Development as a Process of Social Change: An Application to the Fertility Transition *

Kaivan Munshi[†] Jacques Myaux[‡]

December 2002

Abstract

This paper provides an explanation for two features of the development process that have been observed in many different settings: the slow response to external interventions, and the wide variation in the response to the same intervention. One interpretation of these stylized facts is based on the idea that individual behavior is often socially regulated in a traditional economy. When the economic environment changes, individuals may be uncertain about the new social equilibrium that will emerge in their community. This uncertainty is resolved over time, as individuals interact sequentially with each other, and the community gradually finds its way to a new equilibrium. We apply this view of the development process to the fertility transition in rural Bangladesh, during which contraceptive prevalence changes slowly and there is wide variation in long-run levels of contraception across villages. At the individual level, estimated changes in the contraception decision rule over time match well with our learning-based characterization of social change. Further, women respond strongly to contraceptive prevalence within their own religious group within the village, while cross-religion effects are completely absent, consistent with the view that social interactions among the women, which almost never cross religious boundaries, gradually led to changes in reproductive behavior in these communities.

*We thank Dan Ackerberg, Orazio Attanasio, Abhijit Banerjee, Sam Bowles, Joytsnamoy Chakraborty, Esther Duflo, Stephen Donald, Jan Eeckhout, Andrew Foster, Jeroen van Ginniken, Michael Kremer, Kevin Lang, Nancy Luke, Dilip Mookherjee, Jack Porter, Andrew Postlewaite, Mark Rosenzweig, Frank Schorfheide, T. Paul Schultz, Frank Vella, Peyton Young, numerous seminar participants and two anonymous referees for helpful comments. We are especially grateful to George Mailath for many helpful discussions that substantially improved the paper. Mohammad Ali expertly assembled much of the data, Aldo Colussi, Nauman Ilias, Poorti Marino and Chun-Seng Yip provided superb research assistance. Munshi's research was funded by NIH grant R03-HD35516. We are responsible for any errors that may remain.

†University of Pennsylvania and MIT

‡International Centre for Diarrhoeal Disease Research, Bangladesh

1 Introduction

This paper provides an explanation for two features of the development process that have been observed in many different settings: the slow response to external interventions, and the wide variation in the response to the same intervention. Consider, for example, the fertility transition, which will be the focus of this paper. Although fertility rates have declined rapidly throughout the developing world over the past few decades, this trend masks enormous variation in the onset of the fertility transition, as well as the subsequent speed of this transition. Countries at similar levels of economic development are often seen to display very different patterns of fertility behavior (Bongaarts and Watkins 1996). Long delays and wide differentials in the response to family planning programs have also been frequently observed, both across countries as well as within countries (see, for instance, Bulatao 1998, Cleland et al. 1994, NRC 1993).

One interpretation of the patterns described above is based on the idea that many aspects of individual behavior are socially regulated in a traditional economy. While such social regulation has advantages of its own, the drawback is that it may prevent individuals from responding immediately to new economic opportunities. Social norms are typically seen to emerge in environments characterized by multiple equilibria, to keep the community in a preferred equilibrium (Kandori 1992). Changes in the economic environment could reopen the possibility for such multiple equilibria, which would explain the slow response to external interventions, as well as the differential response to the same external stimulus, as each community gradually converges to a new equilibrium.

Our view of development as a process of social change is not new. As Ray (1998, p.323) remarks, also in the context of the fertility transition, "The very strength of such [traditional] norms becomes a weakness when the environment of the society begins to change. Accepted, appropriate practice over many centuries may now become inappropriate, but once this happens, social practice is often slow to alter. It becomes necessary to coordinate on some new norm, but such coordination requires many people to move in unison."

We will apply this view of the development process to the fertility transition in rural Bangladesh. The International Centre for Diarrhoeal Disease Research, Bangladesh (ICDDR, B) launched a Maternal Child Health - Family Planning (MCH-FP) project in 1978, covering seventy villages in Matlab *thana*, Comilla district. The MCH-FP project is quite possibly the most intensive family planning program ever put in place: all households in the *intervention area* have been visited by a Community

Health Worker (CHW) once every two weeks since the inception of the project in 1978, and contraceptives are provided to them free of cost. Despite these economic incentives, and the sustained pressure on the households to change their behavior, we still see long delays in the adoption of contraceptives. The family planning program was already well established in the intervention area by the time our data begin in 1983. Nevertheless, contraception levels continued to increase steadily over the sample period (1983- 93), from 40% in 1983 to 63% in 1993, with an accompanying decline in total fertility rates from 4.5 children per woman to 2.9 children over that period. Wide variation in long-run contraceptive prevalence is also observed across villages in the intervention area.

Most societies have traditionally put norms into place to regulate fertility. In this Bangladeshi setting, the traditional norm was characterized by early and universal marriage, followed by immediate and continuous child-bearing. Religious authority provided legitimacy and enforced the rules that sustained this equilibrium. In such a social environment, the unexpected availability of modern contraceptives would have opened up the possibility for new equilibria, in which a sufficient fraction of the women in the village ignored the religious sanctions and began to regulate fertility.

The point of departure for our simple model of fertility change, following the exogenous economic intervention, is a social uncertainty: the individual does not know what level of contraceptive prevalence will ultimately be sustained in her community. This uncertainty is slowly resolved over time as women in the village interact sequentially with each other from one period to the next, which explains the gradual change in contraceptive prevalence that we see in the data, as well as the convergence to different levels of contraceptive use across communities.

While this simple model explains the broad stylized facts that we described above, it also generates implications for changes in the contraception decision rule at each point in time during the transition from the traditional equilibrium to the new equilibrium. In general, the contraception decision in any time period is determined by the individual's lagged decision and the level of contraceptive prevalence in the community. The individual places substantial weight on her own lagged decision, and there is a high level of state dependence, during the initial stages of the transition process. This weight shifts to neighbors' decisions as the transition progresses, before returning once more to the individual's own lagged decision as the community settles down to the new equilibrium.[1]

[1]An important methodological feature of our analysis is that a role for neighbors' decisions in the individual's decision function is derived endogenously and not simply assumed, as in the social interactions literature (Glaeser, Sacerdote and Scheinkman 1995, Topa 2001). Observed changes in the individual decision rule over the transition period will later allow us to distinguish our model of social change from these alternative models, which assume a time-stationary decision rule.

We test these implications of the model with a unique data set, which includes contraceptive use information as well as demographic and socioeconomic characteristics for *all* the women residing in the intervention area over an eleven year period (1983-93). We find that neighbors' past decisions play an important role in determining the individual's contraception decision, conditional on her own past decision. The changes in the estimated decision rule over the transition period also match the predictions of the model.

While these results are very encouraging, they do not by themselves conclusively establish the presence of an underlying social change. As Manski (1993) points out, a spurious correlation between the individual's decision and her neighbors' past decisions could be obtained when unobserved determinants of the contraception decision are correlated across individuals within the village and over time. For example, neighbors' decisions could simply proxy for changes in economic opportunities or the effectiveness of the MCH-FP project itself.

Our strategy in this paper, to provide additional support for the view that there was a social aspect to the individual decisions in this setting, takes advantage of the institutional background that we will provide in Section 2. Female mobility in Bangladesh has traditionally been severely restricted by the institution of *purdah*. Young married women will rarely leave the homestead (*bari*), and when they do it will typically be to visit extended family or kin. While the two major religious groups in rural Bangladesh, Hindus (who constitute 18% of the population in our villages) and Muslims, share a common language and a common Bengali culture, female interactions almost never cross religious boundaries even within the village. Consistent with this view, we present the striking result in Section 5 that while individuals respond strongly to contraceptive prevalence within their own religious group in the village, cross-religion effects are entirely absent in the data. In contrast, when we partition the village by other variables, such as age or education, we consistently observe large and significant cross-group effects.

We will also show in Section 5 that omitted determinants of the individual's contraceptive decision must be completely uncorrelated across religious groups within the village to spuriously generate the within-religion and cross-religion patterns that we just described. We will argue in that section that standard omitted variables, such as unobserved program effects or economic change, which complicated the interpretation of the estimated contraception decision rule above, are unlikely to satisfy this condition. For example, while health inputs and information-signals supplied by the MCH-FP project may have varied across religious groups within the village, it is difficult to imagine that they were uncorrelated across these groups. After all, it is the same agency, and the same CHW, that is providing these inputs. Overall the empirical evidence, over time as well as across religious groups within the village, supports our learning-based characterization of social change and, in particular, the view that social restrictions delayed changes in reproductive behavior.[2]

The paper is organized in six sections. Section 2 describes the institutional setting, paying special attention to the social restrictions that prevented the immediate adoption of contraceptives in the intervention area. Section 3 describes the village level data: we see a gradual change in contraceptive prevalence over time as well as wide variation in long-run contraceptive prevalence across the villages in the sample. Section 4 presents a simple model of social change that is consistent with these aggregate patterns. The individual's (optimal) decision rule is also derived in this section. Section 5 subsequently presents the data and the estimation results. We begin by verifying the predicted changes in the individual decision rule over time. Subsequently we turn to the cross-religion results, which provide additional support for the view that religion-specific social interactions gave rise to the changes in reproductive behavior that we see in the data. Section 6 concludes the paper.

2 The Institutional Setting

Our primary objective in this section is to describe the social restrictions that prevented the rapid spread of contraception in the intervention area. By making modern contraceptives available for the first time, the MCH-FP project ran counter to the practice of early and universal marriage followed by immediate and continuous child-bearing, followed traditionally throughout rural Bangladesh (Arthur and McNicoll 1978). Not surprisingly, the MCH-FP project faced strong opposition from community elders and local religious leaders, who were responsible for safeguarding the traditional norms. Apart from this social opposition, the spread of contraception was also hindered by the institution of *purdah*, which promotes the seclusion of women and rigidly segregates labor activities along gender lines (Amin 1997). *Purdah* severely restricts the mobility of young married women, which would have reduced social interactions between them, slowing down the process of social change and with it, the diffusion of modern contraceptives.

[2]While this paper provides a particular characterization of the process of social change, the observed changes in reproductive behavior could also have been motivated by changes in preferences (as in Pollak, 1976). We do not attempt to distinguish between these alternative models in this paper. We take the view that the preferred model is one that can explain all the observed patterns in the data, in the most parsimonious fashion, while maintaining the classical economic assumptions that the individual's preferences are exogenously determined and stable (see Postlewaite (1998) for an elaboration of this view).

Neither the Koran nor the Hindu religious scriptures take a firm position on contraception. Thus the attitude of the community in Bangladesh will typically follow the view of the local religious leader (Amin, Diamond and Steele 1997). Regulation of reproduction was seen to be contrary to religious principles by local leaders, and religious views opposing any fertility control persist in the intervention area to this day. Simmons (1996) describes conversations with women in the intervention area who retained their religious objections a decade after the institution of the MCH-FP project. As one of the women put it, "God has given the mouths and he will feed them. Therefore, we have nothing to do with it (contraception)."

While appeal to a "higher order" can have a powerful influence on individual behavior, stronger incentives may be required when the optimal action, from the individual's perspective, deviates significantly from customary behavior. Thus, norms are often associated with social sanctions, which may take the form of peer pressure, ridicule and even ostracism in extreme cases. We will see below that women using contraceptives in the intervention area did in fact face considerable overt pressure, although from our conversations with MCH-FP project field-staff this social opposition has weakened in recent year.

The MCH-FP project enlisted young married women residing in the local area as Community Health Workers (CHWs), to visit each household in their designated area (roughly the size of a village) twice a month providing contraceptives and health inputs. Since their job required them to venture outside the home without an appropriate chaperon, the CHWs directly violated the rules of *purdah*. Simmons, Mita and Koenig (1992) interviewed a small sample of CHWs in 1987-88 to study the social pressure that they faced at the outset of the MCH-FP project. The response from the community was initially extremely hostile. As one worker put it, "Many people used to say many bad things. They used to say that a *khanedarjal* (devil) has come to the village in order to destroy the women. They would not even look at us, or used to tease us whenever we passed their houses."

The hostile community response to the CHWs only reinforced the opposition to contraception in general. One woman interviewed by Simmons, Mita and Koenig describes the villagers' initial reaction in the following manner: "They used to think that she is doing the people harm. Why should she stop fertility? Let the embryo develop. God is the creator and he is giving birth. So she has no right to give tablets, injections, and other methods and stop pregnancy."

A further explanation for the religious opposition to fertility regulation is based on the idea that modern contraceptives reduce the risk that a woman's extramarital relations will be revealed. This is very important in a conservative society where such relations are severely punished when detected.[3] We would expect that the religious establishment, which enforces the moral code, would oppose the introduction of modern contraceptives on the grounds that they would encourage promiscuity and artificially alter the "natural order." This is indeed what appears to have happened; religious leaders in the intervention area went so far as to link contraceptive use, even by married women, to promiscuity in the early years of the MCH-FP project.

Although the religious resistance that we have described may be very persistent, its effectiveness weakens as a greater fraction of the community gradually deviates from the traditional social rules. Since this introduces a strategic aspect to the contraception decision, information about neighbors' decisions is clearly valuable. Unfortunately, access to such information is severely restricted by *purdah*. Patrilinearly related households are grouped together in homesteads or *baris* in Bangladesh, and groups of *baris* in turn form a village. Women in rural Bangladesh are confined to the *bari* and the area immediately surrounding it, and their contacts with the world outside of the family are extremely limited (Schuler and Hashemi 1994).

Cain, Khanam, and Nahar (1979), in a well known study, collected detailed time allocation information for men and women in one village (Char Gopalpur) in Mymensingh District of Bangladesh between 1976 and 1978. They found that both men and women work roughly 8.3 hours per day, but that men allocated 85% of their time to income-earning activities, while women allocate 81% of their time to home production. All the activities grouped under home production, with the possible exception of firewood collection, take place within the *bari*. A disproportionate share of women's income-earning time is allocated to handicrafts and hut construction, which also occur within the *bari*. Cain, Khanam, and Nahar (1979, pp. 428) thus conclude that "The physical limits of the market for a woman's labor are described by a circle with a radius 200-400 meters, with her homestead [*bari*] at the center of the circle."

These spatial patterns of female work activity do not appear to have changed appreciably over time, and more recent studies such as Amin's (1997) survey of two villages in Mohanpur Thana in 1991 show that women's work opportunities continue to be severely limited in rural Bangladesh. Opportunities

[3]Illegitimate birth is severely punished in the intervention area. Koenig et al. (1988) found that no pregnancy outside wedlock in Matlab resulted in a live-birth over the 1976-85 period; either due to abortion of the foetus, death of the mother as a result of abortion, or death of the mother from violence related to the pregnancy. In another study carried out in Matlab, Fauveau et al. (1989) show that despite the low occurrence of illicit coitus, 8% of all maternal deaths take place among unmarried women. 70% of these deaths result from unsuccessful abortion, while the balance is attributed to violence.

for women to travel outside of work are also extremely limited in Bangladesh. For instance, while the women in Cain, Khanam, and Nahar's study have complete responsibility for preparing meals at home, they do not themselves go to the local market to make purchases.

To assess the general level of female mobility in rural Bangladesh, Schuler, Hashemi, and Riley (1997) conducted a survey of 1,300 married women under age 50 in 1992 in which respondents were asked whether they had *ever* gone to the market, a medical facility, the movies, and outside the village. Each respondent was given one point for each place she had visited accompanied by someone else, and two points for each place she had visited alone. The mean score for the women in the sample was just 2.1 (out of a maximum of 8), reflecting the extremely low levels of mobility among the women. Note that while *purdah* is generally associated with Muslim societies, this concept of seclusion applies to both Hindus and Muslims in Bangladesh. It has been suggested that this is because the specific construction of *purdah* in Bangladesh, and its connections with gender-demarcated work patterns, are peculiar to Bengali culture (Rosario 1992).

About the only opportunity for young married women to travel in rural Bangladesh is on suitably chaperoned social visits to other *baris* in the village. But *baris* are geographically dispersed, and during the monsoon months they are often only accessible by boat. The infrequent interaction with friends, married sisters, and sisters-in-law, nevertheless slowly disseminates information through the community. A young woman in Mita and Simmons (1995) describes how her peer-group discussed "how many children we would have, what method would be suitable for us ... whether we should adopt family planning or not, all these topics ... We used to know from people that they used (contraceptives). If a couple takes any such method, the news somehow spreads." This slow diffusion of information matches well with the gradual change in contraceptive prevalence over time that we will observe in the next section, which will in turn motivate the learning-based model of social change that follows in Section 4. Women venture outside the *bari* to meet their kin, and Hindus and Muslims never inter-marry in rural Bangladesh, so female interactions occur exclusively within religious groups. This observation will also help explain the absence of cross-religion social effects, within the village, that we later observe in Section 5.

3 Aggregate Patterns in the Data

We now describe two important features of the data, which will motivate the model of social change that we present later in Section 4: the gradual increase in contraceptive prevalence over time and the sorting among the villages to different long-run levels of contraceptive prevalence.

We begin by describing the gradual change in contraceptive use over time. Contraceptive-use information for all eligible women, 15-49 years, married and capable of conceiving, is available at two points in each year (June 30 and December 31) over an eleven year period from 1983 to 1993. Note that only women capable of conceiving enter the data set at each point in time. Thus, while a woman will typically appear in the data set over many periods, she will not appear in those periods in which she is pregnant, nursing, or otherwise incapable of conceiving. Average contraceptive prevalence, measured as the proportion of all eligible women who use contraceptives, is presented in each year over the sample period in Figure 1.[4] Since we are using panel data, and women are more likely to use contraceptives as they get older, the nonparametric estimates account for age effects by differencing out the individual's age from a first-stage parametric regression.[5] Contraceptive prevalence increases slowly but steadily over time, although it does begin to flatten out after 1990.

Insert Figure 1 here.

Figure 1 also plots contraceptive prevalence separately for Hindus and Muslims (these are the dashed lines in the Figure). While Hindus maintain higher levels of contraceptive prevalence throughout the sample period, the gap between the two communities remains roughly unchanged. It is interesting to note that while the aggregate trajectories for Hindus and Muslims may track together, we will later observe absolutely no local interaction between these religious groups, within the village.

Women enter the data set when they marry (typically in their early twenties) and exit at the age of 49. There is thus wide variation in the year of birth among the women in the sample, which

[4] Figure 1, Figure 2, and Figure 3 below utilize the Epanechnikov kernel function to smooth the data. We use a fairly narrow band-width in these Figures, and essentially the same pattern (particularly in Figure 1 and Figure 2) would be obtained without smoothing. We use annual (December 31) data in Figure 1 and Figure 2 to be consistent with Figure 3, which uses annual village-level data from a different source. The change in contraceptive prevalence that we describe in Figure 1, and later in Figure 2, would be completely unaffected if we used higher frequency six-monthly data instead.

[5] The two-step estimation procedure is implemented as follows: In the first step, regress contraceptive use on the individual's age, age squared, and a full set of year dummies. Then difference out age and age squared, multiplied by the appropriate coefficients, from the contraceptive use variable. To preserve levels, then add on mean age for the entire sample and its squared value, multiplied by the age and age squared coefficients. In the second step, nonparametrically regress the modified contraceptive use variable on time.

covers the 1983-93 period. Figure 2 refines the patterns that we saw in Figure 1 by studying changes in contraceptive prevalence over time, separately by cohort. Four cohorts are defined, consisting of women born before 1948 (cohort 1), women born between 1948 and 1958 (cohort 2), women born between 1958 and 1968 (cohort 3), and women born after 1968 (cohort 4). The patterns that we saw earlier in Figure 1 continue to be obtained, particularly in the middle cohorts (cohort 2 and cohort 3), which together account for 75% of the observations in the data set. Notice also that the youngest cohort (cohort 4) converges to its long-run contraceptive prevalence relatively rapidly. Later in Section 4 we will describe a model of social change in which communities gradually find their way to a new reproductive equilibrium following the introduction of modern contraceptives. One possible explanation for the differences across cohorts is that cohort 4 may have entered the system at a point in time when the uncertainty about the new equilibrium had been largely resolved, allowing the members of this group to learn very quickly about the social regime in their community.[6]

Insert Figure 2 here.

Turning next to the second stylized fact that we wish to describe, Quah (1997) suggests using a simple transition matrix to study sorting among the villages to different long-run levels of contraceptive prevalence. While individual level data, with information on age and religious affiliation, are available from 1983 to 1993, aggregate village level contraceptive prevalence (the proportion of eligible women in each village who use contraceptives in any given year) is available over a longer period, starting from the inception of the MCH-FP project in 1978. The transition matrix in Table 1 thus allows us to study changes in the distribution of contraceptive prevalence from 1978 to 1993.

The numbers to the left of the box in Table 1 describe the distribution of contraceptive prevalence across the 70 villages in the intervention area in 1978, while the numbers above the box describe the corresponding distribution in 1993. The mean of the distribution increases from 0.27 to 0.55 over this period, consistent with the pattern that we saw earlier in Figure 1.[7] Notice, however, that the shape of the distribution, measured by the standard deviation and the inter-quartile range (the difference between the 0.25 and the 0.75 quantiles), is roughly the same in 1978 and 1993.

Insert Table 1 here.

[6]We use the same two-step procedure to difference out the age effects in Figure 2 as we did earlier in Figure 1. However, the first step regression now includes the individual's age, age squared, and a full set of year dummies for each cohort.

[7]The 1993 contraceptive prevalence in Table 1 differs slightly from the corresponding 1993 statistic in Figure 1 because we are computing the (unweighted) mean across villages, rather than across individuals, in the Table.

While the shape of the distribution may not have changed over time, this stability could still mask mobility within the distribution, as villages re-sort, leaving the overall distribution intact. To study such sorting, we turn to the cells within Table 1, which cover all possible transition possibilities in this simple system. For example, the number in the top left hand cell represents the probability that a village which began in the bottom quartile of the distribution in 1978 will remain in the same quartile in 1993. More generally, the numbers along the diagonal of the matrix represent the probability that villages remain in the same quartile that they began in. In the extreme case without state dependence, all the numbers in the transition matrix would be 0.25. Conversely, with complete state dependence, the diagonals would be one and all other cells would be zero. While the diagonal cells, and the cells (horizontally and vertically) adjacent to the diagonal cells, tend to be somewhat larger than 0.25 in Table 1, there is nevertheless a high level of mobility: the probability of remaining in the same quartile is 0.27 on average, and never exceeds 0.33.[8]

The intra-distributional mobility that we have just described suggests that initial conditions will not completely describe a village's position, within the contraceptive prevalence distribution, in 1993. Looking down any column in Table 1, we observe a fairly substantial contribution from each row, as expected. This tells us that villages are being drawn from across the 1978 distribution to fill each segment (quartile) of the 1993 distribution.

The transition matrix in Table 1 uses information at two points in time, 1978 and 1993, to describe the sorting of villages. We now proceed to study this sorting in more detail, separately tracking villages in the top and the bottom quartile of the 1993 distribution, over the entire sample period in Figure 3. Since we are studying intra-distributional sorting, we are most interested in each village's position *relative* to the other villages in the sample, at each point in time. It will therefore be convenient to difference out the mean contraceptive prevalence, computed using all 70 villages, from the village's own contraceptive prevalence (at each point in time) in the discussion that follows.[9]

[8]By way of comparison, Quah (1993) constructs a 5 x 5 transition matrix describing the change in the distribution of real per capita GDP for 118 countries over a 23 year period (1962-1984). With no state dependence the probabilities along the diagonals would be 0.20, but in fact these probabilities are as high as 0.60 on average.

[9]This is easily implemented in two steps by regressing contraceptive prevalence in each village-year on a full set of year dummies, and then nonparametrically regressing the residuals on time, separately for the two groups of villages. We could alternatively compute a rank statistic for each village at each point in time. Since the overall shape of the contraceptive prevalence distribution is unchanged over time, the two approaches are equivalent in our case. Note that we do not account for an age effect in the first step regression, as we did earlier in Figure 1 and Figure 2. But this does not affect our results because the data consists of married women, 15-49 years old and capable of conceiving, at each point in time. We have refreshment of the sample and so the average age in the village hardly varies over time; the difference between the 0.05 quantile and the 0.95 quantile of the mean-age distribution using all village-years in the sample is just two years. While not reported, if we were to include mean-age in each village-year in the first step

We begin by studying the trajectory for villages in the top quartile in 1993 in Figure 3. Following the previous discussion, we would expect these villages to be drawn from across the 1978 distribution, and their 1978 (differenced) mean is indeed close to zero in Figure 3. Over time, this group of villages gradually separates itself from the rest of the sample, ultimately converging to a relatively high level of contraceptive prevalence.

Insert Figure 3 here.

A symmetric, albeit weaker, trajectory is obtained for villages occupying the bottom quartile in 1993. While this group of villages is already below the sample average in 1978, it continues to drop further and further below the average contraception level over time. The dispersion in contraceptive prevalence within the group also narrows, although the convergence is not as pronounced as what we saw earlier with the high contraceptive prevalence group.

The two groups of villages start fairly close together in Figure 3, judging by the wide overlap in their 95% confidence bands in that year.[10] Subsequently they start to diverge, and the difference in (mean) contraceptive prevalence increases from 2% in 1978 to nearly 16% in 1993. The difference in contraceptive prevalence between the two groups of villages in 1993 works out to a corresponding difference in the total fertility rate of more than one child. This is a very large difference in fertility, emphasizing the wide variability in the local response to external interventions that motivates this paper. We saw in Figure 1 and Figure 2 that contraceptive prevalence started to flatten out around 1990, so we would expect to see very little change after our sample period ends in 1993. Indeed, while contraceptive prevalence increased from 40% to 63% over the 1983-93 period, it had only reached 70% by 1999 (the last year for which official data are available). While it is possible that the large differences in contraceptive prevalence across the villages in 1993 are caused by a highly serially correlated mean-reverting process, the stability in contraceptive prevalence in the 1990s suggests that these differences will persist at least in the medium-term.[11]

Note that the sorting in Figure 3 is not obtained mechanically, by our choice of the final year 1993 as the period in which to group villages. We saw earlier that the shape of the distribution,

regression (this is the only age statistic available at the village level) we would obtain estimates that are almost identical to what we see in Figure 3. This would also be true if we ignored the age effect in the first step regression for Figure 1, but not for Figure 2 where the average age of each cohort does grow over time.

[10] Pointwise confidence intervals are computed using a method suggested by Härdle (1990). Bootstrapped confidence intervals that allow for serial dependence in the differenced contraceptive prevalence are similar to the confidence intervals reported in Figure 3.

[11] We are grateful to an anonymous referee for bringing this point to our attention.

measured by the standard deviation and the inter-quartile range, was constant over time. Without intra-distributional sorting among the villages, we would have simply seen two horizontal lines in Figure 3, corresponding to the top and bottom quartile groups, without any intra-quartile convergence.

Note also that differences in religious composition among the villages cannot explain the divergence across villages in Figure 3. Recall from Figure 1 that the gap between the Hindus and the Muslims remains roughly constant over time, so villages with a disproportionate number of Hindus would remain near the top of the contraceptive prevalence distribution throughout the sample period if that were the case. By a similar argument, village size cannot explain the sorting in Figure 3, and we verified that village size is uncorrelated with contraceptive prevalence in 1993 in any case.[12] The sorting in Figure 3 must be due to differences across villages, that vary over time.

We propose a particular explanation for these time-varying differences, based on local village-level social interactions, in this paper. But it should be clear that Figures 1-3 could be explained by other village-level differences that change over time. For example, suppose that the Community Health Workers (CHWs) vary in their ability to (slowly) persuade women to use contraceptives, and suppose further that the gap between the workers grows over time. Now the gradual increase in contraceptive prevalence, as well as the difference across villages, could be explained by underlying heterogeneity in the quality of the MCH-FP project.

We do not attempt to distinguish between these alternative interpretations of the aggregate patterns at this point in the analysis. Our immediate objective is to present two stylized facts that will motivate the model of social change that follows in Section 4. First, contraceptive prevalence increased gradually in the intervention area, only flattening out in 1990 twelve years after the MCH-FP project. Second, there is wide variation in contraceptive prevalence across villages in 1993, which appears to be persistent. In addition, initial conditions cannot completely describe a village's final level of contraception. The individual level analysis, which controls for both age effects and time effects, and accounts for unobserved village-year effects with the cross-religion results, will be used to rule out alternative explanations for these stylized facts in Section 5.

[12]There are on average 330 eligible women in each village. But not all of these women appear in each year over the sample period; the average number of women in each village-year is 190. Contraceptive prevalence is regressed on this statistic in 1993 to verify that there is no link between contraceptive prevalence and village size.

4 A Simple Model of Social Change

Our first objective in this section is to present a model of decentralized social change that can explain the two stylized facts that we described above. The point of departure for the model is a social uncertainty following an exogenous economic intervention: the individual does not know the equilibrium that her community will ultimately converge to. We will see that this uncertainty is gradually resolved as individuals interact sequentially with each other over time.

There are only two types of individuals in our simple model, which is constructed so that only two possible equilibria can emerge in the long-run. No one regulates fertility prior to the intervention. While this continues to remain a potential equilibrium, we show that a new equilibrium in which a sufficient fraction of the community regulates fertility could also emerge. Much of this section is devoted to studying the process by which some communities gradually make the transition from the traditional equilibrium to the modern equilibrium, while other remain where they were.

The model can explain both the gradual change as well as the differences in long-run contraceptive prevalence across communities that we described in Section 3. It also generates implications for changes in the individual's optimal decision rule over the transition period, which we will successfully verify in Section 5.

4.1 Individual Payoffs and Social Equilibria

Each community consists of a continuum of individuals in our simple model of reproductive behavior with social regulation. An individual chooses from two actions at the beginning of each period: the traditional (t) action corresponding to unchecked fertility and the modern (m) action, which refers to fertility control. Subsequently she is randomly matched with a member of the community.

When reproductive behavior is socially regulated, the individual's payoff from a particular action depends not only on the intrinsic utility that she derives from that action, but also on the social pressure or sanctions that go with it. In our framework, which closely matches Kandori's (1992) characterization of social norms, the individual's payoff depends on her own action, as well as her partner's action, which determines the social sanction that she will face in that period.

Since there are two possible actions, and the individual matches with a single partner in each period, we must consider payoffs corresponding to four combinations of actions:

$$V_i(m, m) = U_i$$

$$V_i(m,t) = U_i - l$$

$$V_i(t,t) = 0$$

$$V_i(t,m) = g.$$

V_i is individual i's payoff at the end of the period, where the first term in parentheses refers to the individual's own action, and the second term refers to her partner's action. U_i is the intrinsic utility that the individual derives from the modern action. There are two types of individuals in our simple model, conformists and reformists, with reformists comprising a fraction P of the community. Conformists have very strong religious convictions, and we take it that they have internalized the religious opposition to reproductive control that we described in Section 2. Thus they derive lower intrinsic utility from the m action than the reformists; $U_i = v$ for the conformists, and $U_i = w > v$ for the reformists.

l and g refer to the punishment and rewards that have been put in place, perhaps by the religious establishment, to regulate reproductive behavior. When a woman who chooses the modern action meets another woman who continues to follow the traditional action, she faces some sort of social censure, which is presumably connected to the religious restrictions that we described in Section 2. The reward g may be associated with enhanced social standing, possibly within a very restricted peer group, for having punished a deviator. Notice that there are no social sanctions when two deviators meet each other.

We impose the following conditions on the payoffs prior to the external intervention: $v > 0$, $w - l < 0$, $w < g$. Under these conditions it is easy to verify that a unique equilibrium is obtained in each period, in which both conformists and reformists choose the t action.[13]

Subsequently we introduce an external intervention, which would correspond to the MCH-FP project in this application. We suppose that the availability of modern contraceptives reduces the inconvenience associated with fertility control, increasing the individual's intrinsic utility from the m action by an amount S. The conditions on payoffs in the post-intervention regime are the same as what we described above, with one important exception: $v + S > 0$, $w + S - l < 0$, $w + S > g > v + S$.

It is easy to verify that the traditional equilibrium, without fertility control, continues to be sustainable after the intervention. A new modern equilibrium can also be sustained if the proportion

[13]Each individual's payoff in this equilibrium is zero. No individual will deviate from this equilibrium since $v - l < w - l < 0$. Further, it is easy to verify that an equilibrium in which any group of individuals chooses the m action is unstable, since any member of that group would do better by deviating to the t action.

of reformists in the community P is sufficiently large. In this equilibrium all the reformists choose m and all the conformists choose t. A reformist will not deviate from this equilibrium if the expected payoff from choosing m exceeds the expected payoff from choosing t

$$P(w+S)+(1-P)(w+S-l) \geq Pg. \tag{1}$$

Simplifying the expression above, a necessary condition to sustain the modern equilibrium is obtained as $P \geq P^* = \frac{l-(w+S)}{l-g}$.[14] Communities with $P \geq P^*$ must choose between two equilibria, while only the traditional equilibrium can be supported in communities with $P < P^*$.

4.2 Social Uncertainty

The preceding discussion provides an explanation for the divergence across communities, to different social equilibria, following an external intervention. To explain the gradual transition to the long-run equilibrium in each community, we now introduce a social uncertainty. The basic source of uncertainty in our model is that the proportion of reformists P is not known to begin with, since each individual's type is private information, and both conformists and reformists chose the same traditional action prior to the intervention. To simplify the equilibrium dynamics we assume that there are two types of communities: stable communities with $\underline{P} < P^*$ reformists and unstable communities with $\overline{P} > P^*$ reformists. We will see below that information about P is gradually revealed over time as individuals interact with each other, with unstable communities moving to the modern equilibrium while stable communities remain where they were.

While we focus on uncertainty about social fundamentals (the underlying social structure of the community), a model based on strategic uncertainty could also deliver the aggregate patterns that we see in the data. Suppose, for example, that all the communities are unstable, with $\overline{P} > P^*$. We are still left with a coordination problem, since both the traditional and the modern equilibrium can be sustained in these communities. The standard approach to model this coordination problem would be to perturb the system by exogenously switching a fraction of the community to the m -action in period 0. If we assume that individuals mimic their partner's action (in the next period) with a fixed

[14]Note that $l-(w+S) < l-g$ since $w+S > g$. We also assumed $w+S-l < 0$ above. This ensures $0 < P^* < 1$. All the conformists choose t and all the reformists choose m in the modern equilibrium. It is easy to verify that no conformist ever wishes to deviate and choose m since $v+S-l < 0$, $v+S < g$. Note that a modern equilibrium in which only some of the reformists choose m is unstable. Take the case of a community with $P > P^*$, where only P^* of the individuals (all of whom are reformists) choose m. Clearly, any reformist choosing t in this situation would prefer the m action, since the proportion of deviators has reached the P^* threshold.

probability, then we are essentially describing the beginning of a contagion. It is well known that if the initial perturbation is sufficiently large, then the community will "tip over" to the modern equilibrium, if not it will return to the traditional equilibrium after a temporary deviation.[15] Thus the contagion model can explain both the gradual change in behavior as well as the deviation across communities that we observed in Section 3.

There are two reasons why we prefer our model, based on social fundamentals, to this alternative model, based on strategic uncertainty. First, one of the important objectives of this paper is to explain why different communities respond so differently to the *same* external intervention. For example, every attempt was made to standardize the MCH-FP program across villages, yet we see tremendous variation in contraceptive use across communities that otherwise look fairly similar. We believe that differences in the underlying social structure, measured by the P parameter in the model, may explain much of the variation that is typically observed in developing countries. While we believe that coordination plays an important role in the development process as well, notice in contrast that the contagion model must rely on differences in the initial perturbation (the external intervention) to generate different long-run equilibria across communities.

A second advantage of our model is that social learning about P occurs in a Bayesian setting (as in Banerjee 1992, 1993, Bikhchandani, Hershleifer and Welch 1992). While previous research on social norms (Kandori 1992, Okuno-Fujiwara and Postlewaite 1995, Ellison 1994) has studied how patterns of cooperative behavior can be sustained when any two individuals in the community do not interact repeatedly with each other, to the best of our knowledge this is the first attempt, using a Bayesian framework, to investigate how the community actually moves from one social equilibrium to the other. The individual's (optimal) decision rule is derived endogenously in our case, and is not simply assumed as in the contagion models described above, or the social interactions literature (Glaeser, Sacerdote and Scheinkman 1996, Topa 2001). Moreover, our model generates implications for changes in the individual decision rule over the transition period, which we will subsequently verify, in contrast with these alternative models which assume a static decision rule.

4.3 Equilibrium Dynamics

The analysis of the equilibrium dynamics proceeds in three steps. We begin by describing the exogenous perturbation, associated with the MCH-FP project, that is needed to initiate the transition.

[15]For example, see Eshel, Samuelson and Shaked (1998), or the references cited in Banerjee (1993).

Subsequently we describe the change in the distribution of beliefs α, the probability that $P = \overline{P}$, in both stable and unstable communities over the course of the transition. We will show that the distribution of beliefs shifts gradually until $\alpha = 0$ for all individuals in the stable communities and $\alpha = 1$ for all individuals in the unstable communities in the long-run. We finally proceed to map these changes in beliefs into changes in actions (the proportion of m's in the community), deriving the dynamic path from the traditional to the modern equilibrium in the unstable communities, as well as the return to the traditional equilibrium (after some temporary deviation) in the stable communities.

4.3.1 Initiating the Transition

All communities are in the traditional equilibrium prior to the external intervention. Immediately following the intervention, they continue to remain in that equilibrium. Since the ICDDR,B has exogenously increased the payoffs from the m action, it is evidently interested in making sure that the reformists in the unstable communities take advantage of the new opportunities that are available. To achieve this objective it employs Community Health Workers (CHWs) to persuade individuals to choose the modern action.

We make the following assumptions about the exogenous perturbation. First, the CHW remains permanently in place in our model. When describing the contagion model in the previous section, we discussed how an initial perturbation in period 0 was required to move unstable communities to the modern equilibrium. In a Bayesian setting we will see that beliefs about P, and subsequent decisions, change relatively slowly, so it is necessary to keep the CHWs in place for multiple periods.

While we keep each CHW permanently in place, we maintain the temporary nature of the perturbation by assuming that her ability to influence the individuals that she comes in contact with is temporary. Thus our second assumption is that the CHW visits a fraction θ of the community in each period, drawn at random, and persuades any reformist that she meets to switch to the m action, but for a *single* period only. Since there is a continuum of individuals in each community, this implies that a constant fraction θP of the community, where $P = \overline{P}$ in unstable communities and $P = \underline{P}$ in stable communities, deviates *exogenously* in each period. We will see that this exogenous deviation provides the seed for subsequent *endogenous* deviation in the unstable communities, which ultimately moves them to the new social equilibrium.

Finally, our third assumption is that the value of θ is common knowledge. If the actions in the community (a fraction θP of the individuals choose m) were also common knowledge, then P would be revealed in the first period itself. Instead, we follow the standard set up in the social norms literature in which each individual matches with a single partner in each period. The decision whether or not to use contraceptives has serious implications for the household's future welfare, and while this decision may change over time, it is difficult to imagine that it will change very frequently. There is also a technological constraint on the frequency with which women can change their fertility behavior; the dominant method of contraception in the intervention area - injectables - is effective for a period of three months. Given the severe restrictions on female mobility that we described in Section 2, and the possibility that not all meetings will result in conversations about contraception, it is plausible that the frequency of decision-making and the frequency of social interactions have similar magnitudes in this setting. These low frequencies imply that the sequence of matches over time will only gradually reveal the proportion of reformists in the community.[16]

4.3.2 The Change in Beliefs

Once we have described the exogenous perturbations to the system, the next step in the analysis is to derive the evolution of beliefs over time. We make the standard rational expectations assumption that the individual correctly predicts the proportion of m's that will be realized in stable and unstable communities at each point in time over the transition. The only source of uncertainty for the individual is the type of community that she belongs to. Let $\alpha_t \in [0, 1]$ be the individual's belief about the state of the world, the probability that $P = \overline{P}$, in period t. The individual uses Bayes' Rule to update her belief from one period to the next, based on her partner's action and the proportion of m's in stable and unstable communities that she knows will be realized in that period. Thus different individuals will have different beliefs at each point in time, depending on their particular history of matches.

It will be convenient, in the discussion that follows, to assume a continuous distribution of beliefs in the community at each point during the transition process, although this assumption is relaxed in the simulations that we present later.[17] Let the distribution of beliefs among the *reformists* in period

[16] If the CHW observes all the individual decisions, then P would be revealed to the external agency in the first period itself. Immediate withdrawal by the external agency would signal in turn whether a community was stable or unstable. In our setting, the ICDDR,B must maintain a permanent standardized program across all the 70 villages to satisfy the research objective of the MCH-FP project. More generally the external agency may prefer to maintain a long-term presence even in the stable communities in an effort to change preferences (P) or the social rules (l, g), which we treat as stable in our model.

[17] This assumption is clearly inconsistent with the matching process specified in our model. For example, starting with a degenerate belief distribution in period 0, we would have a bimodal distribution in period 1, and in general it would be many periods before the distribution started to fill out. The assumption that the distribution of beliefs is continuous will, however, simplify the analysis that follows considerably, and we will see that our analytical results match the simulations that we present later (which allow for a discrete distribution) very well.

t be characterized by c.d.f. $\overline{F}_t$, $\underline{F}_t$, in unstable and stable communities respectively.[18]

While the value of P may not be revealed immediately, no one is systematically misinformed through their social interactions in our model. Thus we would expect that in the long run, beliefs in the unstable communities would pile up at $\alpha = 1$, and in the stable communities at $\alpha = 0$.[19] This process, in which the mass of the distribution shifts to the right in the unstable communities, and in the opposite direction in the stable communities, is described in Figure 4.

Insert Figure 4 here.

Let the ability distribution in period t span the range $[\alpha_{Lt}, \alpha_{Rt}]$, where we know from Figure 4 (and will see below) that the support of the distribution is the same in both stable and unstable communities. Next, denote beliefs α'_{Lt}, α'_{Rt} such that a single match moves the individual's belief from α'_{Lt} to α_{Lt}, and similarly from α'_{Rt} to α_{Rt}. A belief α is said to lie in the tail of the distribution if $\alpha \in [\alpha_{Lt}, \alpha'_{Lt}]$ or $\alpha \in [\alpha'_{Rt}, \alpha_{Rt}]$. The dynamics that we describe in Figure 4 can then be expressed as follows:

Proposition 1 *At any belief α, except in the tail of the belief distribution, the flow of beliefs to the right will dominate the flow to the left in unstable communities, $\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha) < 0$, whereas the direction of the flow is reversed in stable communities $\underline{F}_{t+1}(\alpha) - \underline{F}_t(\alpha) > 0$.*

The proof is reported in the Appendix. Since the individual matches with a single m or t in each period, the change in beliefs from one period to the next is very restricted. Applying Bayes' Rule, we can define a neighborhood around α, bounded by beliefs $\alpha(L)$ and $\alpha(R)$, which is relevant when determining $\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha)$, $\underline{F}_{t+1}(\alpha) - \underline{F}_t(\alpha)$. Only individuals with beliefs in $[\alpha(L), \alpha]$ in period t can shift to the right of α (with a single match) in period $t+1$. Similarly, it is only individuals in $[\alpha, \alpha(R)]$ in period t who can shift to the left of α in period $t+1$. Deriving the probability of shifting and the size of each of these neighborhoods in stable and unstable communities, it is easy to verify Proposition 1.

As noted, we make the standard rational expectations assumption that the individual correctly predicts the proportion of m's in any unstable community $\overline{x}_t$, as well as the corresponding proportion

[18] There is no need to characterize the distribution of beliefs among the conformists since they will choose the traditional action in any case.

[19] The process of learning that we describe in this paper is conceptually related to Banerjee's (1993) characterization of a rumor process. In his model, the delay before individuals meet reveals the state of the world. In our case, individuals match every period; it is the sequence of partners' decisions that ultimately reveals the type of community that the individual belongs to.

in any stable community $\underline{x}_t$, when she updates her belief about the type of community that she belongs to, in any period t. Proposition 1 is obtained without placing any other restrictions on $\overline{x}_t$, $\underline{x}_t$: for both $\overline{x}_t > \underline{x}_t$ as well as $\overline{x}_t < \underline{x}_t$. Later in this section we will proceed to map the changes in beliefs that we have just derived into changes in actions, to generate $\overline{x}_t$, $\underline{x}_t$, and complete the characterization of the equilibrium dynamics.

The only complication that is introduced when deriving these changes in beliefs during the transition is that the differences across communities do not apply to beliefs in the tail of the distribution in any period. For example, consider the lowest belief in the support of the distribution in period t, α_{Lt}. Some individuals just to the right of this belief will certainly shift to the left of it in period $t+1$, depending on whom they match with. Since there are no beliefs to the left of this minimum belief, in period t, the net flow *must* be to the left in both stable and unstable communities. Similarly, net flows must be to the right of the maximum belief α_{Rt} in both types of communities. This restriction on the change in beliefs is also saying that the support of the distribution of beliefs must be spreading over time in both types of communities, just as we described in Figure 4. This observation will come in useful below when deriving the change in actions over the transition.

4.3.3 The Change in Actions

Once we have described how beliefs change over time in the two types of communities, the final step in characterizing the equilibrium dynamics is to map beliefs into actions. Specifically, we want to derive the change in the proportion of m's, $\overline{x}_t$, $\underline{x}_t$, in unstable and stable communities over the transition.

It is convenient to begin with a degenerate distribution of beliefs α_0 in period 0, in both stable and unstable communities, such that no reformist deviates endogenously. Thus we only observe exogenous deviation in the first few periods: the (constant) proportion of m's in the unstable communities is given by $\overline{x}_t = \theta\overline{P}$, with a corresponding proportion $\underline{x}_t = \theta\underline{P}$ in the stable communities. While contraceptive prevalence might be constant in these early periods, the distribution of beliefs within each community will spread out over time as different individuals are faced with a different sequence of matches.

To derive the evolution of individual beliefs during these early periods without endogenous deviation apply Bayes' Rule to an individual with belief α_t in period t who matches with an m in that period. Her belief α_{t+1} in the subsequent period is then expressed as:

$$\alpha_{t+1} = Pr(P = \overline{P} \mid m) = \frac{\alpha_t(\theta\overline{P})}{\alpha_t(\theta\overline{P}) + (1-\alpha_t)(\theta\underline{P})}. \tag{2}$$

Since the term in the denominator of equation (2) is a weighted average of $\theta\overline{P}$ and $\theta\underline{P}$, it is easy to verify that $\alpha_{t+1}/\alpha_t > 1$. As the individual matches with m's in the community, her belief that $P = \overline{P}$ grows. The right-hand support of the distribution in any period t is thus defined by the beliefs of the individuals in the community who have matched with a continuous sequence of m's up to that period, and so will shift steadily over time.

A reformist will choose the m action in any period, without persuasion from the CHW, if the expected probability of matching with an m exceeds P^* (from equation (1)). The expected probability of matching with an m in these early periods is simply $\alpha(\theta\overline{P})+(1-\alpha)(\theta\underline{P})$, where α is the individual's belief that $P = \overline{P}$. This expected probability is evidently increasing in α. As long as $\theta\overline{P} > P^*$, there exists a threshold belief α^*, for which the individual is indifferent between the t and the m action, satisfying the following condition:

$$\alpha^*(\theta\overline{P}) + (1-\alpha^*)(\theta\underline{P}) = P^*. \tag{3}$$

If the individual's belief that $P = \overline{P}$ exceeds α^*, then the expected probability of matching with an m will exceed P^*, and she will deviate endogenously. If not, she will only choose m when she meets the CHW. Following the discussion above, the support of the belief distribution will shift steadily over time in the early periods until ultimately the right-hand support reaches α^*. The first wave of endogenous deviators will now appear, at the same time in both stable and unstable communities.[20]

To describe actions in the community after the first wave of endogenous deviators appears, it will be convenient initially to fix the threshold belief α^* to be the same in every period. The proportion of reformists that has crossed the threshold in any period t is simply $1-\overline{F}_t(\alpha^*)$, $1-\underline{F}_t(\alpha^*)$ in the unstable and stable communities. The outcomes that we are interested in, $\overline{x}_t$, $\underline{x}_t$, measure the proportion of m's in the *entire* community, so we need to normalize by the proportion of reformists in each case. Taking into account the exogenous deviators who have yet to cross the belief threshold as well, $\overline{x}_t = \overline{P}[1-(1-\theta)\overline{F}_t(\alpha^*)]$, $\underline{x}_t = \underline{P}[1-(1-\theta)\underline{F}_t(\alpha^*)]$. From Proposition 1 we know that beliefs shift to the right in unstable communities $\overline{F}_{t+1}(\alpha^*) < \overline{F}_t(\alpha^*)$, whereas the direction of the flow is reversed in stable communities $\underline{F}_{t+1}(\alpha^*) > \underline{F}_t(\alpha^*)$. It follows directly from the expressions for $\overline{x}_t$, $\underline{x}_t$

[20]Note the coordinating assumption $\theta\overline{P} > P^*$ embedded above in equation (3), which is necessary to support this initial endogenous deviation. Just as with the contagion model, a sufficiently large perturbation (θ) is required to jolt the community to a new equilibrium. Without this coordinating assumption, the equality in that equation would never be satisfied, the first wave would never emerge, and both stable and unstable communities would remain in the traditional equilibrium.

above that $\overline{x}_{t+1}(\alpha^*) > \overline{x}_t(\alpha^*)$, $\underline{x}_{t+1}(\alpha^*) < \underline{x}_t(\alpha^*)$: the proportion of m's is monotonically increasing (decreasing) in unstable (stable) communities.

While the result that we have just derived holds for α^* almost everywhere in the distribution of beliefs, recall that Proposition 1 cannot be applied when α^* lies in the tail of the distribution. This would be the case when the first wave of endogenous deviators appears, since we noted earlier that the right-hand support of the distribution α_R just reaches α^* at that point. Beliefs must flow to the right in both types of communities immediately following the first wave, so $\underline{x}_t$ must increase temporarily as well. However, once the support of the distribution has shifted sufficiently to the right of α^*, the result that we derived above begins to apply and $\underline{x}_t$ will decline monotonically thereafter.

The discussion up to this point treated the threshold belief as constant over time. More generally, this threshold belief would be derived endogenously in each period, and α^* would be replaced by α_t^* in the expressions above. Once we allow α_t^* to change over time, we cannot describe the details of the equilibrium dynamics without characterizing the distribution of beliefs at each point in time. We can, however, continue to say something about the early periods, as well as the long-run equilibria, in both types of communities. We know from Proposition 1 that the mass of the distribution of beliefs in the unstable communities will pile up at $\alpha = 1$, so as long as α_t^* remains strictly in the interior of the unit interval, the proportion of m's will begin at $\theta\overline{P}$ and end up at $\overline{P}$. Similarly, the stable communities will begin at $\theta\underline{P}$, increase temporarily, and then ultimately return to where they began.

To fully characterize the equilibrium dynamics, we solve the following system of equations:

$$\alpha_t^*\overline{x}_t + (1-\alpha_t^*)\underline{x}_t = P^* \tag{4}$$

$$\overline{x}_t = \overline{P}[1-(1-\theta)\overline{F}_t(\alpha_t^*)] \tag{5}$$

$$\underline{x}_t = \underline{P}[1-(1-\theta)\underline{F}_t(\alpha_t^*)]. \tag{6}$$

While the model cannot be solved analytically, it is fairly easy to simulate the equilibrium dynamics, since the system of equations (4) - (6) can be solved independently in each period. Starting with a degenerate distribution of beliefs in the first period, we only need to keep track of each individual's beliefs over time, rather than the entire history of matches, so the simulated distribution of beliefs can be used directly to solve iteratively for α_t^*, $\overline{x}_t$, $\underline{x}_t$ in each period. These simulation results, presented in Figure 5 below, are very robust to the choice of parameter values and match the general predictions of the model:[21]

Proposition 2 *After an initial delay, the proportion of m's in unstable communities increases over time, starting at $\theta\overline{P}$ and converging to $\overline{P}$ in the long-run. After the same delay, and a temporary increase, the proportion of m's in stable communities begins to decline and ultimately returns to $\theta\underline{P}$.*

Insert Figure 5 here.

The model with two types can be readily extended to multiple types. For example, take the case with three types: conformists, reformists, and super-reformists. Suppose that reformists deviate above a threshold probability P^*_W, while super-reformists deviate above a corresponding probability $P^*_Z < P^*_W$. It is easy to verify that three equilibria can emerge in this case: the proportion of reformists and super-reformists is greater than P^*_W and both types deviate in unstable (WZ) communities, the proportion of super-reformists is greater than P^*_Z and only those types deviate in unstable (Z) communities, or the community is stable and no one deviates in the long-run. The simulation results with three types of communities are presented in Figure 6 below, continuing to match the aggregate patterns predicted by the model.[22]

Insert Figure 6 here.

4.4 The Individual Decision Rule

Uncertainty about the social fundamentals is resolved over time as individuals interact with each other, and we saw above that communities gradually separate to different long-run equilibria. In the discussion that follows, we will derive the individual decision rule during this transition. The discussion in this section restricts attention to reformists, since conformists always choose the t action.

[21] Parameter values are set at $\theta = 0.75$, $\underline{P} = 0.375$, $\overline{P} = 0.625$, $P^* = 0.45$, $\alpha_0 = 0.26$ with 400 individuals in each community, for the simulations reported in Figure 5. Note that the temporary increase in contraceptive prevalence in the stable community is hardly discernable on account of the parameter values that we have chosen for the simulation. In a previous version of the paper we set $\alpha_0 = 0.4$ to generate a sharp increase, at the same time, in both types of communities.

[22] To characterize the equilibrium dynamics we now need to solve a system of five equations, analogous to equations (4) - (6) above. This is a slightly more difficult model to solve because two beliefs are associated with the individual at each point in time: α^Z and α^{WZ}, the probability that she belongs to an unstable(Z) and an unstable(WZ) community, respectively. We must identify a marginal reformist and super-reformist in each period, whose beliefs leave them just indifferent between the m and the t action, instead of solving for a belief threshold α^* explicitly, as in Figure 5. It nevertheless remains fairly easy to simulate the equilibrium dynamics, since the system of equations can be solved independently in each period. Parameter values are set at $\theta = 0.75$, $P^Z_Z = 0.6$, $P^Z_W = 0.1$, $P^{WZ}_Z = 0.6$, $P^{WZ}_W = 0.2$, $\underline{P}_Z = 0.4$, $\underline{P}_W = 0.1$, $P^*_Z = 0.45$, $P^*_W = 0.75$, $\alpha^Z_0 = 0.05$, $\alpha^{WZ}_0 = 0.05$, with 400 individuals in each community for the simulations reported in Figure 6. Here subscripts refer to the individual's type, while superscripts refer to the type of community.

Returning to the model with two types, the individual's decision in period t is determined by her belief, relative to the threshold belief α_t^*. If her belief lies to the right (left) of α_t^*, she will choose the m (t) action. The individual's belief in period t is in turn determined by her belief in period $t-1$, augmented by the change in this belief through the social interaction in that period. It is easy to see that matching with an m will shift her belief to the right, by returning to equation (2) and replacing $\theta\overline{P}$, $\theta\underline{P}$ with $\overline{x}_{t-1}$, $\underline{x}_{t-1}$. We saw that $\overline{x}_{t-1} > \underline{x}_{t-1}$ in Figures 5-6, and this result is obtained without exception with all the parameter values that we experimented with in those Figures. This implies that $\alpha_t/\alpha_{t-1} > 1$ when the individual matches with an m.

There is a high level of state dependence in this system, since it is only individuals with beliefs in a left window $[\alpha(L)_{t-1}^*, \alpha_{t-1}^*]$ or a right window $[\alpha_{t-1}^*, \alpha(R)_{t-1}^*]$, around the threshold belief α_t^*, who can change their actions from period $t-1$ to period t . Individuals in $[\alpha(L)_{t-1}^*, \alpha_{t-1}^*]$ choose the traditional action in period $t-1$, but will switch to the modern action if they match with an m. Similarly, individuals in $[\alpha_{t-1}^*, \alpha(R)_{t-1}^*]$ will switch from the modern to the traditional action if they match with a t. Individuals with beliefs outside $[\alpha(L)_{t-1}^*, \alpha(R)_{t-1}^*]$ will not change their actions, regardless of whom they match with in period $t-1$.[23]

The preceding discussion tells us that the individual's decision in any period will be determined by her belief at the beginning of the previous period, and her social interaction in that period. Since beliefs are unobserved by the econometrician, we will proceed to derive the individual's decision rule in terms of her lagged *decision* in the discussion that follows.

An additional difficulty that arises when deriving the individual decision rule is that the response to neighbors' decisions will vary within the community at each point in time, depending on each individual's location in the belief distribution (relative to the threshold belief). We will consequently derive the decision rule for a representative individual, drawn randomly from the community, at each point in time during the transition.

Let $y_{it} = 1$ if individual i chooses the m action in period t , and let $y_{it} = 0$ if she chooses the t action. The probability that $y_{it} = 1$, conditional on $y_{it-1} = 0$, is the product of two probabilities; the probability of lying in the left window (conditional on $y_{it-1} = 0$), and the probability of matching with an m. Since the individual is selected randomly from the community, the probability that her belief will lie in the left window is simply $\Delta F(L)_{t-1}/F_{t-1}(\alpha_{t-1}^*)$, where $\Delta F(L)_{t-1} \equiv F_{t-1}(\alpha_{t-1}^*) - F_{t-1}(\alpha(L)_{t-1}^*)$. Note that we no longer distinguish between stable and unstable communities, to simplify the exposition.

[23] $\alpha(L)_{t-1}^*$, $\alpha(R)_{t-1}^*$ can be derived using a straightforward application of Bayes' Rule, as in the Appendix.

Further, since individuals match randomly with each other within the community, the probability of matching with an m in period $t-1$ is simply x_{t-1}, the proportion of m's in that period.

$$Pr(y_{it} = 1 \mid y_{it-1} = 0) = \frac{\Delta F(L)_{t-1}}{F_{t-1}(\alpha^*_{t-1})} x_{t-1}. \tag{7}$$

By the same type of argument, it is easy to derive the corresponding expression for the probability that $y_{it} = 1$, conditional on $y_{it-1} = 1$. The probability that the randomly selected individual will occupy the right window, conditional on having chosen m in the previous period, is expressed as $\Delta F(R)_{t-1}/1 - F_{t-1}(\alpha^*_{t-1})$, where $\Delta F(R)_{t-1} \equiv F_{t-1}(\alpha(R)^*_{t-1}) - F_{t-1}(\alpha^*_{t-1})$. Recall that individuals who chose m in period $t-1$ will continue with that action unless they lie in the right window and match with a t. Thus, we have

$$Pr(y_{it} = 1 \mid y_{it-1} = 1) = 1 - \frac{\Delta F(R)_{t-1}}{1 - F_{t-1}(\alpha^*_{t-1})}(1 - x_{t-1}). \tag{8}$$

Combining equation (7) and equation (8), we derive a simple expression for the individual decision rule in terms of the lagged decision and lagged contraceptive prevalence in the community:

$$Pr(y_{it} = 1 \mid y_{it-1}) = \frac{\Delta F(L)_{t-1}}{F_{t-1}(\alpha^*_{t-1})} x_{t-1} + \left[1 - \left\{\frac{\Delta F(L)_{t-1}}{F_{t-1}(\alpha^*_{t-1})} x_{t-1} + \frac{\Delta F(R)_{t-1}}{1 - F_{t-1}(\alpha^*_{t-1})}(1 - x_{t-1})\right\}\right] y_{it-1}. \tag{9}$$

While the lagged decision y_{it-1} and lagged contraceptive prevalence x_{t-1} are observed by the econometrician, the coefficients in the decision rule are expressed in terms of unobserved underlying beliefs. These beliefs, however, were previously simulated, prior to generating $\overline{x}_{t-1}$, $\underline{x}_{t-1}$ in Figure 5. Restricting attention to the unstable community, we present the simulated decision rule parameters $\Delta F(L)_{t-1}/F_{t-1}(\alpha^*_{t-1})$, $\Delta F(R)_{t-1}/1 - F_{t-1}(\alpha^*_{t-1})$, as well as the expression in braces in equation (9) which is just the weighted average of these parameters, in Figure 7.[24]

Insert Figure 7 here.

We see in Figure 7 that both decision rule parameters start close to zero, increase initially, and then decline steadily over time. Since Figure 5 and Figure 7 are both generated by the same underlying

[24] We use the same set of parameters in this simulation as in Figure 5, and we implicitly assume in this case that the simulated beliefs approximate a continuous distribution. Note that these simulations are unstable in the early periods when $1 - F_{t-1}(\alpha^*_{t-1})$ is close to zero, and in the last few periods when $F_{t-1}(\alpha^*_{t-1})$ is close to zero. Figure 7 thus covers period 10 to period 50 only.

social learning process, it will be useful to compare these figures when explaining the patterns that we have just described.

To begin with, notice that the decision rule parameters in Figure 7 start to increase at a point in time when the endogenous deviators first start to appear in Figure 5. Prior to that time, the belief windows are empty, $\Delta F(L)$, $\Delta F(R)$ are both zero, and $F = 1$. Once the first reformists cross the belief threshold α^*, we move into the growth phase in Figure 5. F starts at one at the beginning of this phase, and then steadily declines, while $\Delta F(L)$, $\Delta F(R)$ increase. While $\Delta F(L)/F$ must certainly increase in the growth phase, we see in Figure 7 that $\Delta F(R)/1 - F$ increases as well, presumably because the rise in $\Delta F(R)$ is sufficiently steep.

Once the growth phase is over and a substantial proportion of the reformists have crossed the belief threshold, the community will gradually settle down and converge to the new equilibrium, as in Figure 5. F continues to decline steadily towards zero (remember that this is an unstable community) during this slow growth phase. Since there are fewer and fewer reformists left to cross the belief threshold, the belief windows will also shrink over this period.[25] While $\Delta F(R)/1 - F$ must certainly decline in this period, towards zero, we see in Figure 7 that $\Delta F(L)/F$ declines slowly as well.

Returning to equation (9), the preceding discussion and Figure 7 tells us that the coefficient on lagged contraceptive prevalence will start at zero, increase during the growth phase, and then decline slowly as the community gradually converges to the new equilibrium. The coefficient on the lagged decision will display the opposite pattern; it is close to one initially, declines during the growth phase, and then slowly begins to rise again.[26]

Following this discussion, the decision rule for a representative individual, drawn randomly from the community at each point in time over the transition can be described as follows:

Proposition 3 *The individual's contraception decision in any period is determined by her own lagged decision and the lagged level of contraceptive prevalence in the community. The individual places substantial weight on her own lagged decision, and there is a high level of state dependence, during the initial stages of the transition. This weight shifts to neighbors' decisions as the transition progresses, before returning once more to the individual's own lagged decision as the community settled down to*

[25]While we do not report these simulations separately, $\Delta F(L)$ and $\Delta F(R)$ are both increasing during the growth phase in Figure 5, before declining steadily to zero as contraceptive prevalence converges to its long-run level.

[26]These changes in the coefficients on x_{t-1}, y_{it-1} distinguish our Bayesian model from the alternative contagion model, which assumes that individuals mimic their neighbors' actions with a fixed probability (ϕ). With the contagion model, $Pr(y_{it} = 1 \mid y_{it-1} = 0) = \phi x_{t-1}$, $Pr(y_{it} = 1 \mid y_{it-1} = 1) = 1 - \phi(1 - x_{t-1})$, and hence, $Pr(y_{it} = 1 \mid y_{it-1}) = \phi x_{t-1} + (1 - \phi) y_{it-1}$.

the new equilibrium.

An alternative (equivalent) specification of the decision rule separates out the $y_{it-1} \cdot x_{t-1}$ interaction term in equation (9). The coefficient on this term, $\frac{\Delta F(R)_{t-1}}{1-F_{t-1}(\alpha^*_{t-1})} - \frac{\Delta F(L)_{t-1}}{F_{t-1}(\alpha^*_{t-1})}$, is almost always negative in Figure 7. We also see from the Figure that the (modified) coefficient on y_{it-1}, $1 - \Delta F(R)_{t-1}/1 - F_{t-1}(\alpha^*_{t-1})$, retains its earlier properties; it is large initially, declines in size as contraceptive prevalence grows, and then increases once more as the transition draws to a close. The coefficient on x_{t-1} is unchanged. Later in the estimation section we will verify that our results are robust to the inclusion of the interaction term in the contraception regression.

Finally, note that while the simulations in Figure 7 are specific to the unstable community in a model with two types, we would expect Proposition 3 to apply more generally to any community in which contraceptive prevalence increases steadily over time, in a model with many types. Later in Section 5 we will verify that the predicted changes in the learning rule are indeed obtained in villages with very different levels of long-run contraceptive prevalence.

4.5 Beyond the Model: Allowing the Social Regime to Change over Time

The model assumes that the social regime, characterized by the rewards g and the punishment l, is stable over time. Change in contraceptive prevalence is associated with learning about the social fundamentals P alone. However, it is easy to imagine that changes in contraceptive prevalence could in turn induce changes in the underlying social regime. Political scientists (for example, Lohmann 1994) have documented the process by which small political protests can generate increasingly large cycles of protests, which ultimately may lead to changes in the political regime. In the same way, initial deviation from the traditional equilibrium in our villages could have induced further deviation, ultimately resulting in changes in the social rules.[27]

To understand the effect of such changes in the social regime on our learning model, it is convenient to consider the case with three types that we discussed in Section 4.3. Suppose that we are in an unstable (Z) community, where only super-reformists deviate, to begin with; $P^*_Z < P < P^*_W$. If deviation

[27]While we focus on interactions among the women in this paper, contraception decisions will also in general depend on the attitudes of the men in the community. Since women bear most of the costs of child-bearing and child-rearing, they are typically at the forefront of efforts to regulate fertility in most traditional societies. It may thus be convenient to think of male attitudes as being aligned with those of the local leaders and the religious establishment. Moreover, since there are few restrictions on male mobility, we might expect to see relatively little variation in male attitudes within the village. In terms of the model, these (common) male attitudes may then be seen to determine, in part, the social rules l and g.

by the super-reformists induces changes in the social rules (lowers P^*_W), then the community could be endogenously transformed into an unstable (WZ) community in which both reformists and super-reformists subsequently deviate. It is important to understand, however, that there is a stochastic element to such regime change; an increase in contraceptive prevalence could trigger a fundamentalist backlash from the religious establishment, or it could lead to a relaxing of religious restrictions. Successive regime changes thus slow down the learning process, since information obtained from past interactions is now no longer as useful as information from recent interactions about the current state of the social regime. Convergence in this case would evidently be slower than in a community that was unstable (WZ) to begin with, and where the social rules did not change over time.

At the individual level, the objective continues to be to learn the type of community at each point in time, since the individual does not account for the effect of her actions on the social regime in a large community. Matching with an m increases the probability that the community is (more) unstable, which increases the likelihood that an individual drawn at random will subsequently choose an m, as before. P^* is fixed initially (at low levels of contraceptive prevalence), and so F will decrease, while $\Delta F(L)$, $\Delta F(R)$ increase, but we conjecture that subsequent changes in the social regime will extend the growth phase that we saw in Figure 7. Ultimately P^* must settle down to its new level, and $\Delta F(L)$, $\Delta F(R)$ will decline as before. We would thus expect changes in P^* to stretch out the learning process, otherwise leaving the individual's decision rule qualitatively unchanged. At the aggregate level, this stretching out of the learning process, together with the infrequent social interactions among the women, might help explain the very gradual change in contraceptive prevalence that is observed in our communities.

While direct evidence of social regime shifts is difficult to obtain, it appears from our conversations with MCH-FP project field-staff that the religious opposition to contraceptive use has weakened in many villages in recent years. Recall also from Figure 2 that Cohort 4, which entered the system at a time when contraceptive prevalence in the intervention area was starting to settle down, converges much faster to the steady state than the cohorts before it (in about five years). While other cohort-based explanations might be available, this observation is consistent with the idea that the youngest cohort was able to learn more quickly about the state of the world because its members were interacting in a relatively stable social regime.

5 Individual Level Empirical Analysis

This section begins with a brief description of the data in Section 5.1. Thereafter we estimate the individual decision rule in Section 5.2, paying particular attention to changes in this rule over time. While the empirical results match well with the predictions of the model, we show in Section 5.3 that they could be spuriously generated by unobserved determinants of the individual's contraception decision that are correlated within the village. We consequently proceed to partition the village by religion in Section 5.4, and by age and education in Section 5.5. Strong within-religion effects are obtained, while cross-religion effects are completely absent. In contrast, strong cross-group effects are consistently obtained when the village is partitioned by age or education. These empirical results, taken together, allow us to rule out a spurious role for neighbors' decisions in Section 5.6, providing additional support for the presence of social interactions.

5.1 The Data

Descriptive statistics for the women in our sample are presented in Table 2. Recall that the sample consists of all women, 15-49 years, married and capable of conceiving, residing in the intervention area at each point in time over the 1983-93 period. We provide these statistics for the full sample, as well as for sub-samples in which the women are divided by religion and by level of education, since these variables will later be used to partition the village in the contraception regression.

Insert Table 2 here.

Starting with individual characteristics in Panel A, we see that the women are on average 29 years old, with 2.4 children. Note that the number of children under-estimates the ultimate family-size since many households will continue to produce children in the future. The women in our sample have roughly two years of education, while their husbands have on average one more year of schooling. Looking across columns in Panel A we see qualitatively similar statistics for Hindus and Muslims, as well as for illiterate versus literate women.

Next we turn to the occupation of the household head in Panel B, using data from a Socioeconomic Census that was conducted by the ICDDR,B in Matlab *thana* in 1982. Starting with the full sample in Column 1, the traditional occupations, farming and fishing, maintain their importance, followed by business, which in this setting refers essentially to petty trade. Looking across columns, the occupational choices are roughly comparable, with one exception: 26% of the Hindus and less than 2%

of the Muslims are fishermen. Comparing Column 1 with Columns 4 and 5 we also see that fishermen are disproportionately illiterate. From our conversations with MCH-FP field-staff it appears that the fishermen tend to be socially conservative and may form a distinct group, which we will take account of later in the estimation section.

Finally we look at asset ownership, also obtained from the 1982 Census, in Panel C. Land, cows and boats are the main assets, and once more we see similar patterns across religious groups and education categories.

While we could reject the hypothesis that the means across religious groups are equal for most of the variables in Table 2, these statistics are generally of comparable magnitude. The two religious groups display qualitatively similar demographic characteristics, occupational patterns, and asset ownership, yet we will later see what appears to be absolutely no interaction, with regard to contraceptive use, within the village.

We complete this section by reporting average contraceptive prevalence, for the full sample as well as for the different groups of women in Panel D. Contraceptive prevalence is roughly 55% over the sample period, and it is about 5 percentage points higher for the Hindus and the literate women, relative to their respective comparison groups (these differences are statistically significant).[28]

5.2 Estimation Results: The Individual Decision Rule

The individual decision rule that we derived in the previous section cannot be taken directly to the data. It is necessary to control for additional determinants of the individual's decision in the contraception regression.

The model assumes implicitly that the individual would choose to use contraceptives in every period, if the social environment was favorable. In practice, women who have not achieved their desired family-size will periodically discontinue using contraceptives to produce children, even when social restrictions are absent. Since the frequency of such temporary discontinuation will be strongly (negatively) correlated with the woman's age, we include age and age squared as additional controls in the contraception regression. Recall that we controlled for such age effects in Figure 1 and Figure 2 as well. The coefficient on the woman's age is positive, the coefficient on age squared is negative,

[28]The number of observations in Panel D is larger than the number of observations in the regressions that we report later because we compute all the statistics in Table 2 over the full 1983-93 sample period. In contrast, we must drop the first year (1983) in the regressions since the lagged decision and lagged contraceptive prevalence were shown to determine the individual's current decision in Section 4. Annual (December 31) data are used to compute these statistics.

and both coefficients are very precisely estimated, in all the regressions that we report in Tables 3-5.

The model assumes that the only change in the intervention area is caused by the introduction of modern contraceptives. But we will see in Section 5.3 that unobserved changes in the MCH-FP project or the economic environment over the sample period could also (spuriously) generate a role for neighbors' past decisions (x_{t-1}) in the contraception regression; neighbors' decisions simply proxy for the unobserved changes in this case. Since contraceptive prevalence information is available across multiple villages, we will control for secular changes by including time effects in all the regressions that we estimate. This is an improvement on the aggregate analysis reported in Section 3 where we could not distinguish between the effect of exogenous changes in the economic environment or the MCH-FP project and the effect of underlying social interactions on contraceptive prevalence in the villages. However, we are still unable to control for village-specific changes over time, and we will return to this point in the next section.

We will also see in Section 5.3 that unobserved individual characteristics could generate a spurious role for neighbors' past decisions in the contraception regression, to the extent that these characteristics are correlated within the village. Since we have panel data, we will check the robustness of all our results by including individual fixed effects in the contraception regression. Fixed effects control for observed time-invariant characteristics, such as religion and education, as well as for the woman's unobserved type; recall from Section 4 that while reformists change their behavior over the transition period, the conformists never use contraceptives.

Once we have settled on an appropriate specification for the contraception regression, the next step in the empirical analysis is to decide on the data frequency for the empirical analysis. We have contraception data at two points in the year, June 30 and December 31. Most of the regressions in this paper use the full sample (six-monthly data), but we will check the robustness of our results with annual data as well. The major disadvantage of using relatively high frequency six-monthly data is that we are left with a narrow belief window in each period, $\Delta F(L)$, $\Delta F(R)$ are small in equation (9), which results in an extremely high level of state dependence. Annual data permits individuals with beliefs in a wider range around the threshold belief to change their behavior from one period to the next, which increases the response to neighbors' decisions but otherwise leaves all the implications of the previous section intact. We will show that the major results of the paper hold up with annual data, and indeed all the results that we report in this paper are qualitatively unchanged when six-monthly data are replaced with annual data.

Turning to the regressions, an augmented version of equation (9) which includes the control variables described above may be written as,

$$y_{it} = \alpha + \gamma_t y_{it-1} + \beta_t x_{t-1} + X_{it}\eta + \xi_{it} \tag{10}$$

where β_t, γ_t are the coefficients in the decision rule whose properties were described in Proposition 3. x_{t-1} is the proportion of eligible women in the village in period $t-1$ who use contraceptives. X_{it} is a vector of control variables which includes the woman's age, age squared, time effects, and (in some cases) individual fixed effects. ξ_{it} is a mean-zero disturbance term. Since the individual decision rule, equation (9), is linear in variables, it is appropriate to use the linear probability model.[29]

While we characterized the changes in β_t, γ_t over time in Section 4, we must now parametrize these coefficients to test the predictions of the model. Since both β_t and γ_t were highly nonlinear in Figure 7, we begin by specifying the learning rule coefficients as quadratic functions of time, t. Starting with the lagged contraceptive prevalence terms, x_{t-1}, $x_{t-1} \cdot t$, $x_{t-1} \cdot t^2$, the transition to the new reproductive equilibrium was clearly under way when our sample begins in 1983, five years after the inception of the MCH-FP project, so we would expect the coefficient on x_{t-1} to be positive from Figure 7. Thereafter, contraceptive prevalence increased relatively steeply from 1983-90 in Figure 1, which would correspond to the growth phase in Figure 5 and Figure 7, so we would expect $x_{t-1} \cdot t$ to be positive as well since β_t was increasing in that period. While contraceptive prevalence continued to increase throughout the 1990s, and reached 70% by 1999 (the most recent year for which official data are available), the trajectory clearly begins to flatten after 1990 in Figure 1. This corresponds to the range with declining β_t in Figure 5 and Figure 7, so we would expect the coefficient on $x_{t-1} \cdot t^2$ to be negative.

We can apply the same sort of argument to predict the sign of the coefficients on the lagged decision terms, y_{it-1}, $y_{it-1} \cdot t$, $y_{it-1} \cdot t^2$. The coefficient on y_{it-1} will be positive, but less than one. The coefficient on $y_{it-1} \cdot t$ will be negative, to pick up the decline in γ_t during the growth phase (1983-89). And finally, the coefficient on $y_{it-1} \cdot t^2$ will be positive, reflecting the increase in γ_t as the community begins to converge to the new equilibrium (1990-93).

[29] An additional advantage of the linear probability model over the more standard discrete choice models in this application is that the lagged dependent variable, individual fixed-effects and time-varying characteristics can be included in the contraception regression (Maddala 1987). As a robustness check we also experimented with a logit model to estimate equation (10). Estimates of the basic regressions in Table 4 (changes in the learning rule over time) and Table 6 (the cross-religion regression) are qualitatively the same with the logit and the linear probability model.

We test these predictions of the model in Column 1 of Table 3. The estimated coefficients of the individual decision rule perfectly match the predictions of the learning model and are very precisely estimated, providing us with the first empirical result of the paper.

Insert Table 3 here.

We showed in Section 4 that an alternative specification of the decision rule would separate out the $x_{t-1} \cdot y_{it-1}$ interaction term. We consequently proceed to experiment with this alternative specification in Column 2 of Table 3. We mentioned that changes in β_t, γ_t over the transition would be qualitatively unaffected by the inclusion of the interaction term, and that the coefficient on this term would be negative. The estimates in Column 2 match perfectly with these predictions. While the coefficient on $x_{t-1} \cdot y_{it-1}$ is insignificant, it will later be precisely estimated in Table 4, where we experiment with an alternative parametrization of β_t, γ_t.

Subsequently, we include individual fixed effects as additional controls in the contraception regression.[30] While the coefficients on x_{t-1}, y_{it-1} both decline, the estimated coefficients in Column 3 continue to be consistent with the predicted changes in β_t, γ_t over time. Later in this section we will see that individual decisions respond strongly to lagged contraceptive prevalence in the individual's own religious group within the village, while cross-religion effects are completely absent. These results suggest that the religious group *within* the village may be a more appropriate social unit than the entire village. Column 4 reports the learning rule with this alternative social unit, and using the results from Column 1 as the basis for comparison we see little change in the estimated coefficients.

Finally, we experiment with annual data in Column 5 of Table 3. Comparing the estimates in Column 5 with the estimates using six-monthly data in Column 1, notice that the coefficient on the lagged decision declines substantially, while there is a corresponding increase in the coefficient on x_{t-1}. This is precisely what we would expect to see since the belief windows, $\Delta F(L)$, $\Delta F(R)$ must be wider when lower frequency data are utilized, which implies a lower level of state dependence and a larger response to neighbors' decisions, and indeed the same patterns will be obtained in the regressions that we report later in the paper. Nevertheless, the changes in β_t, γ_t over time are qualitatively the same with annual and six-monthly data.

[30]It would not make much sense to include fixed-effects and the lagged dependent variable if we did not observe switches, in both directions, between contraceptive use and non-use from one year to the next. Computing sample statistics for all possible transition probabilities over the 1983-93 period, we obtain: non-use to non-use 33%, non-use to use 13%, use to non-use 10%, and use to use 44%.

We complete the estimation of the individual decision rule by considering an alternative parametrization of β_t, γ_t, as spline functions of time, in Table 4. The sample is partitioned into three equal time periods; period 1: 1984-86, period 2: 1986-89, period 3: 1990-93, and β_t, γ_t are now estimated separately in each of these time periods. The 1984-89 period in Figure 1 corresponds to the growth phase in Figure 5 and Figure 7, so we would expect β (γ) to be increasing (decreasing) from period 1 to period 2. Subsequently the contraceptive prevalence trajectory flattened over the 1990-93 period in Figure 1, implying a decline (increase) in β (γ) from Figure 5 and Figure 7. This tells us, in turn, that the β estimate will decline from period 2 to period 3, or at least increase less sharply than it did from period 1 to period 2. Similarly, the γ estimate will increase, or at least flatten out, from period 2 to period 3.

Starting with the basic specification in Column 1 of Table 4, we see that the estimated coefficients coincide with the predictions of the model once again.[31] The difference between the period 1 and period 2 coefficients on x_{t-1}, as well as the difference between the corresponding period 2 and period 3 coefficients, is statistically significant at the 5 percent level, matching the inverted-U pattern for β_t that we obtained in Figure 7. The coefficient on the lagged decision in period 1 is much larger than the corresponding coefficient on x_{t-1}, declines substantially from period 1 to period 2, and then remains stable over the last two periods, consistent once more with the broad predictions of the model.

Insert Table 4 here.

Subsequently we repeat all the regressions that we reported in Table 3. We include the $x_{t-1} \cdot y_{it-1}$ interaction as an additional regressor in Column 2, individual fixed effects are included as control variables in Column 3, the religious group within the village is treated as the social unit in Column 4, and finally we experiment with annual data in Column 5. The basic pattern that we saw in Column 1 continues to be obtained, across all the alternative specifications in Table 4.[32]

We close this section by describing a number of robustness checks that we conducted, to ensure that

[31] We follow the convention that the time controls in the contraception regression should match the specification of the interaction terms. Accordingly, when x_{t-1}, y_{it-1} were interacted with t, t^2 in Table 3, t and t^2 were included as time controls. Similarly, time-period dummies are included as additional controls when x_{t-1}, y_{it-1} are interacted with the three time-period dummies in Table 4.

[32] One explanation for the decline in the point estimates when fixed effects are included is that we are controlling for unobserved heterogeneity, which may have previously biased the coefficient on the lagged decision as well as the response to neighbors' decisions. However, it is also well known that within-estimation to eliminate fixed-effects gives rise to inconsistent estimates when the lagged dependent variable is included as a regressor. Correlation of the order $(1/T)$, where T is the number of time periods, is created between the lagged dependent variable and the residuals in the transformed model (Hsiao 1986). Since $T = 6$ in our data, this bias could be significant.

the empirical results are robust to alternative samples, specification of the contraception regression, and construction of the contraception variable.

The individual decision rule in Section 4 was derived for reformists. The regression results reported up to this point, however, do not distinguish between individuals of different types. Table 5, Column 1 thus discards all individuals who never use contraceptives from the sample.[33] Estimates with this reduced sample are similar to what we obtained earlier in Table 4, Column 1.

Insert Table 5 here.

The model in Section 4 derived changes in contraceptive prevalence for a single cohort of individuals. However, since individuals move in and out of the sample, there are actually multiple cohorts in the data set. We saw in Figure 2 that Cohort 2 and Cohort 3, which account for 75% of the observations in the data, display aggregate patterns of contraceptive prevalence that are similar to the full sample. We consequently restrict attention to each of these cohorts in Columns 2-3 of Table 5. Contraceptive prevalence x_{t-1} is now computed within each cohort, to match the set up of the model. While the point estimates are slightly different from what we obtained with the full sample in Table 4, Column 1, the basic pattern in the coefficients over time is unchanged.[34]

The changes in the decision rule over time were derived in Section 4 for a particular (unstable) community in a simple model with two types. We mentioned in that section that the basic changes in β_t, γ_t would be obtained in a more general model with many types, for any community in which contraceptive prevalence increased steadily over time. To verify the robustness of our results, we consequently proceed to estimate the individual decision rule separately for villages in the top and the bottom quartile of the 1993 contraceptive prevalence distribution in Columns 4-5 of Table 5. The pattern observed with the full sample of villages in Table 4, Column 1 continues to be obtained with each restricted group of villages.

The learning rule derived in Section 4, and estimated in this section, includes a single lag for both individual decisions (y_{it-1}) and neighbors' decisions (x_{t-1}). But an alternative model without social interactions could well have included additional lags on the individual's own decision. A spurious social effect would then be obtained if neighbors' decisions proxy for these longer lags.

[33] 11% of the individuals in the data never use contraceptives, and the corresponding statistics for villages in the top and the bottom quartile of the 1993 contraceptive prevalence distribution are 10% and 14% respectively.

[34] Similar results are obtained when contraceptive prevalence is computed at the level of the village, but separate regressions are estimated for each cohort.

We check the robustness of our results to these longer lags by including the twice-lagged individual decision in Table 5, Column 6. The coefficients on the (once) lagged decision and lagged contraceptive prevalence in the village are hardly affected by the inclusion of these additional variables. And although the coefficient on the twice-lagged decision is precisely estimated, it is small in magnitude (roughly 0.02) with no discernable pattern over the three time periods.

And finally, the robustness tests conclude with an alternative construction of the contraception variable, to allow for the possibility that women could temporarily discontinue using contraceptives to have a child, in Table 5, Column 7. Up to this point in the empirical analysis we have implicitly assumed that any switch from use to non-use reflects a change in the individual's belief. There will clearly be some periods of non-use in which the individual will be trying to have a child, even though her beliefs are to the right of the threshold belief at that time. Most women will conceive within a year, and so to account for such spurious episodes of non-use we now treat the individual as an adopter in a given period if she used contraceptives at any point over the past 1.5 years (with six-monthly data we are aggregating over four periods). The lagged decision and lagged contraceptive prevalence in the village are re-computed using this alternative contraception variable.[35] We are now ignoring many real switches from use to non-use, particularly in the early years of the transition, and not surprisingly the point estimates do change in Column 7. But the basic pattern of learning weights over time remains unchanged with the alternative construction of the contraception variable.

5.3 The Identification Problem

In the previous section we briefly mentioned the problems for consistent estimation that could arise when neighbors' past decisions proxy for unobserved determinants of the individual's contraception decision. We now proceed to study this identification problem more formally.

Begin with a modified version of the individual decision rule, equation (10), that includes unobserved determinants of the contraception decision, but treats β_t, γ_t as being stable over time.

$$y_{it} = \alpha + \gamma y_{it-1} + \beta E_{t-1,v}(y_{it-1}) + X_{it}\eta + C_t^v + \xi_{it} \tag{11}$$

where $E_{t-1,v}(y_{it-1})$ is the expected level of contraceptive prevalence in the village in period $t-1$, which corresponds to x_{t-1} in the model. C_t^v is any unobserved determinant of the contraception decision that varies across villages and over time. In this setting C_t^v reflects unobserved individual

[35] We are grateful to an anonymous referee for suggesting this alternative estimation procedure.

characteristics that are correlated within the village, as well as changes in the economic environment or the MCH-FP project that are village specific.

To see the identification problem that arises in this case, lag equation (11) by one period and take expectations across individuals within the village,

$$E_{t-1,v}(y_{it-1}) = \alpha + (\gamma + \beta)E_{t-2,v}(y_{it-2}) + E_{t-1,v}(X_{it-1})\eta + C^v_{t-1}. \quad (12)$$

It is easy to verify from equation (11) and equation (12) that $E_{t-1,v}(y_{it-1})$ will proxy for C^v_t when the omitted variable is auto-correlated. In general, a role for neighbors' decisions cannot be identified when the omitted variable is correlated across individuals in the village and over time. We avoid this problem to some extent by including individual characteristics in X_{it}; the fixed-effects in particular will control for any time-invariant determinants of the contraception decision. Time effects will also control for secular changes that are common to all the villages. But we cannot control for unobserved village-specific variation over time.[36]

While we assumed that the decision rule was stable in equation (11), our model actually placed restrictions on changes in this rule over time, which we successfully verified in the previous section. These results are very useful, providing strong support for the presence of an underlying learning process, but they do not by themselves conclusively establish a role for neighbors' decisions in the contraception regression.[37]

Our solution in this paper, to provide additional support for the presence of social interactions, takes advantage of the institutional background that we presented in Section 2. We noted in that section that the very infrequent female social interactions in these villages will almost never cross religious boundaries. Thus, we would expect individuals to respond to contraceptive prevalence within their own religious group within the village, while cross-religion effects should be completely absent. These restrictions across religious groups will turn out to be very useful in ruling out potential omitted variables that could have generated a spurious role for neighbors' past decisions in the contraception

[36] The health component of the MCH-FP project was very successful, with child mortality rates (per thousand live births) declining from 170 deaths in 1983 to 85 deaths in 1993. We included village-specific child mortality in the contraception regressions in Tables 3-4, to check the robustness of our results. While these results are not reported here, the coefficient on mortality is negative and significant as expected, but the other coefficients in the regression are completely unchanged. While this test allows us to control for village-specific change along one dimension, we cannot control for other unobserved determinants of contraceptive behavior that could vary across villages and over time.

[37] Brock and Durlauf (2001) show that the nonlinearity inherent in the discrete choice model improves prospects for the identification of social interactions in general. In our case, the individual decision rule is characterized by a particular nonlinearity, which is based on the discrete nature of the contraception decision, the matching process, and the Bayesian updating.

regression.

To test these additional restrictions, we will later estimate separate contraception regressions for each religion, allowing for within-religion and cross-religion effects. Muslim (M) individuals are identified by an i subscript, with a j subscript for Hindus (H). $E_{t-1,Mv}(y_{it-1})$ represents the average adoption among the eligible Muslims in village v in period $t-1$. $E_{t-1,Hv}(y_{jt-1})$ represents the corresponding statistic for Hindus in the village. Note that we now allow for religion-specific omitted variables C_t^{Mv}, C_t^{Hv} in the contraception regressions.

$$y_{it} = \alpha_M + \gamma_M y_{it-1} + \beta_{MM} E_{t-1,Mv}(y_{it-1}) + \beta_{MH} E_{t-1,Hv}(y_{jt-1}) + X_{it}\eta_M + C_t^{Mv} + \xi_{it}. \tag{13}$$

$$y_{jt} = \alpha_H + \gamma_H y_{jt-1} + \beta_{HM} E_{t-1,Mv}(y_{it-1}) + \beta_{HH} E_{t-1,Hv}(y_{jt-1}) + X_{jt}\eta_H + C_t^{Hv} + \xi_{jt}. \tag{14}$$

Following the discussion above, we would expect that $\beta_{MM} > 0$, $\beta_{HH} > 0$, $\beta_{MH} = \beta_{HM} = 0$. We will see later that the estimation results precisely match these predictions, providing strong support for the view that social interactions, within each religious group, gave rise to changes in contraceptive use. As noted, these results are also useful in ruling out the possibility that neighbors' decisions simply proxy for unobserved determinants of the contraception decision C_t^{Mv}, C_t^{Hv}. In the discussion that follows we derive conditions that these omitted variables must satisfy to spuriously generate the observed within-religion and cross-religion effects. Later in Section 5.6 we will argue that the omitted variables that we would expect to encounter in this setting are unlikely to satisfy these conditions.

Consider an alternative model *without* social interactions:

$$y_{it} = \alpha_M + \gamma_M y_{it-1} + X_{it}\eta_M + C_t^{Mv} + \xi_{it} \tag{15}$$

$$y_{jt} = \alpha_H + \gamma_H y_{jt-1} + X_{jt}\eta_H + C_t^{Hv} + \xi_{jt}. \tag{16}$$

$E_{t-1,Mv}(y_{it-1})$, $E_{t-1,Hv}(y_{jt-1})$ no longer belong in the contraception regressions, but could the $E_{t-1,Mv}(y_{it-1})$ term in equation (13) and the $E_{t-1,Hv}(y_{jt-1})$ term in equation (14) simply proxy for the unobserved C_t^{Mv}, C_t^{Hv} terms in this case? They could, for exactly the same reason that we described above for the village-level regression. Taking expectations in equation (15) and lagging one period, $E_{t-1,Mv}(y_{it-1})$ is correlated with C_{t-1}^{Mv}. Similarly, $E_{t-1,Hv}(y_{jt-1})$ would be correlated with C_{t-1}^{Hv}. So spurious $\hat{\beta}_{MM}$, $\hat{\beta}_{HH}$ estimates could be obtained if C_t^{Mv}, C_t^{Hv} are auto-correlated.

But what about the cross-religion effects? $E_{t-1,Mv}(y_{it-1})$ cannot perfectly proxy for C_{t-1}^{Mv} on account of the $E_{t-2,Mv}(y_{it-2})$ and $E_{t-1,Mv}(X_{it-1})$ terms that correspond to y_{it-1} and X_{it} in equation (15), once we have lagged that equation by one period and taken expectations. This leaves room for $E_{t-1,Hv}(y_{jt-1})$ to appear as an additional proxy for C_{t-1}^{Mv}. Cross-religion effects in the Muslim regression can only be absent if $E_{t-1,Hv}(y_{jt-1})$ provides no information about C_{t-1}^{Mv}. A necessary condition for this result is that C_t^{Mv} and C_t^{Hv} should be orthogonal. Similarly, if $E_{t-1,Mv}(y_{it-1})$ has no role to play in the Hindu regression, then C_t^{Mv}, C_t^{Hv} must be orthogonal.

To explain the estimated pattern of weights $\hat{\beta}_{MM} > 0$, $\hat{\beta}_{HH} > 0$, $\hat{\beta}_{MH} = 0$, $\hat{\beta}_{HM} = 0$ without social effects, the omitted variables C_t^{Mv}, C_t^{Hv} *must* be orthogonal within the village. We will argue later in Section 5.6 that it is difficult to imagine that potential omitted variables such as program effects or economic change are completely uncorrelated across religious groups within the same village. As long as the omitted variable is not female-specific, as with the social interactions, there is no reason to expect that its effect would not cross religious boundaries within the village.

5.4 Estimation Results: Partitioning the Village by Religion

Contraceptive prevalence, separately for Hindus and Muslims, together with the individual's lagged decision are now included as determinants of the contraception decision. The contraception regression will be estimated separately for Hindus and Muslims, in most of the specifications that we consider in this section. The regressions in Tables 3-5 allowed us to study changes in the learning rule over time. The empirical analysis that follows partitions the sample by religion, allowing us to study within-religion and cross-religion effects in the village. Since we are no longer concerned with changes in the learning rule over time, the coefficients on the lagged decision and lagged contraceptive prevalence are assumed to be stable over the sample period.[38] Age, age squared, and time-period dummies are included as control variables as usual. The coefficient on the individual's age is positive, the coefficient on age squared is negative, and both these coefficients are very precisely estimated, in all the cross-religion regressions.

The first regression in Columns 1-2 of Table 6 considers all mixed religion villages, and we see that strong within-religion effects are obtained while cross-religion effects are completely absent, both for Hindus and Muslims. While these results are very promising, one cause for concern is that villages

[38]We did, however, verify that the basic result in Table 6: strong within-religion effects and absent cross-religion effects, for Hindus and Muslims, continue to be obtained when we further divide the sample into three time periods as in Tables 4-5.

may be predominantly of one religion or the other. In the extreme case, all the within-religion effects could be obtained from villages that consist exclusively of households belonging to a particular religion, which would leave no room at all for cross-religion effects. Although we do not see this sort of extreme segregation in the data, some villages are dominated by a single religion. We consequently proceed to remove all villages with less than 5% Hindus or Muslims from the sample in Columns 3-4. Subsequently we discard villages with less than 15% Hindus or Muslims in Columns 5-6. The sample size declines substantially over the course of this exercise, and is less than half the size of what we began with. Yet we see that the estimated within-religion and cross-religion effects, for both Hindus and Muslims, remain remarkably stable across the different sample sizes in Table 6.[39] Following the suggestion of a referee, we will restrict the sample to villages with more than 5% Hindus and Muslims in all the cross-religion regressions that follow.

Insert Table 6 here.

We complete Table 6 by running the contraception regression with annual data in Columns 7-8. Comparing the estimates using annual data with the corresponding estimates using six-monthly data in Columns 3-4 we see that the coefficient on the lagged decision declines substantially, while there is a corresponding increase in the within-religion effect, matching the difference between the six-monthly and the annual regressions that we reported earlier in Section 5.2. Nevertheless, the usual pattern of strong within-religion effects and completely absent cross-religion effects continues to be obtained for both Hindus and Muslims.

Subsequently we check the robustness of the results by replacing the three time-period dummies (1984-86, 1987-89, 1990-93) that we used to control for time effects in Tables 4-6 with a full set of bi-annual (six-monthly) dummies in Table 7, Columns 1-2. The estimated coefficients are remarkably robust when compared with the corresponding estimates from the basic specification in Table 6, Columns 3-4. In a related robustness check we include individual fixed effects in the contraception regression in Table 7, Columns 3-4. As in Table 3 and Table 4, all the coefficients decline in size, but the basic within-religion and cross-religion patterns continue to be obtained.

Insert Table 7 here.

[39]In a related robustness test, we also verified that the size of the village, measured by the total number of eligible women, has no effect on the estimated within-religion and cross-religion effects. We also verified that the within-religion and cross-religion results were robust to the tests reported in Table 5; the inclusion of the twice-lagged dependent variable, and the alternative construction of the contraception variable to account for birth-spacing.

Continuing with these robustness tests, while the village is treated as the social unit for much of the analysis in this paper, it may be less important as a social institution in Bangladesh, as compared with the rest of South Asia (Arthur and McNicoll, 1978). On average ten households form one *bari*, and since all the households in a *bari* share the same religion, within-religion effects could simply proxy for underlying *bari*-specific social effects. While these would be social interactions of a sort, they are not necessarily related to the community-level religious restrictions that we are interested in. We consequently include *bari*-level adoption as an additional regressor in the contraception equation. While positive and statistically significant *bari*-effects are obtained in Columns 5-6 of Table 7, the estimated within-religion and cross-religion effects remain very stable when compared with Columns 3-4 of Table 6.

Further, we noted in Table 2 that fishermen, who display relatively low levels of education and tend to be socially conservative, are disproportionately Hindu. One possibility in this case is that the observed within-religion and cross-religion effects simply proxy for an underlying occupation effect. We have already accounted for any characteristics that do not vary over time in the fixed effects regressions, but we nevertheless include the husband's occupation in Columns 7-8 as an additional control. Comparing the coefficients in Columns 7-8 with the coefficients in Table 6, Columns 3-4 we see that the estimates are very robust to the inclusion of the husband's occupation for both Hindus and Muslims.[40]

Finally, suppose that social interactions are absent, but that the Community Health Worker (CHW) can only induce members of her own religious group to change their behavior over time.[41] Changes in within-religion contraceptive prevalence will then proxy for changes in the CHW's persuasive ability, but only among women who share her religion. Within-religion effects will be absent among the women in the village who do not share the CHW's religion. Further, cross-religion effects will be absent for all women. CHW's are typically drawn from the dominant group in the village, so most women in the sample will share the same religion as the local CHW. The unobserved CHW-effect

[40]The husband's occupation includes a set of four dummies- farming, fishing, housework, and business - with "other occupations" included as the base category in the contraception regression. Housework and business are positive and significant for both religions. Farming is positive and significant while fishing is negative and significant for the Hindus. The coefficients on farming and fishing are imprecisely estimated for the Muslims. In a previous version of the paper we experimented with an alternative test that dropped fishermen from the sample. The basic pattern of within-religion and cross-religion effects continued to be obtained with this restricted sample for both religions.

[41]The assumption that the CHW completely ignores the members of the other religious group, or that they completely ignore her, is very strong. Later in Section 5.6 we will argue that it is difficult to imagine that she would provide orthogonal inputs to the two religious groups within the village or, more generally, that she would have an orthogonal impact on the two groups.

that we have just described could, in principle, have generated both the within-religion effects, as well as the absence of cross-religion effects, that we see in the data.

To rule out this possibility, individuals from both religions are pooled together in Columns 9-10 of Table 7. Column 9 restricts attention to individuals that share the same religion as the CHW and Column 10 restricts attention to individuals who do not share her religion. Strong within-religion effects continue to be obtained with both groups of individuals, ruling out the alternative CHW-based explanation, while the cross-religion effects are absent as usual.

5.5 Estimation Results: Alternative Partitions of the Village

We complete the empirical analysis by partitioning the village, separately by age category and by level of female education, in Table 8. What within-group and cross-group effects would we expect to see with these alternative partitions of the village? There are two independent mechanisms that generate particular within-group and cross-group patterns in this case.

First, while we assume for simplicity that individuals match randomly within the village, peer groups tend to consist of individuals with similar characteristics in practice. We would thus expect individuals to place more weight on their own group in the contraception regression, everything else being equal.

Second, in our Bayesian setting the weight that an individual places on her partner's decision in each period will depend on the amount of information that she receives, from that interaction, about the nature of the social equilibrium that will ultimately prevail. Partition the community into G groups with different observed characteristics, and apply Bayes' Rule to derive the posterior belief when any individual with belief α matches with an m from some group g:

$$Pr(P = \overline{P} \mid m_g) = \frac{1}{1 + \left(\frac{1-\alpha}{\alpha}\right) \underline{x}_{tg}/\overline{x}_{tg}}$$

where $\underline{x}_{tg}$, $\overline{x}_{tg}$ refers to the proportion of m's among the members of the g group in stable and unstable communities respectively. Clearly a lower $\underline{x}_{tg}/\overline{x}_{tg}$ implies a stronger response to the m-match. In general, a group whose behavior is less sensitive to the type of community will wield less social influence. In the extreme case, if $\underline{x}_{tg} = \overline{x}_{tg}$, then it is easy to verify from the expression above that the posterior belief will be equal to the prior belief, no information will be provided through the social interaction, and the members of that group will carry zero weight.

Earlier in Section 4 we saw that super-reformists chose the m-action in both unstable (Z) and unstable (WZ) communities, while the reformists chose the m-action in unstable (WZ) communities only. While these types are private information in our model, we would expect such patterns of behavior to hold for groups with different observed characteristics as well; groups with a greater propensity to choose the m-action (with a lower P^* in the model), should in general be less sensitive to the type of community they belong to. Since older women and more educated women are more likely to use contraceptives at each point in time over the sample period, our simple model predicts that they should have *less* influence in the contraception regression, everything else being equal.

As usual within-group and cross-group contraceptive prevalence, and the lagged decision, are included as regressors in the contraception equation. Control variables include the woman's age, age squared, and time-period dummies.[42] The village is divided into two age categories in Columns 1-2 of Table 8, using the median age in the sample (30 years) as the cut off. Young women typically use contraceptives for birth-spacing, so there will be switches in the early years between use and non-use, even for individuals that have crossed the belief threshold. The level of state dependence, measured by the coefficient on the lagged decision, is consequently increasing with age. Turning to the social interactions, while both young and old women put more weight on their own group, cross-group effects are substantial and statistically significant.

Insert Table 8 here.

To further explore these results, we partition the village using the 0.75 age quantile (37 years), computed using the full sample of women, as the cut off. Essentially the same pattern that we saw above continues to be obtained in Columns 3-4. As we noted in Section 2, social interactions tend to occur within peer groups of the same age, which would explain the dominance of the within-group effect in Columns 1-4.

Turning to the second partitioning variable, the village is divided into literate and illiterate women in Columns 5-6 of Table 8. The coefficient on the lagged decision is now roughly the same, for both groups of women. Within-group effects are larger while cross-group effects are substantial and statistically significant, both for literate and illiterate women, just as we saw above when the village was partitioned by age. But notice now, that the literate women place relatively more weight on the

[42] Since we are now partitioning the sample by age (and education is highly correlated with age), the age effects in Table 8 are difficult to interpret. While we observe the usual pattern with a positive coefficient on age and a negative coefficient on age squared in some cases, the opposite pattern is also obtained.

other group in Column 6. This pattern is accentuated when we partition the village using the 0.75 education quantile (four years of schooling), computed using the full sample of women, as the cut off. Both groups of women put more weight on the less educated women in Column 7 and Column 8. This result is consistent with the view that social interactions occur freely across these groups, and more importantly that less educated women (who have a lower propensity to use contraceptives) provide more information through their social interactions. We are used to thinking of educated women as being more influential in the community, presumably due to their higher social status.[43] The results that we have just described, go in the opposite direction to this view but are perfectly consistent with the predictions of the model, providing additional independent support for our learning-based theory of social change.

The results that we have presented in Table 8 are very different from what we saw earlier, when the village was partitioned by religion. While within-group effects tend to be larger than cross-group effects, these cross-group effects are nevertheless substantial and statistically significant in the Table. In sharp contrast, while strong within-religion effects were also obtained, cross-religion effects were completely absent across all the specifications that we experimented with in Table 6 and Table 7.

5.6 Alternative Explanations for the Estimation Results

We conclude the empirical analysis by discussing alternative explanations for the results described in this section. We will discuss unobserved program effects, economic change, and learning about new contraceptive technology below.

1. Program Effects: Cross-sectional variation in the MCH-FP project is captured by the individual fixed-effects in the contraception regression. Secular changes in the program are accounted for by the time-period dummies. However, changes across villages and over time could generate a spurious role for neighbors' decisions in the contraception regression. For example, the Community Health Worker (CHW) may become more effective in persuading individuals to adopt contraceptives with experience. If CHWs improve at different speeds, then neighbors' decisions could simply proxy for an unobserved CHW-effect.

To rule out the CHW-effects, and program-effects more generally, we appeal to the within- religion

[43] Educated women could also be more influential if they had access to superior information, through the media for instance. While this argument would make sense if the women were learning about a new contraceptive technology, it does not apply when the community is learning about local social fundamentals, since there is no reason why educated women should have better information about the social equilibrium that will ultimately prevail.

and cross-religion results. Recall that omitted variables must be orthogonal across religious groups within the village to spuriously generate the patterns that we saw in the data. While project-inputs and information-signals may have varied across religious groups within the village, it is difficult to imagine that they were uncorrelated across these groups. Remember that it is the same agency, and the same CHW, that is providing these inputs. The absence of cross-religion social effects thus rules out an important potential source of bias. Note that this result does not imply that the MCH-FP project was ineffective in bringing about change. What this says is that there was little cross-village variation in project inputs. The program-effects would then be picked up by the time effects in the contraception regression.

2. Economic Change: It is well known that fertility rates have historically responded to changes in income and returns to human capital investment (Schultz 1988, Rosenzweig 1990). There has, however, been little economic change in the intervention area over time. Matlab *thana* was chosen as the setting for the early cholera-vaccine trials in the 1960s by the ICDDR,B precisely because it was so isolated. We saw in Table 2 that traditional agriculture and fishing continue to be the main occupations, and no new technologies for these activities have been introduced. Apart from fishing, there does not appear to be a religious aspect to occupational choice among the men, and we controlled for the husband's occupation in Table 7, Columns 7-8, in any case. As discussed in Section 2, women, from both religions rarely work outside the home in rural Bangladesh.

Moreover, local segregation by occupational activity, along religious lines, is not observed in the data. The choice of occupational activities is positively correlated across religions *within* the village. Computing the share of each occupational activity by religion within the village, the correlation coefficient across religions is found to be 0.48 for farming, 0.66 for fishing and 0.37 for business. Thus even if economic change did occur, it is very likely to have been correlated across religious groups within the village. While female social interactions may not cross religious boundaries, there are no such restrictions on male interactions; Hindu and Muslim men mix freely in the marketplace and in other public areas. We can therefore use our cross-religion result once more to rule out economic change as a source of spurious correlation.

3. Learning about Contraceptive Technology: The model of social change that we present in this paper assumes that both preferences (P) and social rules (l, g) are exogenous and stable. However, within this framework, uncertainty about the performance of the new contraceptive technology (measured by S in the model) could still generate many of the patterns that we see in the data.

The MCH-FP project introduced modern contraceptives in the intervention area for the first time. While the main contraceptive method promoted by the ICDDR, B - injectables - appears to have been a well understood and well established technology by the time our sample begins in 1983, there was some initial uncertainty as to whether a woman who had previously used injectables for birth-spacing would be able to conceive in the future.

Both neighbors' decisions, as well as their experiences, will typically provide information about a new technology. For instance, a neighbor's (unexpected) decision to use a particular contraceptive method reveals that she must have received a favorable signal about its performance. This signal extraction process would generate a link between the individual's decision and her neighbors' past decisions (as in Munshi 2001). Uncertainty as to whether the new injectible technology was reversible would also have begun to be resolved when women that used injectables for birth-spacing discontinued and subsequently conceived.

We have already argued that external information signals and program inputs will be correlated across religious groups within the village, so the signal extraction process described above cannot explain the complete absence of cross-religion effects that we see in the data. However, information generated within the community, through individuals' experiences with the new technology, could give rise to these patterns if information generated internally fails to cross religious boundaries within the village.

While there is a wealth of anecdotal evidence suggesting that religious restrictions led to substantial delays in the adoption of contraceptives in the intervention area, we are unaware of any study that points to persistent concerns about the performance of the new technology as an explanation for these delays. In general, we were unable to uncover independent supplemental evidence supporting the alternative view that individuals were learning about a new technology. For instance, it is possible to identify women in the data who used injectables for birth-spacing and then subsequently conceived. We found that the number of post-injectable births in the village, or within each religious group in the village, in any period had absolutely no effect on subsequent contraceptive use in the community during the sample period. This absence of an experience effect contrasts with the results from our own previous research on the adoption of new agricultural technology (Munshi 2001), where we find that neighbors' experiences (crop yields) had a strong effect on the individual grower's subsequent acreage decision in the Indian Green Revolution.

6 Conclusion

Our objective in this paper is to explain two common features of the development experience: the slow response to external interventions, and the wide variation in the response to the same intervention.

One explanation for these stylized facts is based on the idea that individual behavior is often socially regulated in a traditional economy. When the economic environment changes, individuals may be uncertain about the new equilibrium that will emerge in their community. Limited initial deviation from the traditional equilibrium induces further deviation, as individuals interact sequentially with each other, allowing the community to gradually find its way to the new equilibrium.

We apply this view of the development process to the demographic transition in rural Bangladesh. We find that contraceptive prevalence changes slowly, and that groups of villages converge to very different long-run equilibria. At the individual level, changes in the estimated decision rule over time match well with the predictions of our learning-based theory of social change. To provide additional support for this view, we partition the village by religion. Individuals respond strongly to contraceptive prevalence within their own religious group within the village, while cross-religion effects are completely absent, consistent with the idea that social interactions among the women in each religious group gradually led to changes in reproductive behavior in these communities.

The MCH-FP project is quite possibly the most intensive family planning program ever put in place. Community Health Workers meet each woman at her home once every two weeks, in an attempt to circumvent the restrictions on female mobility associated with *purdah* in rural Bangladesh. But we see at the end that there appears to be no substitute for the social interactions among the women. And since these are very infrequent, contraceptive prevalence ultimately changed very slowly in the intervention area. With hindsight, a program that encouraged women to meet at the primary health clinic, instead of delivering services to their homes, might have been more effective despite the initial resistance and delays in adoption that would almost certainly have occurred.

The general point that we hope the reader will take away from this paper is that individual decisions are not always independent choices in a traditional society. The inflexibility that comes from the interaction between individual behavior and social institutions may in fact explain a large part of the inertia and the weak response to new opportunities that has been a pervasive feature of the development experience. Development interventions that pay attention to the underlying social institutions would in that event be significantly more effective in bringing about change than their predecessors.

7 Appendix: Proof of Claim 1

We make the usual rational expectations assumption that each individual correctly predicts the proportion of m's in both stable communities $\underline{x}_t$ and unstable communities $\overline{x}_t$, at each point in time during the transition, but is unsure about the type of community that she belongs to. We place no other restrictions on $\overline{x}_t$, $\underline{x}_t$, and the shifts in beliefs will be derived below for both $\overline{x}_t > \underline{x}_t$ and $\overline{x}_t < \underline{x}_t$.

Case 1: $\overline{x}_t > \underline{x}_t$

Consider any belief α in the support of the distribution. Applying Bayes' Rule, define a neighborhood $[\alpha(L), \alpha(R)]$ around α, which is relevant to determining $\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha)$, $\underline{F}_{t+1}(\alpha) - \underline{F}_t(\alpha)$. Begin by deriving an expression for $\alpha(L)$, which lies to the left of α. Since there are more m's in an unstable community at each point in time, $\overline{x}_t > \underline{x}_t$, matching with an m will shift beliefs to the right. Thus an individual with this belief in period t who matches with an m just reaches α in the subsequent period

$$\alpha = Pr(P = \overline{P} \mid m) = \frac{\alpha(L)\overline{x}_t}{\alpha(L)\overline{x}_t + (1-\alpha(L))\underline{x}_t}.$$

Solving for $\alpha(L)$ we obtain:

$$\alpha(L) = \frac{\alpha\underline{x}_t}{(1-\alpha)\overline{x}_t + \alpha\underline{x}_t} = \frac{\alpha\underline{x}_t}{X_t}$$

where X_t is a weighted average of $\overline{x}_t$ and $\underline{x}_t$, with the weight α on $\underline{x}_t$. Since $\overline{x}_t > \underline{x}_t$, it is easy to verify that $\alpha(L) < \alpha$. From the preceding expression we obtain:

$$\alpha - \alpha(L) = \frac{\alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)}{X_t}. \tag{17}$$

A similar exercise allows us to solve for $\alpha(R)$. An individual with belief $\alpha(R)$ in period t who matches with a t just reaches α in the subsequent period. To derive an expression for $\alpha(R)$, we apply Bayes' Rule as before:

$$\alpha = Pr(P = \overline{P} \mid t) = \frac{\alpha(R)(1-\overline{x}_t)}{\alpha(R)(1-\overline{x}_t) + (1-\alpha(R))(1-\underline{x}_t)}.$$

Solving for $\alpha(R)$ as above, and then for $\alpha(R) - \alpha$, we finally obtain:

$$\alpha(R) - \alpha = \frac{\alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)}{1 - X_t}. \tag{18}$$

The flow to the right of α is determined by the measure of individuals in $[\alpha(L), \alpha]$ who match with an m in period t. Similarly the flow to the left of α is determined by individuals in $[\alpha, \alpha(R)]$ who match with a t. Deriving the net flows is relatively straightforward in this case because the range of beliefs $\alpha - \alpha(L)$ and $\alpha(R) - \alpha$ is extremely narrow: a *single* m (or t) will shift beliefs from $\alpha(L)$ (or $\alpha(R)$) to α. We can therefore assume that the distribution of beliefs is approximately uniform in $[\alpha(L), \alpha(R)]$.Starting with the unstable communities, the net flow to the right is given by:

$$\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha) = -[\overline{f}_t(\alpha)(\alpha - \alpha(L))]\overline{x}_t + [\overline{f}_t(\alpha)(\alpha(R) - \alpha)](1 - \overline{x}_t) \tag{19}$$

where $\overline{f}_t(\alpha)$ is the density of the distribution at α. Substituting expressions for $\alpha - \alpha(L)$ from equation (17) and $\alpha(R) - \alpha$ from equation (18), equation (19) can thus be simplified as

$$\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha) = -\overline{f}_t(\alpha) \cdot \frac{\alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)}{X_t} \cdot \overline{x}_t + \overline{f}_t(\alpha) \cdot \frac{\alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)}{1 - X_t} \cdot (1 - \overline{x}_t). \tag{20}$$

Collecting terms and simplifying equation (20),

$$\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha) = \frac{\overline{f}_t(\alpha) \cdot \alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)(X_t - \overline{x}_t)}{X_t(1-X_t)}.$$

Recall that X_t is a weighted average of $\overline{x}_t$ and $\underline{x}_t$. Since $\overline{x}_t > \underline{x}_t$, $\overline{F}_{t+1}(\alpha) - \overline{F}_t(\alpha) < 0$.

Turning to the stable communities, note that the expressions for $\alpha - \alpha(L)$ and $\alpha(R) - \alpha$ are unchanged. The expression corresponding to equation (20) is therefore

$$\underline{F}_{t+1}(\alpha) - \underline{F}_t(\alpha) = -\underline{f}_t(\alpha) \cdot \frac{\alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)}{X_t} \cdot \underline{x}_t + \underline{f}_t(\alpha) \cdot \frac{\alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)}{1-X_t} \cdot (1 - \underline{x}_t).$$

Simplifying as before,

$$\underline{F}_{t+1}(\alpha) - \underline{F}_t(\alpha) = \frac{\underline{f}_t(\alpha) \cdot \alpha(1-\alpha)(\overline{x}_t - \underline{x}_t)(X_t - \underline{x}_t)}{X_t(1-X_t)}.$$

Since $\overline{x}_t > \underline{x}_t$, $X_t - \underline{x}_t > 0$ and, hence, $\underline{F}_{t+1}(\alpha) - \underline{F}_t(\alpha) > 0$.

Case 2: $\overline{x}_t < \underline{x}_t$

This case is very similar to the previous case except that matching with an m will now shift beliefs to the left, since there are fewer m's in an unstable community at each point in time, $\overline{x}_t < \underline{x}_t$. Thus, to derive the expression for $\alpha(L) < \alpha$, the individual must match with a t to reach α. Similarly, an individual with belief $\alpha(R) > \alpha$ who matches with an m just reaches α. Once we have derived these beliefs, it is straightforward to derive the new flow; note that matching with a t now shifts beliefs to the right.

References

[1] Amin, Sajeda (1997). "The Poverty-Purdah Trap in Rural Bangladesh: Implications for Women's Roles in the Family," *Development and Change*, Vol.28.

[2] Amin, Sajeda, Ian Diamond and Fiona Steele (1997). "Contraception and Religiosity in Bangladesh," in *The Continuing Demographic Transition*, eds. G.W. Jones, R.M. Douglas, J.C. Caldwell and R.M. D'Souza, Clarendon Press, Oxford.

[3] Arthur, Brian W. and Geoffrey McNicoll (1978). "An Analytical Survey of Population and Development in Bangladesh," *Population and Development Review*, Vol. 4, No.1.

[4] Banerjee, Abhijit V. (1992). "A Simple Model of Herd Behavior," *The Quarterly Journal of Economics*, Vol.117, Issue 3.

[5] Banerjee, Abhijit V. (1993). "The Economics of Rumours" *Review of Economic Studies*, Vol.60.

[6] Bikhchandani, Sushil, David Hershleifer and Ivo Welch (1992). "A Theory of Fads, Fashion, Custom, and Cultural Change as Informational Cascades," *Journal of Political Economy*, Vol.100, No.5.

[7] Bongaarts, John and Susan Cotts Watkins (1996). "Social Interactions and Contemporary Fertility Transitions," *Population and Development Review*, Vol.22, No.4.

[8] Brock, William and Steven Durlauf (2001). "Discrete Choice with Social Interactions," *Review of Economic Studies*, Vol. 68.

[9] Bulatao, Rodolfo (1998). *The Value of Family Planning Programs in Developing Countries.* RAND: Santa Monica, California.

[10] Cleland, John, James Phillips, Sajeda Amin and G.M. Kamal (1994). *The Determinants of Reproductive Change in Bangladesh.* The World Bank: Washington D.C.

[11] Ellison, Glenn (1994). "Cooperation in the Prisoner's Dilemma with Anonymous Random Matching," *The Review of Economic Studies*, Vol.61(3).

[12] Eshel, Ilan, Larry Samuelson and Avner Shaked (1998). "Altruists, Egoists, and Hooligans in a Local Interaction Model," *American Economic Review*, Vol.88, No.1.

[13] Fauveau, V., B. Wojtyniak, M.A. Koenig, J. Chakraborty and A.I. Chowdhury (1989). "Epidemiology and Cause of Deaths Among Women in Rural Bangladesh," *International Journal of Epidemiology* , Vol. 18, No. 1.

[14] Glaeser, E., B. Sacerdote and J. Scheinkman (1996). "Crime and Social Interactions," *Quarterly Journal of Economics*, Vol. 106.

[15] Härdle, Wolfgang (1990). *Applied Nonparametric Regression*, Econometric Society Monographs No.11, Cambridge University Press, Cambridge.

[16] Hsiao, Cheng (1986). *Analysis of Panel Data*, Econometric Society Monographs No.11, Cambridge University Press, Cambridge.

[17] Kandori, Michihiro (1992). "Social Norms and Community Enforcement," *Review of Economic Studies*, Vol. 59.

[18] Koenig, M.A., V. Fauveau, A.I. Chowdhury, J. Chakraborty and M.A. Khan (1988). "Maternal Mortality in Matlab, Bangladesh: 1976-85," *Studies in Family Planning*, Vol. 19, No. 2.

[19] Lohmann, Susanne (1994). "The Dynamics of Informational Cascades: The Monday Demonstrations in Leipzig, East Germany, 1989-91," *World Politics*, Vol.47.

[20] Maddala, G.S. (1987). "Limited Dependent Variable Models Using Panel Data," *Journal of Human Resources*, Vol.22.

[21] Manski, Charles F. (1993). "Identification of Endogenous Social Effects: The Reflection Problem," *Review of Economic Studies*, Vol.60.

[22] Mita, Rezina and Ruth Simmons (1995). "Diffusion of the Culture of Contraception: Program Effects on Young Women in Rural Bangladesh," *Studies in Family Planning*, Vol. 26, No. 1.

[23] Munshi, Kaivan (2001). "Social Learning in a Heterogeneous Population: Technology Diffusion in the Indian Green Revolution," University of Pennsylvania, mimeo.

[24] National Research Council (1993). *Factors Affecting Contraceptive Use in Sub-Saharan Africa.* National Academy Press: Washington D.C.

[25] Okuno-Fujiwara, Masahiro and Andrew Postlewaite (1995). "Social Norms and Random Matching Games," *Games and Economic Behavior*, Vol.9.

[26] Pollak, Robert A. (1976). "Interdependent Preferences," *American Economic Review*, Vol.66.

[27] Postlewaite, Andrew (1998). "Social Status, Norms and Economic Performances: The Social Basis of Interdependent Preferences," *European Economic Review*, Vol.42.

[28] Quah, Danny T. (1993). "Empirical cross-section dynamics in economic growth," *European Economic Review*, Vol.37.

[29] Quah, Danny T. (1997). "Empirics for Growth and Distribution: Stratification, Polarization, and Convergence Clubs," *Journal of Economic Growth*, Vol.2.

[30] Ray, Debraj (1998). *Development Economics*, Princeton University Press: Princeton, New Jersey.

[31] Rosario, Shanti (1992). *Purity and Communal Boundaries: Women and Social Change in a Bangladeshi Village.* Zed Press: London.

[32] Rosenzweig, Mark R. (1990). "Population Growth and Human Capital Investments: Theory and Evidence," *Journal of Political Economy*, Vol.98.

[33] Schuler, Sidney Ruth and Syed M. Hashemi (1994). "Credit Programs, Women's Empowerment, and Contraceptive Use in Rural Bangladesh," *Studies in Family Planning*, Vol25, No.2.

[34] Schuler, Sidney Ruth, Syed M. Hashemi and Ann P. Riley (1997). "The Influence of Women's Changing Roles and Status in Bangladesh's Fertility Transition: Evidence from a Study of Credit Programs and Contraceptive Use," *World Development*, Vol25, No.4.

[35] Schultz, T. Paul (1988). "Economic Demography and Development: New Directions in an Old Field," in *The State of Development Economics*, Gustav Ranis and T. Paul Schultz (eds.), Basil Blackwell, Oxford.

[36] Simmons, Ruth (1996). "Women's Lives in Transition: A Qualitative Analysis of the Fertility Decline in Bangladesh," *Studies in Family Planning*, Vol. 27, No. 5.

[37] Simmons, Ruth, Rezina Mita and Michael A. Koenig (1992). "Employment in Family Planning and Women's Status in Bangladesh," *Studies in Family Planning*, Vol. 23, No. 2.

[38] Topa, Georgio (2001). "Social Interactions, Local Spillovers and Unemployment," *Review of Economic Studies*, Vol. 68.

Figure 1: Contraceptive Prevalence Over Time

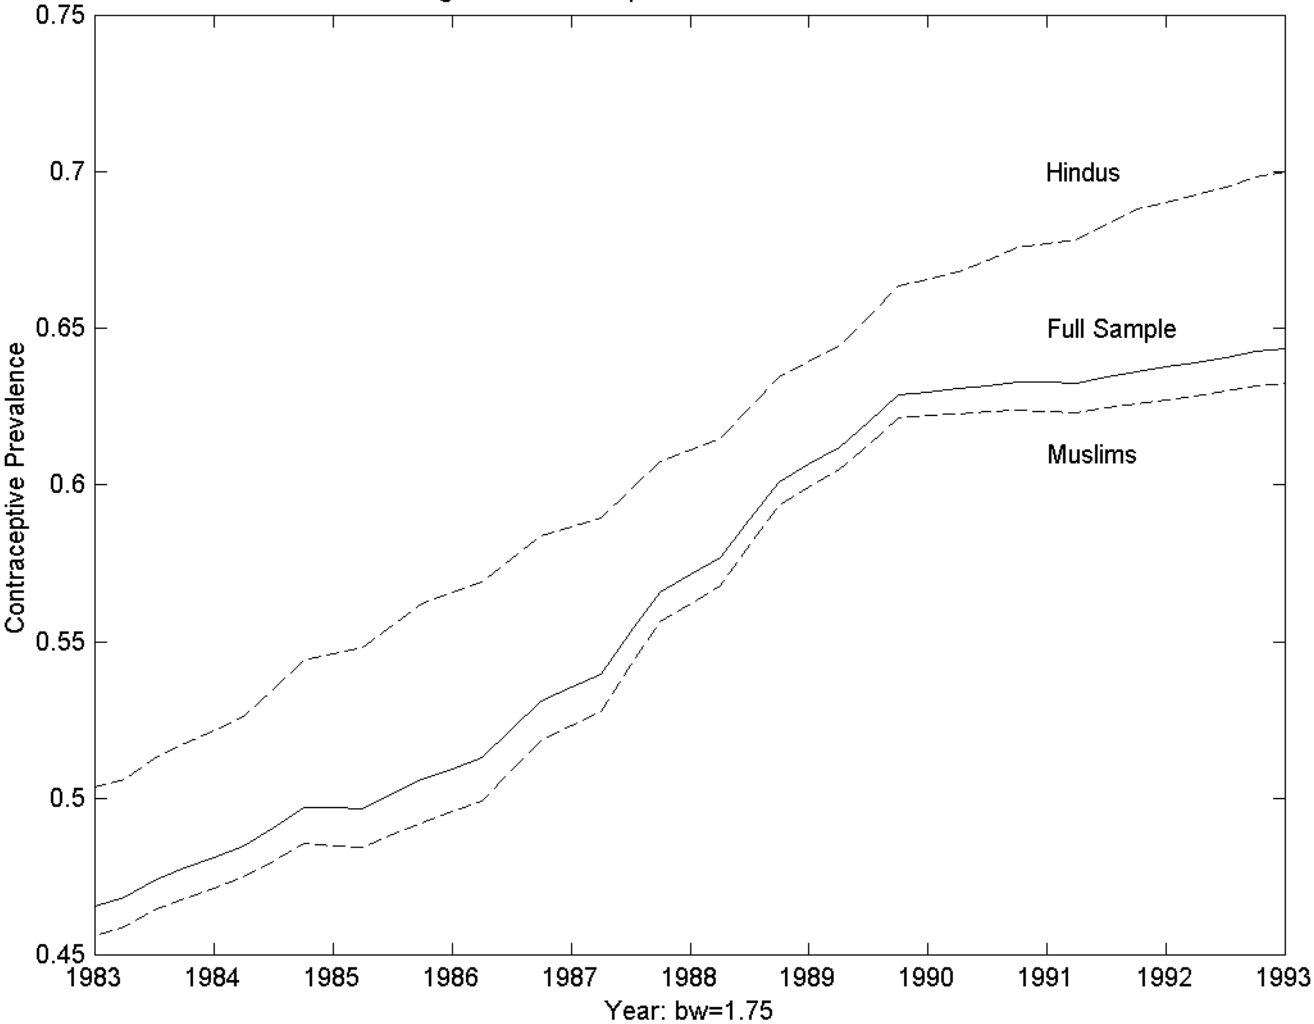

Figure 2: Contraceptive Prevalence Over Time - By Cohort

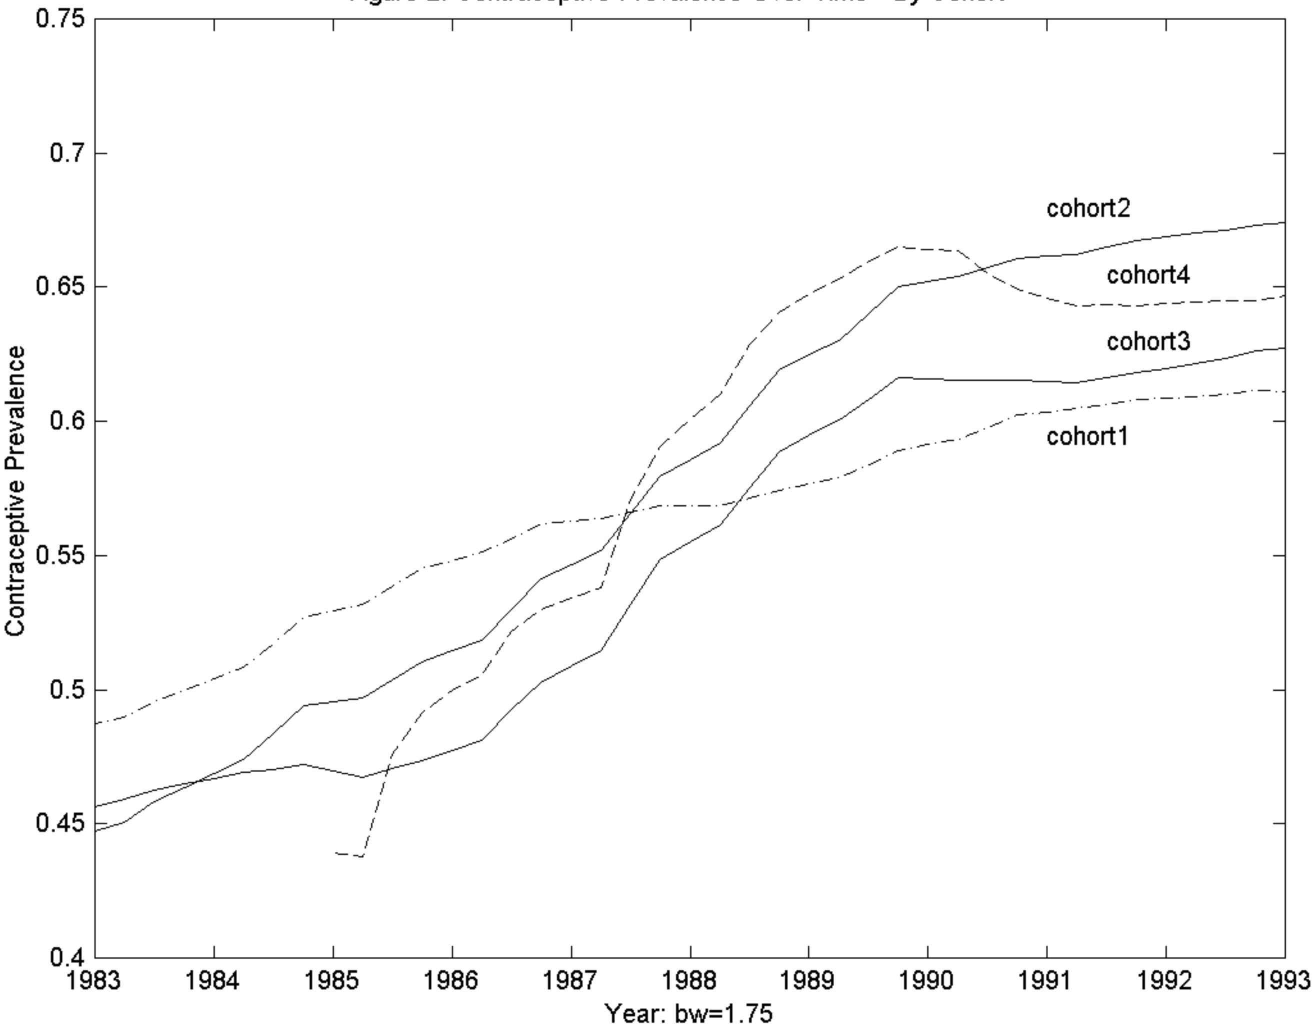

Figure 3: Contraceptive Prevalence Across Villages

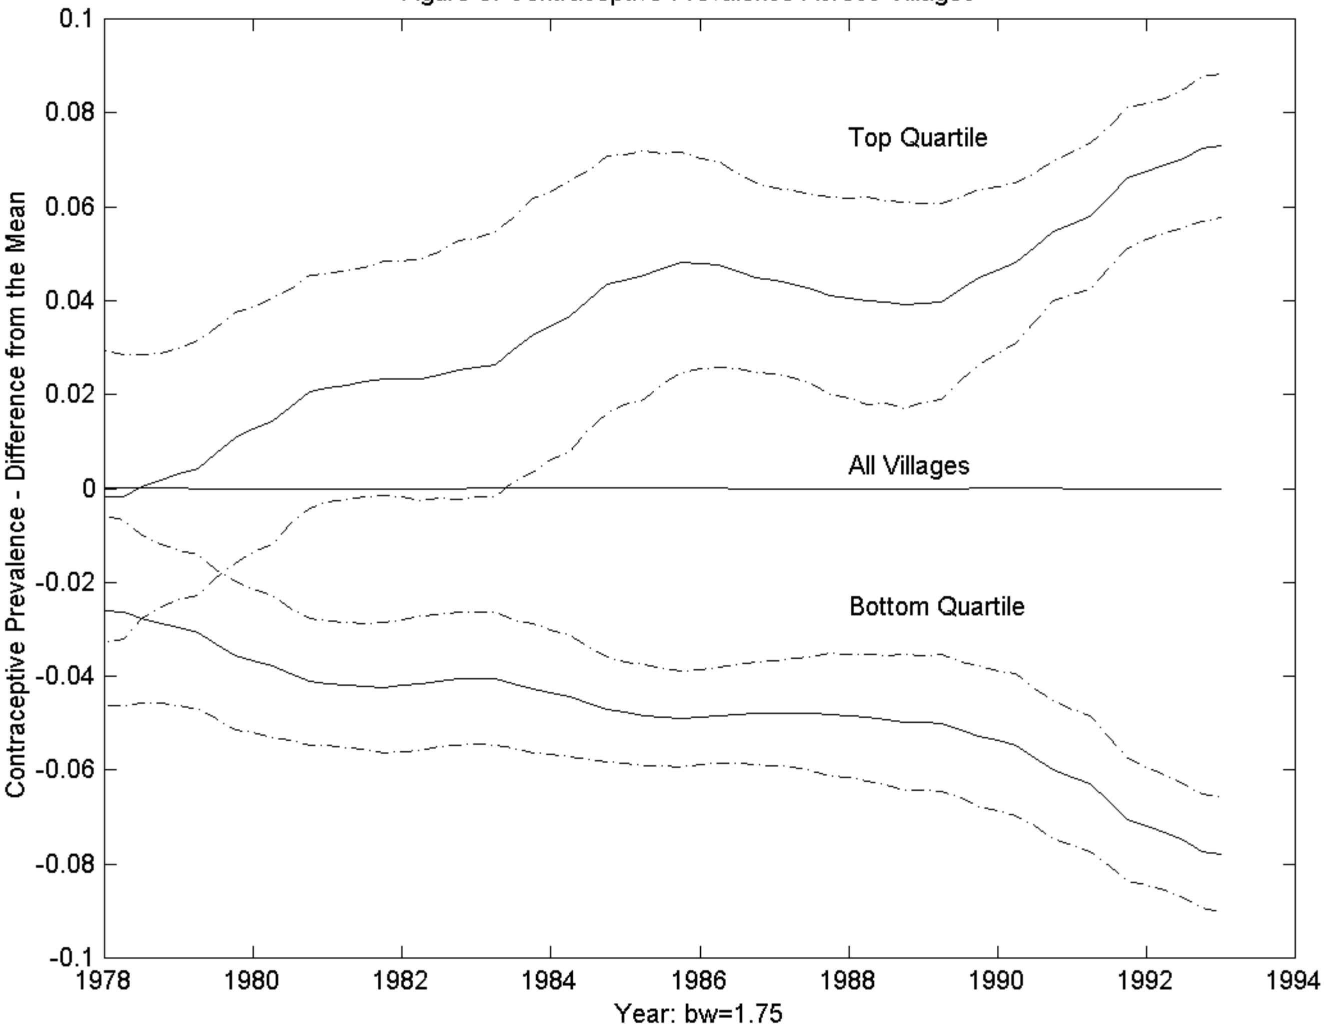

Figure 4: Change in Beliefs over Time

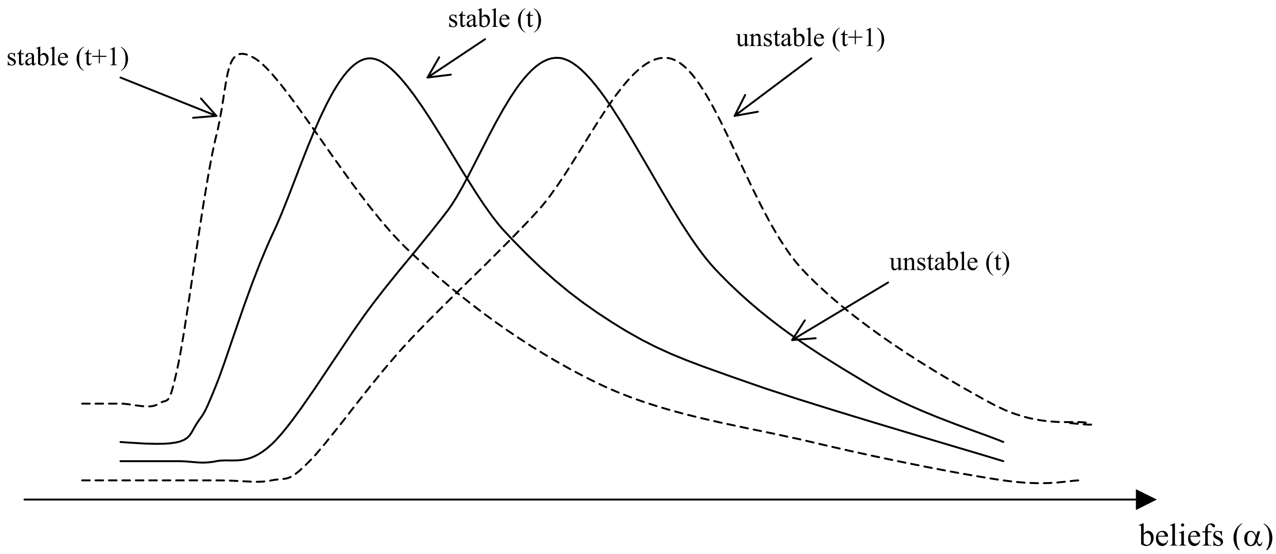

Figure 5: Simulated Contraceptive Prevalence - Two Types

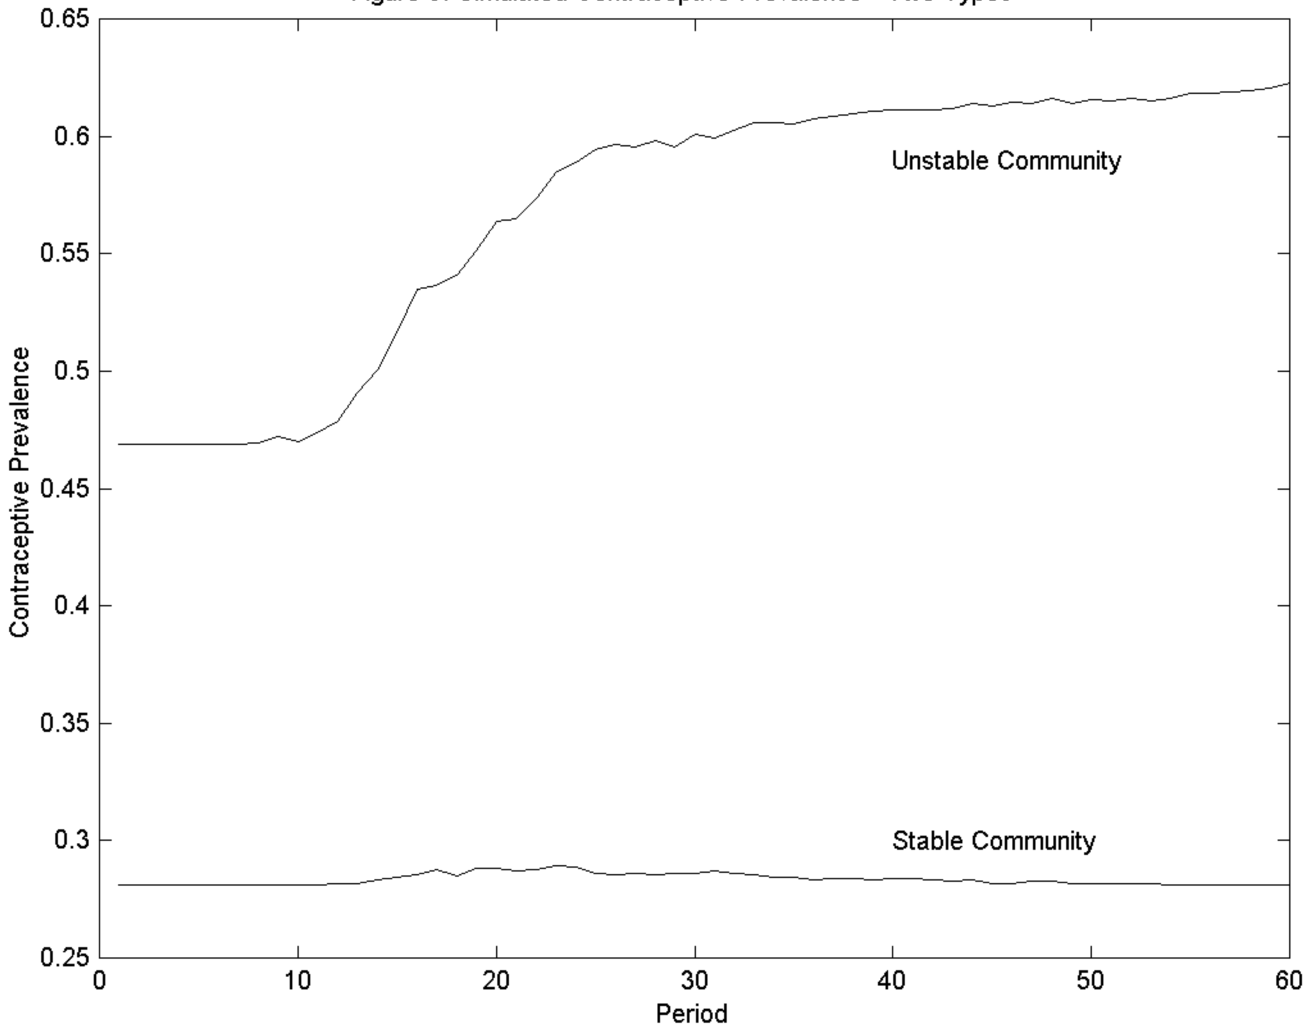

Figure 6: Simulated Contraceptive Prevalence - Three Types

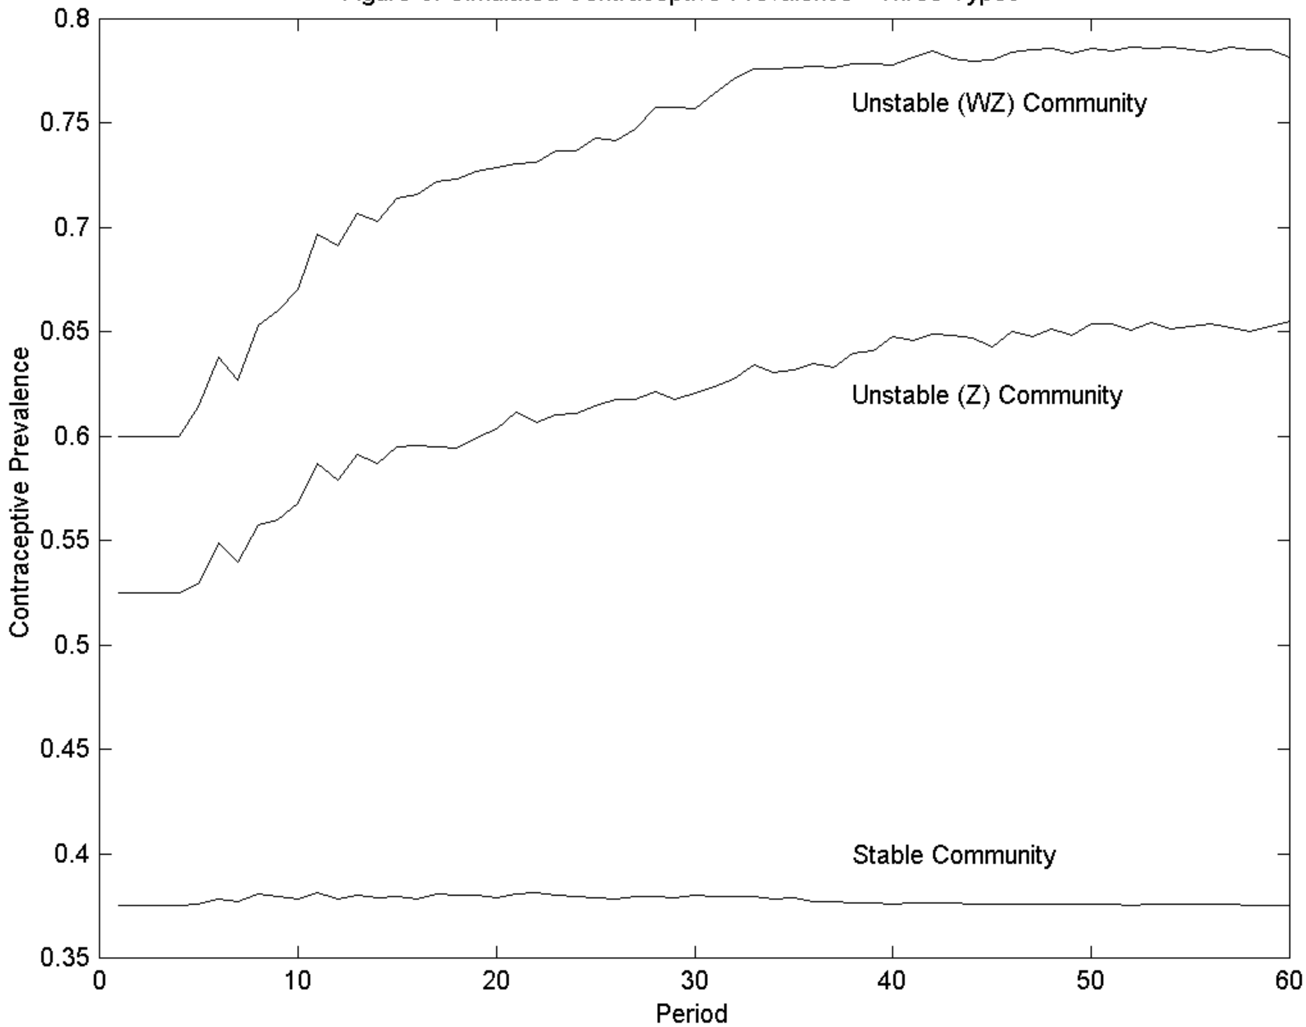

Figure 7: Simulated Decision Rule Parameters

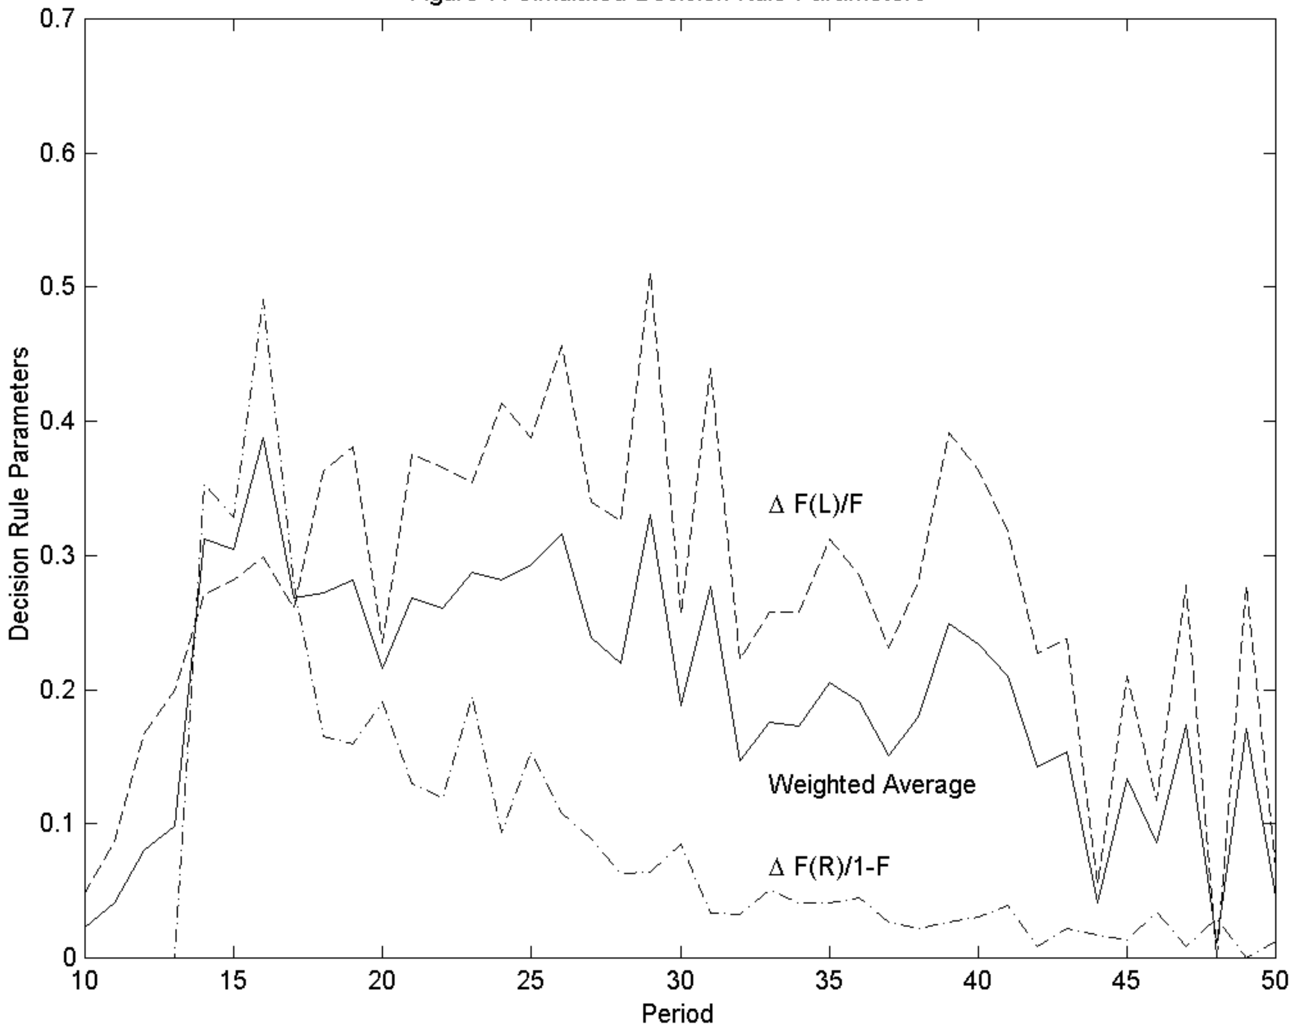

Table 1: Inter-quartile Transition Probabilities

1993 Distribution
mean 0.55
std.dev. 0.06

0.25 quantile 0.51 | .50 quantile 0.56 | 0.75 quantile 0.60

1978 Distribution
mean 0.27
std.dev. 0.09

0.25 quantile 0.21
0.50 quantile 0.27
0.75 quantile 0.31

0.33	0.33	0.06	0.28
0.35	0.29	0.24	0.12
0.22	0.22	0.22	0.33
0.12	0.12	0.53	0.24

Table 2: Descriptive Statistics

	Full Sample	Hindus	Muslims	Illiterate	Literate
	(1)	(2)	(3)	(4)	(5)
Panel A : Individual Characteristics					
Age	29.44	29.91	29.34	30.49	27.75
	(8.01)	(8.00)	(8.01)	(8.18)	(7.44)
Number of children	2.41	2.18	2.45	2.57	2.14
	(1.99)	(1.79)	(2.03)	(2.05)	(1.86)
Education	2.12	1.48	2.26	0.00	5.53
	(3.12)	(2.68)	(3.19)	--	(2.55)
Husband's education	3.21	3.07	3.24	1.53	5.91
	(4.00)	(3.81)	(4.04)	(2.62)	(4.34)
Panel B : Occupation of Household Head (%)					
Farming	34.48	23.45	36.88	30.32	41.16
Fishing	5.80	26.18	1.37	8.07	2.15
Business	6.75	8.37	6.40	6.30	7.47
Housework	10.46	6.81	11.26	10.00	11.21
Other	42.51	35.20	44.10	45.31	38.01
Total	100.00	100.00	100.00	100.00	100.00
Panel C : Asset Ownership					
Land (hectares)	1.00	0.72	1.06	0.82	1.29
	(2.55)	(1.39)	(2.74)	(2.41)	(2.74)
Cows	1.06	0.81	1.11	0.91	1.28
	(1.57)	(1.42)	(1.59)	(1.46)	(1.70)
Boats	0.55	0.63	0.54	0.55	0.56
	(0.61)	(0.76)	(0.57)	(0.61)	(0.60)
No. of Observations	21,570	3,847	17,723	13,288	8,282
Panel D : Contraceptive Prevalence					
Probability of using contraceptives	0.55	0.59	0.54	0.53	0.57
	(0.50)	(0.49)	(0.50)	(0.50)	(0.50)
No. of Observations	144,186	26,414	117,772	91,727	52,459

Note : Means (standard deviations) in Panel A, Panel C and Panel D.
The individual is the unit of observation in Panels A-C. The individual-year is the unit of observation in Panel D.
All statistics in this table are computed over the full 1983-93 sample period.

Table 3: Contraception Decision Rule - Polynomial Function of Time

	Dependent variable: Contraception				
	Village as social unit	With interaction	Fixed effects	Religion-village as social unit	Annual data
	(1)	(2)	(3)	(4)	(5)
Lagged contraceptive prevalence	0.138	0.156	0.064	0.143	0.159
	(0.028)	(0.031)	(0.033)	(0.024)	(0.048)
Lagged prevalence*time	0.026	0.027	0.060	0.019	0.067
	(0.012)	(0.012)	(0.013)	0.011	(0.021)
Lagged prevalence*time squared	-0.002	-0.002	-0.006	-0.002	-0.006
	(0.001)	(0.001)	(0.001)	(0.001)	(0.002)
Lagged decision	0.781	0.797	0.592	0.780	0.602
	(0.005)	(0.012)	(0.007)	(0.005)	(0.008)
Lagged decision*time	-0.022	-0.021	-0.028	-0.022	-0.034
	(0.002)	(0.002)	(0.003)	(0.002)	(0.004)
Lagged decision*time squared	0.001	0.001	0.002	0.001	0.002
	(0.0002)	(0.0002)	(0.0004)	(0.0002)	(0.0005)
Lagged prevalence*lagged decision	--	-0.039	--	--	--
		(0.028)			
R squared	0.533	0.533	0.512	0.533	0.298
Number of observations	252,833	252,833	252,833	252,833	121,136
Box-Pearson Q statistic	0.002	0.002	0.021	0.002	0.004

Note: Standard errors in parentheses.

Standard errors are robust to heteroskedasticity and correlated residuals within each village-year.

$Q \sim X_1^2$ under H_0: no serial correlation. The critical value above which the null is rejected at the 5 percent significance level is 3.84.

Time is year minus 1983.

Column 1: contraceptive prevalence measured at the level of the village (also for Columns 2,3,5).

Column 2: regression includes lagged prevalence*lagged decision interaction term.

Column 3: regression with individual fixed effects.

Column 4: contraceptive prevalence measured seprately by religious group within the village.

Column 5: regression with annual data.

Time, time-squared, age, age squared are included as control variables in all the regressions.

Table 4: Contraception Decision Rule - Spline Function of Time

	Dependent variable: Contraception				
	Village as social unit	With interaction	Fixed effects	Religion-village as social unit	Annual data
	(1)	(2)	(3)	(4)	(5)
Lagged contraceptive prevalence*period 1	0.151	0.190	0.147	0.155	0.186
	(0.019)	(0.023)	(0.024)	(0.016)	(0.039)
Lagged contraceptive prevalence*period 2	0.257	0.301	0.237	0.228	0.382
	(0.021)	(0.027)	(0.027)	(0.017)	(0.039)
Lagged contraceptive prevalence*period 3	0.182	0.233	0.095	0.167	0.283
	(0.024)	(0.030)	(0.031)	(0.021)	(0.033)
Lagged decision*period 1	0.756	0.792	0.553	0.755	0.577
	(0.003)	(0.012)	(0.005)	(0.003)	(0.006)
Lagged decision*period 2	0.693	0.735	0.484	0.692	0.487
	(0.004)	(0.014)	(0.005)	(0.004)	(0.007)
Lagged decision*period 3	0.682	0.732	0.473	0.681	0.478
	(0.003)	(0.016)	(0.005)	(0.003)	(0.006)
Lagged prevalence*lagged decision	--	-0.082	--	--	--
		(0.026)			
R squared	0.533	0.533	0.525	0.534	0.298
Number of observations	252,833	252,833	252,833	252,833	121,136
Box-Pearson Q statistic	0.003	0.002	0.019	0.002	0.005

Note: Standard errors in parentheses.

Standard errors are robust to heteroskedasticity and correlated residuals within each village-year.

$Q \sim X_1^2$ under H_0: no serial correlation. The critical value above which the null is rejected at the 5 percent significance level is 3.84.

Sample is divided into three time periods - period 1: 1984-86, period 2: 1987-89, period 3: 1990-93.

Column 1: contraceptive prevalence measured at the level of the village (also for Columns 2,3,5).

Column 2: regression includes lagged prevalence*lagged decision interaction term.

Column 3: regression with individual fixed effects.

Column 4: contraceptive prevalence measured seprately by religious group within the village.

Column 5: regression with annual data.

Time-period dummies, age, age squared are included as control variables in all the regressions.

Table 5: Contraception Decision Rule - Robustness Tests

	Dependent variable: Contraception						
	Without never-users	Cohort 2	Cohort 3	Villages in bottom quartile	Villages in top quartile	Twice-lagged decision	Accounting for birth-spacing
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Lagged contraceptive prevalence*period 1	0.141	0.088	0.165	0.112	0.181	0.125	0.024
	(0.022)	(0.020)	(0.029)	(0.043)	(0.031)	(0.020)	(0.012)
Lagged contraceptive prevalence*period 2	0.259	0.163	0.269	0.169	0.281	0.230	0.153
	(0.022)	(0.022)	(0.029)	(0.043)	(0.052)	(0.019)	(0.018)
Lagged contraceptive prevalence*period 3	0.180	0.098	0.123	0.121	0.050	0.160	0.103
	(0.026)	(0.019)	(0.029)	(0.059)	(0.051)	(0.025)	(0.017)
Lagged decision*period 1	0.706	0.765	0.642	0.764	0.760	0.731	0.860
	(0.004)	(0.004)	(0.006)	(0.005)	(0.006)	(0.005)	(0.003)
Lagged decision*period 2	0.635	0.742	0.599	0.704	0.680	0.686	0.789
	(0.004)	(0.005)	(0.006)	(0.007)	(0.010)	(0.005)	(0.004)
Lagged decision*period 3	0.604	0.767	0.633	0.683	0.665	0.673	0.767
	(0.004)	(0.005)	(0.004)	(0.006)	(0.008)	(0.004)	(0.003)
R squared	0.472	0.613	0.420	0.534	0.558	0.536	0.691
Number of observations	225,500	85,749	108,149	87,158	55,584	227,764	265,322
Box-Pearson Q statistic	0.073	0.057	0.365	0.005	0.000	0.0004	0.000

Note: Standard errors in parentheses.

Standard errors are robust to heteroskedasticity and correlated residuals within each village-year.

$Q \sim X_1^{\,2}$ under H_0: no serial correlation. The critical value above which the null is rejected at the 5 percent significance level is 3.84.

Sample is divided into three time periods - period 1: 1984-86, period 2: 1987-89, period 3: 1990-93.

Column 1: drop individuals who never use contraceptives from the sample.

Column 2: restrict the sample to individuals born between 1948 and 1958 (Cohort 2 in Figure 2). Contraceptive prevalence is measured within the cohort.

Column 3: restrict the sample to individuals born between 1958 and 1968 (Cohort 3 in Figure 2). Contraceptive prevalence is measured within the cohort.

Column 4: restrict the sample to villages in the bottom quartile of the 1993 village contraceptive prevalence distribution.

Column 5: restrict the sample to villages in the top quartile of the 1993 village contraceptive prevalence distribution.

Column 6: include twice-lagged individual adoption. The coefficients (std.errors) for the three time periods are: 0.029(0.004),0.018(0.004),0.026(0.004).

Column 7: use a 1.5 year moving window (period t to t-3) to compute adoption - if there is any contraceptive use in these four periods, then the individual is treated as an adopter (accounts for birth-spacing).

Time-period dummies, age, age squared are included as control variables in all the regressions.

Table 6: Partitioning the Village by Religion

	Dependent Variable: Contraception							
	All villages		More than 5% Hindus/Muslims		More than 15% Hindus/ Muslims		Annual data	
	Muslims	Hindus	Muslims	Hindus	Muslims	Hindus	Muslims	Hindus
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged contraceptive prevalence	0.197	0.138	0.178	0.158	0.198	0.161	0.251	0.247
(own group)	(0.017)	(0.016)	(0.022)	(0.018)	(0.024)	(0.021)	(0.041)	(0.029)
Lagged contraceptive prevalence	0.003	0.005	-0.0002	0.013	-0.006	0.005	-0.006	0.008
(other group)	(0.007)	(0.007)	(0.012)	(0.018)	(0.014)	(0.030)	(0.023)	(0.034)
Lagged contraception	0.698	0.712	0.704	0.710	0.706	0.717	0.504	0.522
	(0.003)	(0.005)	(0.004)	(0.005)	(0.004)	(0.006)	(0.007)	(0.009)
R squared	0.513	0.559	0.520	0.558	0.521	0.565	0.285	0.344
Number of observations	139,875	43,101	79,927	29,771	49,730	20,756	40,198	14,838
Box-Pearson Q statistic	0.000	0.003	0.001	0.003	0.002	0.006	0.005	0.011

Note: Standard errors in parentheses.

Standard errors are robust to heteroskedasticity and correlated residuals within each village-year.

$Q \sim X_1^2$ under H_0: no serial correlation. The critical value above which the null is rejected at the 5 percent significance level is 3.84.

Columns 1-2: Sample includes all mixed-religion villages.

Columns 3-4: Sample restricted to villages with more than 5% Hindus and Muslims.

Columns 5-6: Sample restricted to villages with more than 15% Hindus and Muslims.

Columns 7-8: Annual data (using villages with more than 5% Hindus and Muslims).

Table 7: Partitioning the Village by Religion - Robustness Checks

	Dependent Variable: Contraception									
	Six-monthly dummies		Fixed effects		With bari adoption		With occupations		CHW effect	
	Muslims	Hindus	Muslims	Hindus	Muslims	Hindus	Muslims	Hindus	CHW's religion	Other religion
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)
Lagged contraceptive prevalence	0.183	0.158	0.123	0.138	0.164	0.139	0.171	0.139	0.208	0.154
(own religious group)	(0.021)	(0.018)	(0.034)	(0.031)	(0.023)	(0.019)	(0.022)	(0.019)	(0.019)	(0.018)
Lagged contraceptive prevalence	0.006	0.019	0.022	0.029	-0.001	0.007	-0.004	0.022	-0.011	0.034
(other religious group)	(0.011)	(0.018)	(0.020)	(0.022)	(0.012)	(0.019)	(0.012)	(0.019)	(0.011)	(0.018)
Lagged decision	0.704	0.710	0.490	0.485	0.702	0.709	0.704	0.709	0.706	0.714
	(0.004)	(0.005)	(0.005)	(0.008)	(0.004)	(0.005)	(0.004)	(0.005)	(0.004)	(0.005)
Bari adoption	--	--	--	--	0.026	0.023	--	--	--	--
					(0.005)	(0.008)				
R squared	0.520	0.559	0.514	0.528	0.519	0.558	0.520	0.560	0.531	0.547
Number of observations	79,927	29,771	79,927	29,771	75,297	27,973	75,524	28,779	78,282	36,869
Box-Pearson Q statistic	0.001	0.003	0.015	0.013	0.001	0.001	0.001	0.002	0.004	0.009

Note: Standard errors in parentheses.

Standard errors are robust to heteroskedasticity and correlated residuals in each village-year.

$Q \sim X_1^{\,2}$ under H_0: no serial correlation. The critical value above which the null is rejected at the 5 percent significance level is 3.84.

Columns 1-2: Full set of (six-monthly) time dummies included as control variables.

Columns 3-4: Individual fixed effects included as control variables.

Columns 5-6: Include bari level contraceptive prevalence.

Columns 7-8: Include occupation categories.

Column 9: Individuals that share the same religion as the CHW. Column 10: Individual and CHW have different religions.

Time period dummies (1984-86, 1987-89, 1990-93) are included as control variables in all regressions except Columns 1-2.

Age, age squared are included as control variables in all regressions.

Sample restricted to villages with more than 5% Hindus and Muslims in all regressions

Table 8: Alternative Partitions of the Village - Age and Education

	Dependent Variable: Contraception							
Cut-off:	0.50 quantile age		0.75 quantile age		Literacy		0.75 quantile educn.	
Group:	< 0.50	> 0.50	< 0.75	> 0.75	Illiterate	Literate	< 0.75	> 0.75
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Lagged contraceptive prevalence	0.158	0.112	0.177	0.055	0.133	0.121	0.154	0.082
(own group)	(0.020)	(0.012)	(0.016)	(0.011)	(0.014)	(0.016)	(0.013)	(0.019)
Lagged contraceptive prevalence	0.078	0.025	0.044	0.035	0.058	0.081	0.036	0.121
(other group)	(0.016)	(0.013)	(0.010)	(0.016)	(0.012)	(0.019)	(0.009)	(0.024)
Lagged contraception	0.620	0.793	0.661	0.851	0.714	0.684	0.712	0.675
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.002)	(0.004)
R squared	0.401	0.652	0.466	0.738	0.544	0.495	0.539	0.481
Number of observations	121,050	119,585	182,762	57,873	154,077	86,558	185,092	55,095
Box-Pearson Q statistic	1.208	0.125	0.049	0.222	0.008	0.002	0.007	0.007

Note: Standard errors in parentheses.

Standard errors are robust to heteroskedasticity and correlated residuals in each village-year.

$Q \sim X_1^2$ under H_0: no serial correlation. The critical value above which the null is rejected at the 5 percent significance level is 3.84.

Columns 1 to 4 : Partition the village by age.

Columns 1-2: 0.5 age quantile cut-off. Columns 3-4: 0.75 age quantile cut-off.

Columns 5-8: Partition the village by education.

Columns 5-6: Literacy cut-off. Columns 7-8: 0.75 education quantile cut-off.

Time period dummies (1984-86, 1987-89, 1990-93), age, age squared are included as control variables in all regressions.