviate in a multilateral fashion, this problem disappears. More formally, if a program promoting abandonment can convince enough families to stop cutting simultaneously, this multilateral deviation will shift the population into the basin of attraction for the non-cutting equilibrium (*5, 7*). Provided this deviation is somehow public (*9, 10, 14*), the families who still cut will update their beliefs about the distribution of cutting practices in the group. These families will then realize they would be better off if they also stopped cutting. Cutting quickly disappears at this point as endogenous social dynamics ensure that finding marriage partners from non-cutting families is increasingly easy to do.

Programs based on this model can take various forms, but they have certain common features. Development workers usually invest in local projects at the community level. Along with other development activities, these workers gradually try to convince families to abandon female genital cutting and publicly declare that they have done so. Once the development workers estimate that they have assembled a sufficient number of families in a community, they have a public declaration (*8–10, 14*). This is typically a large celebratory festival in which the families in question publicly declare that they will no longer cut their daughters. The hope is that this public declaration will lead the remaining families who cut to update their beliefs and thus realize that abandoning the practice is now in their own interests. If this happens, cutting should disappear relatively quickly because social interactions among families now work in favor of abandonment.

Surprisingly, given the development funds at stake, no one has ever provided data clearly showing that female genital cutting exhibits the characteristics of a social norm based on coordination. In particular, if coordination is an appropriate theoretical framework, two related predictions hold. First, groups should have extreme cutting rates near zero or one. Second, if cutting practices vary, the variation should be primarily between groups, not within groups. These predictions follow directly from the model outlined above (Fig. 1) in which any mix of strategies within a group is unstable, and thus a group only stabilizes culturally if everyone cuts or if no one cuts. By extension, if behavioral variation exists, it must be between groups, not within groups.

To see if such a pattern can be detected for female genital cutting, we collected data with approximately 2300 girls entering primary school for the first time in 45 different communities in the state of Gezira, Sudan. Children enter primary school in Sudan at the age of six, and school begins in Gezira in late June. In addition, the Gezira state government hires doctors at the beginning of each school year to give basic medical exams to all children entering primary school. To develop a novel and culturally appropriate method for reliably estimating cutting rates by community, we added two simple steps to this medical exam.

Novel methods are necessary (*21*) because existing methods rely heavily on surveys that simply ask parents about cutting (*22*). The risk with such methods is that adult respondents know that governmental and non-governmental workers asking about cutting probably have a social and political agenda of some sort. As a result, respondents might misrepresent their true practices and attitudes in various ways and for various reasons. The net effect of potential biases is not known, and so we have little sense of how reliable existing survey methods are.

We took a different approach. In late June and early July of 2013, nine teams of professional photographers and medical doctors visited 45 schools in the state of Gezira, one school in each of the 45 communities in our study. Importantly, girls in Gezira are typically cut during the summer vacation just before entering primary school or earlier. Henna is almost always applied to a girl's feet when she is cut, and this is the only time in a young girl's life when henna is applied. The presence of henna is thus an important and extremely reliable signal that a young girl has been recently cut, and after being applied residual henna remains on the toenails for several weeks. For this reason, the photographers in our study photographed the feet of nearly all girls entering primary school for the first time in the 45 schools (*17*). In addition to the photographs, the medical doctors conducted the basic medical exams that are normally sponsored by the Gezira state government. We added one question to this exam. Namely, as part of several health-related questions, the doctors asked each girl if she had been "purified" and, if so, when. The word "purified" is a literal translation of the Arabic word commonly used in the local area to describe being cut. We combined data from both the photographs and the doctors' question to code each girl as cut or not and, by extension, to estimate cutting rates for each community (*17*). This, in turn, allowed us to test if cutting rates for the communities in our study exhibited the kind of variation we would expect if cutting practices are coordinated social norms.

Although communities varied substantially in terms of their estimated cutting rates, they did not vary in any way that reveals the signature of social norms that evolved within the context of a coordination game. In particular, apart from a single community with an estimated cutting rate of 100%, all estimated cutting rates are squarely in the interior of the unit interval (Fig. 2). Moreover, when ordered, estimated cutting rates show no trace of the discontinuity they should exhibit if the families in our study were coordinating on cutting practices at the community level (Fig. 2).

To analyze cutting rates formally, we fit three different models to the data set (*17*). The models differ in terms of the assumptions they make about the structure of heterogeneity among communities. The **homogeneous** model assumes that all communities are identical, and thus the variation in cutting rates across communities is simply sampling variation. The **coordination** model assumes that two types of communities exist. Communities that have coordinated on a non-cutting norm have a relatively low cutting rate, while communities that have coordinated on a cutting norm have a relatively high cutting rate. Finally, the **heterogeneous** model assumes that all communities are different, they vary continuously with respect to cutting behavior, and they cannot be sensibly grouped by cutting rate. We fit the three models to our data using maximum likelihood. We then used an information theoretic criterion (AIC_c) to select the best model (*23*). The model selection results are overwhelmingly clear (Table 1). The best model, by an enormous margin, is the heterogeneous model.

Although our observed cutting rates show no trace of coordination at the community level, an important caveat applies. Coordination and associated social norms could have been present in marriage pools that were not the same as the geographically defined communities in our study. If marriage pools subsume multiple communities, this is not an issue for our analysis. A marriage pool that evolves toward a specific norm will take all the communities it includes along with it. In principle, however, a community could also consist of two or more marriage pools, some that coordinate on cutting and others that coordinate on not cutting. If this was the case in the majority of the communities in our sample, the prevalence of interior cutting rates in our data might still conceivably be consistent with the coordination model.

To address this possibility, we surveyed 2458 adults, with a roughly equal mix of men and women, from 2458 randomly sampled households in the 45 communities in our study. Specifically, we asked each participant if her family would refuse a marriage with another family in her community because of the other family's cutting practices. Every community had a strictly interior proportion of subjects answering yes (*17*). This means that every community had families who apparently insisted on similar cutting practices, and every community also had families who did not insist.

To understand the implications of this observation with respect to a coordination game model, we have to consider strategies defined in two dimensions. In the first dimension, families either cut or not. In the second dimension, when negotiating marriages, families either insist on coordinating with respect to cutting behavior or not. This yields a total of four possible strategies.

We developed an expanded coordination game model in which strategies are defined in this way (*17*). The model is similar to the model presented at the beginning of the paper, but it additionally allows families to vary in terms of whether they insist on coordinated cutting practices. It also allows families to assort themselves into distinct marriage pools, one pool that cuts and another that does not. The key conclusion from the model is that a dynamically stable mix of cutting and non-cutting families at the community level does not exist under any conditions (*17*). As a consequence, the dynamical logic of a simple coordination game (Fig. 1) is maintained in a more complicated, more general, and more realistic game with two-dimensional strategies, multiple marriage pools with divergent cutting practices, and assortment within marriage pools. In this expanded coordination game model, the prediction is still that communities should have cutting rates either close to zero or close to one. Strictly interior cutting rates without a discontinuity, as observed in our data, are not compatible with the model. By extension, our data do not support the view that female genital cutting in Sudan is a social norm that evolved within the context of a coordination game.

To further substantiate this conclusion, we conducted another large-scale empirical study based on an application of the implicit association test (*24*) to female genital cutting. Implicit association tests measure associations between target stimuli, which are presented in neutral terms, and valued stimuli, which are necessarily value-laden. In our case, the target stimuli were drawings of girls who were identified as either cut or not cut, while the valued stimuli were audio recordings of positive words and negative words (*17*). Thus, our implicit association test was designed to measure if an individual associates a cut girl with positive words and an uncut girl with negative words or vice versa. Implicit association tests have been repeatedly validated and widely used to study associations related to race, sexual orientation, religion, and other sensitive topics in contemporary societies (*24*). Implicit association tests are much less prone to producing socially desirable responses than traditional survey methods. Therefore, after extensive development and pre-testing, we implemented our implicit association test with 2260 adult participants from 2260 randomly sampled households in the 45 communities where we estimated cutting rates (*17*). As explained below, we used the resulting data to check for coordination with a simple analysis that requires no information whatsoever about the relationship between communities and marriage pools.

Scoring the implicit association test produces a single measure, D, for each participant. D can take any value in the interval $[-2, 2]$. In our case, $D \in [-2, 0)$ indicates negative implicit associations with uncut girls. In contrast, $D \in (0, 2]$ indicates positive implicit associations with uncut girls. Finally, $D = 0$ means the participant has no implicit associations.

Our empirical approach is based on the following reasoning. If cutting practices constitute social norms that evolved within the context of a coordination game, these norms should affect people's attitudes about cutting (*25*). Simply put, people in marriage pools with a cutting norm should have positive attitudes about cutting, and people in marriage pools with a non-cutting norm should have positive attitudes about not cutting. Because our implicit association test is a measure of attitudes about cutting, a person's score from the test should tend to reflect the prevailing norm in that person's marriage pool. More specifically, people from marriage pools with a cutting norm should tend to produce negative values of D, while people from marriage pools with a non-cutting norm should tend to produce positive values of D. Moreover, if all families are part of a marriage pool with one norm or the other, most people should have either positive or negative associations with cutting, and thus values of D near zero should be comparatively rare. By extension, if two disparate normative systems coexisted when we sampled our study population, we should observe a bimodal distribution of D values with one negative mode for individuals from cutting marriage pools and one positive mode for individuals from non-cutting marriage pools. Critically, we do not need to know the relationship between communities and marriage pools to test this prediction. Because D is an individual measure, and because the sign of D has real meaning, we only need to know how D values are distributed over the participants who took our test.

Importantly, the average value of D by community correlates negatively with our estimated cutting rates by community (Spearman's rank correlation, $\rho = -0.493$, $p < 0.001$). This means that communities with relatively high estimated cutting rates had, on average, relatively positive associations with cut girls. Communities with relatively low estimated cutting rates, in contrast, tended to have relatively positive associations with uncut girls. This correlation shows that our implicit association test measured attitudes that both varied across individuals and covaried in the expected way with observed cutting rates at the community level.

The question then, as indicated above, is whether the distribution of D values is bimodal. Far from being bimodal, the distribution of D values is unequivocally unimodal (Fig. 3). This same result holds for all communities if the D distributions are tested separately for each community (*17*). Consequently, either across all communities or within any single community, our implicit association test provides no evidence for two distinct normative systems. Instead, the data lead to the same conclusion as our analysis of variation in cutting rates by community (Table 1 and Fig. 2). Our participants varied, in short, with respect to both their attitudes and their associated cutting practices. They did not vary, however, in a way that produced any trace of a discontinuity that should have been present if two distinct normative systems had evolved within the context of a coordination game.

How can all of our results be so inconsistent with the coordination view of female genital cutting? To address this question, we would like to be very precise about what we have shown, what we have not shown, and what it means. Emphatically, we have *not* shown that coordination with respect to marriageability is irrelevant for all families in our study. Indeed, as discussed above, all communities have at least some families who claim that coordination matters.

Just as emphatically, however, we have shown that a simple story in which cutting families are collectively trapped at a cutting equilibrium of a coordination game cannot be correct. As explained, this is an extremely influential model of cutting, and it makes clear predictions. A group of families goes one way or another. They either cut together, or they have nothing to do with the practice. Every way we look at the problem, our data are inconsistent with this prediction. Indeed, our data suggest that cutting families routinely live door-to-door with non-cutting families.

In principle, the design of development programs and policy interventions should be based on a sound understanding of the relevant decision making forces. The social norms approach to female genital cutting and to behavioral change in other domains is perhaps best characterized by the following key assumptions. People who engage in harmful practices do so because their own social interactions generate a harmful equilibrium, and they are stuck in this equilibrium even to the extent that they are ignorant of alternatives (*6, 11*). By extension, the task of the development agency is to engineer a collective deviation from the behavioral status quo (*5, 7, 9, 10, 14*). If this deviation is sufficiently large and public, individual incentives to coordinate with others will begin to support behavioral change quite apart from the activities of the development agency itself. Thus, coordination means the development agency can use endogenous social forces to promote changes in behavior that are ultimately self-sustaining. This idea is promising in the sense that it suggests how a relatively short-term intervention can lead to long-term improvements in human well-being. If the underlying determinant of behavior is not coordination, however, we have no particular reason to think that such programs will work.

For example, other forces can easily dominate the need to coordinate. If some economic or cultural force unrelated to coordination leads people to place a strong intrinsic value on cutting, people will continue to cut even if they believe a large proportion of families have abandoned the practice. In such a case, coordination-based programs should not be the focus of development interventions. Perhaps people cut because they have come to believe their religion demands it (*26*). Perhaps in some societies, some men have sufficient power in the mating market to demand costly signals of sexual fidelity, and some women have incentives to provide these signals (*26*). Even more generally, families might cut because in some cultures gender is very clearly marked, and people feel that cutting produces feminine women (*26*). These are just three examples of alternate mechanisms, but they illustrate the risk for a norms-based intervention. Specifically, they illustrate the risk that coordination could be largely beside the point for development efforts. Coordination and related social mechanisms do affect behavior in many situations. In a given development context, however, this does not mean that endogenous social forces will suddenly switch from a state in which they obstruct development objectives to a state in which they promote development objectives.

In the final analysis, rigorous empirical research is necessary to understand why harmful practices like cutting persist and, by extension, how to design effective programs that promote endogenous and sustainable behavioral change. Our study is the first large-scale attempt to see if female genital cutting shows the characteristic signature of a social norm based on coordination. Although the coordination model forms the cornerstone for a number of development programs related to cutting, our results suggest that development agencies cannot safely assume that incentives to coordinate play an important role in a family's final decision to cut.

References and Notes

1. World Health Organization, Female genital mutilation, `http://www.who.int/topics/female\_genital\_mutilation/en/` (2014, accessed 17 November 2014).

2. D. Bishai, *et al.*, *Bulletin of the World Health Organization* **88**, 281 (2010).

3. G. Mackie, *American Sociological Review* **61**, 999 (1996).

4. G. Mackie, *Female "Circumcision" in Africa: Culture, Controversy, and Change*, B. Shell-Duncan, Y. Hernlund, eds. (Boulder, CO: Lynne Rienner Publishers, 2000), pp. 253–283.

5. UNICEF, Coordinated Strategy to Abandon Female Genital Mutilation/Cutting in One Generation: A Human-Rights Based Approach to Programming, `http://www.childinfo.org/files/fgmc_Coordinated_Strategy_to_Abandon_FGMC__in_One_Generation_eng.pdf` (2007, accessed 17 November 2014). The Introduction, Chapters 3 and 4, and Annex II detail how the social norms approach and its emphasis on coordination inform programming efforts related to female genital cutting.

6. G. Mackie, J. LeJeune, Social dynamics of abandonment of harmful practices: a new look at the theory, *Innocenti Working Paper No. 2009-06*, UNICEF Innocenti Research Centre (2009).

7. UNICEF Innocenti Research Centre, The dynamics of social change: Towards the abandonment of female genital mutilation/cutting in five african countries, `http://www.unicef-irc.org/publications/618` (2010, accessed 17 November 2014). For a clear formulation of the coordination game logic and its programming implications for female genital cutting and other harmful practices, see Chapters 3, 4, and 5.

8. C. W. Dugger, Senegal curbs a bloody rite for girls and women, `http://www.nytimes.com/2011/10/16/world/africa/movement-to-end-genital-cutting-spreads-in-senegal.html` (2011, accessed 18 November 2014).

9. TOSTAN, Community Empowerment Program, http://www.tostan.org/community-empowerment-program (2013, accessed 17 November 2014).

10. UNFPA-UNICEF, Joint evaluation of the UNFPA-UNICEF joint programme on female genital mutilation/cutting: Accelerating change, http://www.unfpa.org/public/home/about/Evaluation/EBIER/TE/pid/10103 (2013, accessed 17 November 2014). For the key role of community-level programming based on social norms and coordination, with respect to both female genital cutting and other harmful practices, see sections 1.2, 3.1.4, 3.2.3, 3.4, 4.1, and 4.2.4.

11. UNICEF, Female genital mutilation/cutting: A statistical overview and exploration of the dynamics of change, http://www.unicef.org/publications/index_69875.html (2013, accessed 17 November 2014). For the central importance of coordination, social norms, ignorance of alternatives to cutting, and their relation to programs promoting the abandonment of female genital cutting, see Chapters 2 and 3.

12. L. Featherstone, UK to help end female genital mutilation, https://www.gov.uk/government/news/uk-to-help-end-female-genital-mutilation (2013, accessed 17 November 2014).

13. TOSTAN, Annual Reports, 2008–2012, http://www.tostan.org/annual-reports-financials (2013, accessed 17 November 2014).

14. United Kingdom Department for International Development, Towards Ending Female Genital Mutilation/Cutting in Africa and Beyond, devtracker.dfid.gov.uk/projects/GB-1-203024/documents (2013, accessed 17 November 2014). For a detailed account of the U.K.'s focus on community-level development activities based on social norms and coordination, see pp. 5, 9-10, 13, 17-19, 21-22, 24, and 27.

15. UNICEF, Child protection from violence, exploitation, and abuse, http://www.unicef.org/protection/57929_57990.html (2011, accessed 17 November 2014). Explains the general use of the social norms approach for programs related to child protection.

16. B. Shell-Duncan, K. Wander, Y. Hernlund, A. Moreau, *Social Science & Medicine* **73**, 1275 (2011).

17. Materials and methods are available as supporting material on *Science* Online.

18. H. P. Young, *The Journal of Economic Perspectives* **10**, 105 (1996).

19. R. McElreath, R. Boyd, P. J. Richerson, *Current Anthropology* **44**, 122 (2003).

20. S. Bowles, *Microeconomics: Behavior, Institutions, and Evolution* (New York: Russell Sage, 2004).

21. S. Elmusharaf, N. Elhadi, L. Almroth, *BMJ* **333**, 124 (2006).

22. Measure DHS/ICF International, Demographic and health surveys methodology, http://dhsprogram.com/pubs/pdf/DHSQM/DHS6_Module_Female_Genital_Cutting_19Aug2013_DHSQM.pdf (2013, accessed 16 June 2014).

23. K. P. Burnham, D. R. Anderson, *Model Selection and Multimodel Inference: A Practical Information-Theoretic Approach* (New York: Springer-Verlag, 2002), second edn.

24. B. A. Nosek, A. G. Greenwald, M. R. Banaji, *Automatic Processes in Social Thinking and Behavior*, J. A. Bargh, ed. (Psychology Press, 2007), pp. 265–292.

25. S. Bowles, *Journal of Economic Literature* **36**, 75 (1998).

26. E. Gruenbaum, *The Female Circumcision Controversy: An Anthropological Perspective* (Philadelphia: University of Pennsylvania Press, 2001).

27. J. A. Hartigan, P. M. Hartigan, *The Annals of Statistics* **13**, 70 (1985).

Table 1: Model selection results. We fit three models to our observed cutting rates by community. The homogenous model assumes that all communities are identical, and observed variation is due to sampling. The coordination model assumes that two types of communities exist, cutting communities and non-cutting communities. The heterogenous model assumes that communities differ, but they cannot be sensibly grouped in terms of cutting rates. A derivative form of Akaike's information criterion (AIC_c) is our model selection criterion (*23*). The value of AIC_c for a given model depends on the number of observations, and it also involves a penalty that increases in the number of estimated parameters. Model fitting is based on 2265 observations. Akaike weights (w_i) rescale AIC_c values. They sum to one and represent the proportional weight of evidence in favor of each model (*23*). Here they show that the heterogeneous model is by far the best.

Model	Parameters	Observations	AIC_c	w_i
Homogeneous	1	2265	433.615	2.645×10^{-36}
Coordination	3	2265	320.126	1.165×10^{-11}
Heterogeneous	44	2265	269.774	> 0.999

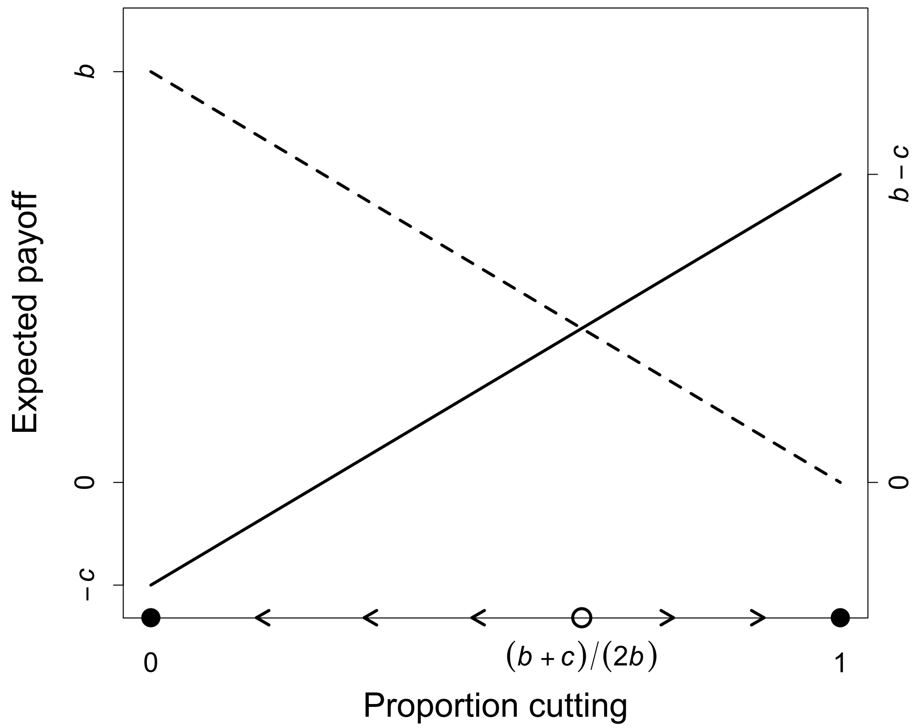

Figure 1: The dynamics of female genital cutting as a social norm based on coordination incentives. The solid line shows the expected payoff of cutting as a function of the proportion of families cutting. The dashed line shows the analogous expected payoff associated with not cutting. The value b is the payoff from a successful marriage, and c is the cost of cutting. Generic dynamics (arrows) follow from any updating process that tends to increase the frequency of the strategy with the highest expected payoff. The model has three steady states. Stable steady states (solid circles) are at 0 and 1. An unstable steady state (open circle) is at $(b + c)/(2b)$, which is the distribution of strategies such that cutting and not cutting have the same expected value. The arrows show that, if the proportion of families cutting is greater than $(b + c)/(2b)$, the proportion of families cutting increases through time to 1. Analogously, if the proportion of families cutting is less than $(b + c)/(2b)$, the proportion of families cutting decreases through time to 0.

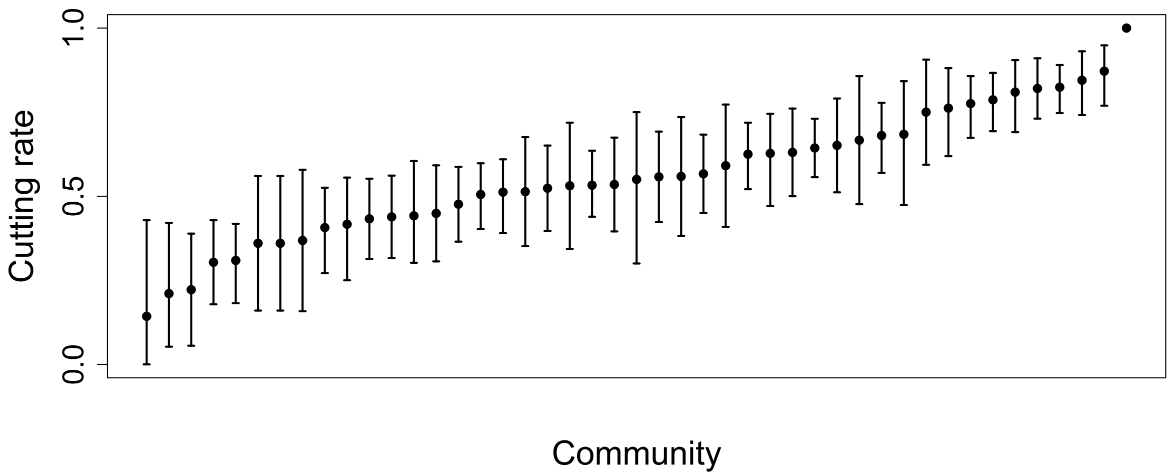

Figure 2: Cutting rates for 45 communities, ordered from lowest to highest, with 95% bootstrapped confidence intervals. The coordination game model (*3, 17*) predicts that cutting rates, when ordered, exhibit a discontinuity separating non-cutting communities from cutting communities. No such discontinuity is present, and a model that assumes continuous variation among communities fits the data far better than a model based on coordination (Table 1).

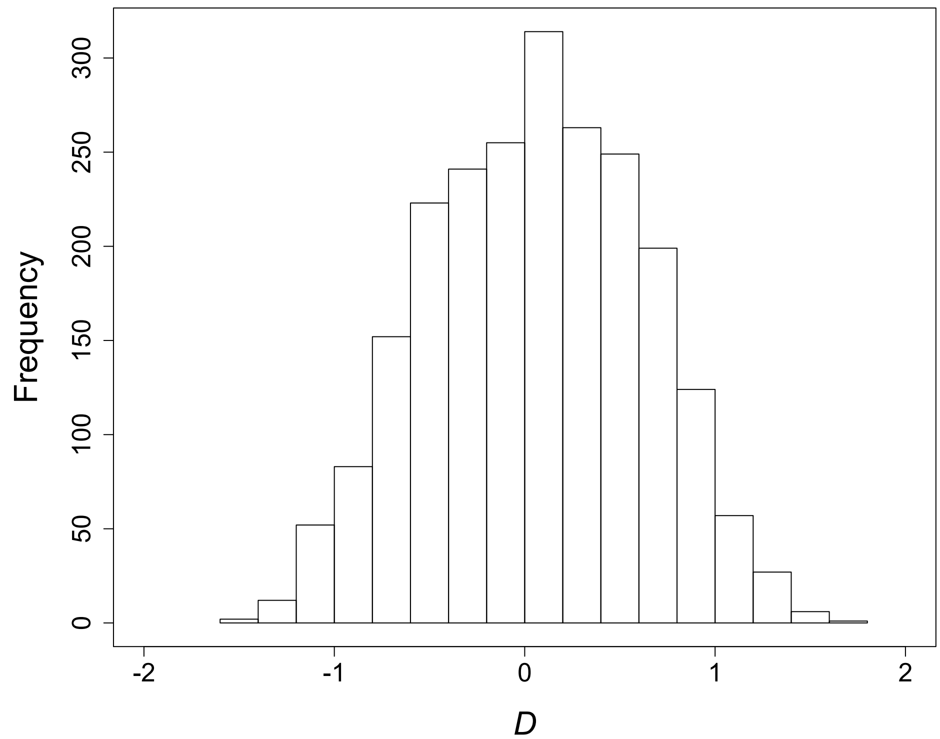

Figure 3: Distribution of scores from the implicit association test. D is a summary score for each participant. A score in $[-2, 0)$ indicates that the participant has positive associations with cut girls and negative associations with uncut girls, while a score in $(0, 2]$ shows the opposite set of associations. If two norms coexisted at the time of our study, the distribution of D scores should be bimodal. It should have one negative mode for people from marriage pools with a cutting norm and one positive mode for people from marriage pools with a non-cutting norm. Using Hartigans' dip test of unimodality (*27*), the dip statistic is $D_n = 0.0035$, and $p = 1$. Consequently, we cannot reject the null hypothesis that our observed distribution of D values has only one mode. The mean value of D is 0.058 (Wilcoxon signed rank test, $p < 0.001$). Data are from 2260 adults, with a nearly equal mix of men and women, from 2260 randomly sampled households in 45 communities.