

NBER WORKING PAPER SERIES

RELIGIOUS MARKET STRUCTURE,
RELIGIOUS PARTICIPATION, AND OUTCOMES:
IS RELIGION GOOD FOR YOU?

Jonathan Gruber

Working Paper 11377
<http://www.nber.org/papers/w11377>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
May 2005

I am grateful to Max Farrell, Mike Frakes, Liz Jiminez, Liz Oltmans and David Seif for excellent research assistance. I am also grateful to Mark Chaves, Roger Finke, Larry Iannaccone, and Dan Olson for educating me on the economics and sociology of religion, and to them, Ed Glaeser, Larry Katz, and seminar participants at Harvard, MIT, NBER the University of Alabama-Birmingham, and George Mason University for helpful comments on this project. The views expressed herein are those of the author(s) and do not necessarily reflect the views of the National Bureau of Economic Research.

©2005 by Jonathan Gruber. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Religious Market Structure, Religious Participation, and Outcomes: Is Religion Good for You?
Jonathan Gruber
NBER Working Paper No. 11377
May 2005
JEL No. H3, N4

ABSTRACT

Religion plays an important role in the lives of many Americans, but there is relatively little study by economists of the implications of religiosity for economic outcomes. This likely reflects the enormous difficulty inherent in separating the causal effects of religiosity from other factors that are correlated with outcomes. In this paper, I propose a potential solution to this long standing problem, by noting that a major determinant of religious participation is religious market density, or the share of the population in an area which is of an individual's religion. I make use of the fact that exogenous predictions of market density can be formed based on area ancestral mix. That is, I relate religious participation and economic outcomes to the correlation of the religious preference of one's own heritage with the religious preference of other heritages that share one's area. I use the General Social Survey (GSS) to model the impact of market density on church attendance, and micro-data from the 1990 Census to model the impact on economic outcomes. I find that a higher market density leads to a significantly increased level of religious participation, and as well to better outcomes according to several key economic indicators: higher levels of education and income, lower levels of welfare receipt and disability, higher levels of marriage, and lower levels of divorce.

Jonathan Gruber
MIT Department of Economics
E52-355
50 Memorial Drive
Cambridge, MA 02142-1347
and NBER
gruberj@mit.edu

Religion plays an important role in the lives of many Americans. Over two-thirds of Americans belong to a church or other religious organization, and this has risen substantially over time. Two-fifths of Americans attend church in a typical week, and 95% profess belief in “the existence of God or a universal spirit” (Iannaccone, 1998). Giving to religious causes accounts for more than two-thirds of all reported individual charitable contributions. Religiosity is not confined to particular income groups, racial groups, or locations in the U.S.: religious adherence and participation is widespread among all demographic groups.

Given this important role, there is relatively little study by economists of the implications of religiosity for economic outcomes. This is certainly not the case in other disciplines. There are hundreds of articles in sociology, psychology, and medicine that overwhelmingly document the positive impacts of religiosity on a wide variety of outcomes. Those few studies by economists have also found that religiosity, and in particular religious participation, is strongly associated with positive outcomes such as reduced frequency of socially deviant activities among youth.

The small literature on this topic in economics likely reflects the heightened sensitivity in our profession to the enormous difficulty inherent in separating the causal effects of religiosity from other factors that are correlated with outcomes. Most factors which determine the religiosity of any given individual are likely correlated with their outcomes through other channels as well. This is true of both short run factors which may cause fluctuations in religious devotion, such as particularly good or bad shocks to personal well-being, and long run factors which cause heterogeneity across individuals, such as correlation between religiosity and ambition or ability.

In this paper, I propose a potential solution to this long standing problem, by drawing on

the literature on the sociology of religion. Starting with the influential work of Finke and Stark (1988), this literature has posited and tested a host of hypotheses about how the structure of religious markets influences religiosity, and in particular religious participation. One factor that has been suggested as a determinant of religious participation is market density, or the share of the population in an area which is of an individual's religion. If a larger percentage of the local population shares your religion, there may be easier access to religious observation (a "distance" effect) or other social interactions that are related to religious participation.

I investigate the impact of market density on both religious participation and economic outcomes. The key insight that makes the analysis possible is that exogenous predictions of market density can be formed based on area ancestral mix. For a Catholic person of Italian extraction, living near persons of Polish extraction or living near persons of Swedish extraction is likely irrelevant for any aspect of life except for the distribution of religious preference in the area; there will be a higher percentage Catholic if the neighboring group is Poles rather than Swedes. Thus, Italians who live near Poles will have a higher religious market density than Italians who live near Swedes, yet they should be otherwise identical.

To model impacts of market density on religious participation, I rely on the General Social Survey (GSS), a nationally representative survey which collects data on religious preference and religious participation. For the outcomes analysis, I use the IPUMS sample of the 1990 Census. Of course, the Census in the U.S. does not ask about religious preference. But I can form a strong prediction of religious preference based on one's ancestral heritage, so that I can model outcomes as a function of the correlation of the religious preference of one's own heritage with the religious preference of other heritages that share one's area.

My results are striking. I find that a higher market density leads to a significantly increased level of religious participation, and as well to better outcomes according to several key economic indicators: higher levels of education and income, lower levels of welfare receipt and disability, higher levels of marriage, and lower levels of divorce. These results are robust to a variety of specification checks.

The paper proceeds as follows. In Part I, I review the literature on religiosity and outcomes, and on religious market structure and religiosity. In Part II, I describe the data sources and empirical strategy for the analysis. Part III presents the main results, and Part IV assesses their robustness to alternative interpretations. Part V concludes.

Part I: Background

Religiosity and Outcomes

There is an enormous literature which documents a positive cross-sectional association between religious participation, or other measures of “religiosity”, and outcomes. The earliest claim of such a relationship is likely Durkheim’s 1897 study of the role of religiosity and suicide rates. Durkheim found that suicide rates were lower in the more socially integrated Catholic communities than in less integrated Protestant areas. He attributed this difference not to Catholic teachings on suicide, but rather to the “tendency of Catholicism to promote integration through ritual and community, and to regulate standards of faith, piety, and moral conduct” (Ellison, Burr & McCall, 1997). Another early contributor was Weber, whose 1908 work (as referenced by Iannaccone, 1998), which claimed that the industrial revolution was made possible by the Protestant Reformation.

There have been a large volume of subsequent studies of the impact of religiosity on many different outcomes, which I can review only incompletely; I provide some sample citations here. One area of significant focus is deviant behaviors, with more frequent religious participation and increased religiosity being associated with lower levels of adult criminal behavior, at both the individual level (e.g. Evans, Cullen, Dunaway, and Burton, 1995) and the area level (Lipford, McCorkmick, and Tollison, 1993; Hull and Bold, 1995), and with lower levels of substance abuse (e.g. National Center on Addiction and Substance Abuse, 2001). There is a parallel literature showing the retarding effects of religious participation and religiosity on deviant behavior of many types among youths (e.g. Bachman et al, 2002; Wallace and Williams, 1987). A notable contribution to this literature among economists is Freeman (1986), who finds that churchgoing among young black males favorably affects the allocation of time, school attendance, work activity, and the (in)frequency of socially deviant activity. Freeman highlights the problems with interpreting causally these types of results, but he notes that “the pattern of statistical results suggests that at least some part of the churchgoing effect is the result of actual causal impact”

Another major area of emphasis of the existing literature is health impacts of religious participation and religiosity on health status. Levin and Vanderpool (1987) and Hummer, Rogers, Nam and Ellison (1999) review this literature, documenting a consistently strong relationship between religiosity, usually proxied by religious participation, and health status/outcomes. Hummer et al. extend this literature by using matched death record data at the individual level to show a negative relationship between religious participation and mortality. Levin’s (1994) more critical review concludes that while these associations are valid, there is no

convincing evidence that they are causal.

There is also a literature documenting a very strong correlation between religiosity and self-reported measures of well-being (Ellison, 1991; Hout and Greeley, 2003). Gruber and Mullainathan (2002) find that the effect on self-reported happiness of moving from never attending to attending weekly is comparable to happiness effect of moving from bottom to top income quartile. In addition, religiosity is associated with more marital stability (e.g. Lehrer and Chiswick, 1993).

There have been fewer studies of measures of the impact of religiosity on more standard economic measures of outcomes. There is a large literature on religious preference and earnings (Tomes, 1985). This literature concludes that there is a large advantage for Jews in terms of both earnings and return to human capital, with mixed rankings for other religions. I am not aware of any work on religious participation and earnings or other income measures.

While this large literature has often recognized the difficulties of assigning causal interpretations to these positive impacts of religiosity, there has not been any fully satisfactory resolution to this problem. Those articles that are most sensitive to this issue try to address it using either rich background controls or longitudinal analysis of changes in religious participation, but neither of these approaches is entirely satisfactory.

Religious Market Structure and Religiosity

The literature on religious market structure and religiosity, which is almost exclusively found in sociology, owes its origins to Adam Smith. In a largely ignored chapter in *The Wealth of Nations*, “Smith argued that self-interest motivates clergy just as it does secular producers; that

market forces constrain churches just as they constrain secular firms; and that the benefits of competition, the burdens of monopoly, and the hazards of government regulation are as real for religion as for any other sector of the economy” (Iannaccone, 1998). This line of argument has formed the basis for a number of analyses by the sociologists Finke and Stark (e.g. Finke and Stark, 1988, 1992), who argued that there is empirical evidence for the proposition that more religious competition promotes religiosity.

A large subsequent literature has evaluated these claims, with mixed results.¹ Recent critical reviews in Chaves and Gorski (2001) and Voas, Olson and Crockett (2002) question the contention that more competition leads to more religiosity. Voas et al., in particular, question the entire empirical framework that has been used to test this contention, highlighting the mathematical linkage between the church membership data used to measure religious competition and the same (or similar) data used to measure rates of religious adherence.

There is a smaller literature which has focused on the question of how religious densities affect religiosity. There are only three studies which directly model church attendance as a function of density; Phillips (1998) and Olson (1998) find that higher density is associated with more attendance, while Perl and Olson (2000) find no significant association. Other studies which model religious charitable giving as a function of religious market density find a negative association (Zaleski and Zech, 1995; Perl and Olson, 2000). This is very consistent with the findings in Gruber (2004), who shows a negative association between religious attendance and religious charitable contributions. If attendance and contributions are substitutes, and more

¹This literature has mostly focused in sociology, with notable exceptions such as Iannaccone (1991) and Iannaccone et al. (1997).

density raises attendance, it may lower contributions. Thus, this set of findings is consistent with a positive effect of religious market density on religious attendance.

Part II: Data and Empirical Strategy

Data

I rely on two data sources for this analysis. The first is the General Social Survey (GSS), a nationally representative survey carried out in most years from 1973 through 2000, with roughly 1500-2000 households interviewed in each year. These data have information on location (PSU, which is roughly equivalent to MSA), religious participation/religiosity, ethnic heritage, demographics, and economic/social outcomes. For this analysis, I rely on all non-hispanic whites age 25 or older. I focus on non-hispanic whites only because there is very strong evidence of racial segregation in church going, so that the density of hispanics or nonwhites in a religion in some area is not likely to be relevant for the religious participation of whites in that area.

There are many possible definitions of religious affiliation, ranging from very general to very detailed. I use a fairly broad categorization for this work, dividing individuals into seven groups: Catholics, Jews, Liberal Protestants, Moderate Protestants, Conservative Protestants, Other, and None. This is the finest breakdown that is feasible, given the empirical strategy described below. The GSS reports data on roughly 150 denominations and subdenominations, and I use the classification from Roof & McKinney (1987) to divide the sample into these

groups.² The questions on religious affiliation have changed over time in the GSS, with a major revision in 1984, when detail was added on a host of Protestant denominations. I therefore rely on data from 1984 forwards for this analysis.

The main dependent variable in the GSS for this analysis is religious attendance. The GSS asks respondents about their frequency of religious attendance, with nine possible responses: never; less than once per year; about once or twice a year; several times a year; about once a month; two to three times a month; nearly every week; every week; and several times a week. For the basic analysis below, I simply use the linear index formed by these responses (with values 0 through 8), thereby implicitly assuming that there are equal impacts of each unit change in the attendance index. Since each unit is roughly twice the previous unit in time terms, this is roughly equivalent to a log scale in time attending. But I show as well the effects of alternative formulations below.

The second is the Integrated Public Use Microdata Sample (IPUMS) of the 1990 census. These data contain information for a 5% sample of the population on a variety of socioeconomic characteristics. I once again use a sample of non-hispanic whites ages 25 and older for this analysis. I use several measures of economic well-being as the dependent variable: the log of household income; the log of the ratio of household income to the poverty line (using the poverty-line normalization for family size); years of education; categories of educational attainment (high school dropout, high school graduate, some college, college graduate); whether

²Liberal Protestants include, for example, Episcopalians, Presbyterians, and the United Congregational Church/Unitarian-Universalists. Moderate Protestants include Methodists, Lutherans, and Non-Southern Baptists. Conservative Protestants include Southern Baptists, Pentecostals, and other Evangelical denominations.

the respondent is working; whether the respondent's spouse is working; whether the family receives welfare payments; whether the respondent is disabled; whether the respondent is married or divorced; and the number of children in the respondent's household.

Empirical Strategy - GSS Religious Attendance Estimates

I initially estimate models of religious attendance as a function of religious market density:

$$(1) \quad R_{ia} = \alpha + \beta \text{RELIGDENS}_{ia} + \gamma \text{ANCESDENS}_{ia} + \delta \text{AGESDENS}_{ia} + \theta \text{ANCES}_{ia} + \lambda \text{AGES}_{ia} + \rho \text{RELIG}_{ia} + \pi \text{AREA}_a + \tau \text{YEAR}_t + \epsilon$$

where: i indexes persons, a indexes areas, and t indexes years
 R is religious attendance
 RELIGDENS is the % of the population in the area that is of individual i 's religion
 ANCESDENS , AGESDENS are parallel measures of density of ancestry group and age/sex group³
 ANCES , AGES , and RELIG are fixed effects for the individual's ancestry, age/sex, and religion
 AREA and YEAR are fixed effects for area and year

This specification provides a “differences-in-differences” estimate of the impact of religious market density on own participation. Market density is simply the percentage of each PSU in the GSS that is of each of the seven religious groups. By including market and religious preference fixed effects, I can control for differences across groups and areas in their religiosity. Area-wide factors which causes differences in religious participation (history, weather, etc.) are

³I divide the data into 28 age/sex groups, using five year age groupings fully interacted with sex. In the GSS, this age/sex density is attached to individuals based on their actual age/sex groups. In the IPUMS data, we match this to individuals before collapsing the data into ancestry/MSA cells, and we then use the mean for the ancestry/MSA cell.

held constant, as are any religion-wide factors that are associated with religious participation (level of “devotion”). Such a regression asks, for example: are Catholics more religious when there is a higher percentage Catholic in an area, relative to Protestants in the same area, and relative to Catholics in other areas with a lower percentage Catholic? A positive effect of religious density on religious participation would be indicated by a positive coefficient on β .

Even in this differences-in-differences framework, there are three potential problems with the analysis. The first is classical measurement error; given the small GSS sample sizes, I am measuring denominational density in an area fairly imprecisely. Second, in addition, there could be systematic measurement error in market density. If individuals who participate more in religious activities are also more likely to report their religious affiliation in a survey (e.g. as opposed to saying that they are not religious), then there will be a direct spurious link between religious participation and religious density. Third, even conditional on these fixed effects, market density may be capturing majority status effects. That is, it could be that it is not being part of a large religious group, but rather being part of a large group per se, that determines religious attendance.

My approach to addressing these problems is to create an instrument for religious density which is based on one’s *ancestry*, and the ancestries of others in the area, rather than on religion per se. In particular, the instrument takes the following form:

$$(2) \quad \text{PRDEN}_{na} = \sum_r [P_r * (\sum_n P_n * \text{ANCES}_{na})]$$

where n indexes ancestries, r indexes religions, and a indexes areas
 PRDEN_{na} is predicted density for ancestry group n in area a
 P_r is the distribution of religious preferences r within each ancestry n
 ANCES_{na} is the density of each ancestral group n in each area a

This measure, PRDEN, is formed in two steps. First, for each ancestry, I multiply the density of that ancestry in an area times the probability that the ancestry is of a given religion. Summing these, the final term in expression (2) gives me the predicted density of each religion in each area; that is, based on the ancestral distribution in an area and the religious distribution of each ancestry, this final term gives the predicted distribution of each religion in each area. Second, I take a weighted average of this expression across religions within an ancestry, to measure the predicted religious density for any ancestral group n in an area a . That is, I am modeling directly the outcome of an ancestral group in an area as a function of the shared religious preferences of other ancestral groups that reside in an area. So Italians that live with Poles and Irish will have a high value; Italians that live Swedes and French will have a low value. The Appendix provides a detailed example to illustrate this methodology.

This approach solves the problems of measurement error noted above, since this predicted religious density is not based on actual stated religious preferences, but rather on the ethnic makeup of the area. By including own ancestral density and age/sex density in the regression, I also address the concern about identifying solely the impacts of majority status. Of course, I am only including own ancestry density linearly in the regression, while it enters nonlinearly into the computation of PRDEN, so this linear control may not be sufficient. To address this, I actually exclude one's own ancestral group when I compute PRDEN. That is, PRDEN measures, given your area and ancestral group, the likelihood that you live with others of your same religion, based on the other (but not own) ancestral groups in that area.

Both the GSS and Census ask respondents for a single ancestry, and, if that is not possible, for a primary and secondary ancestry. I define the primary ancestry as either the single

ancestry or the primary ancestry given. To address issues of segregation, I only use data on those ancestries from Europe or Canada in the analysis. That is, I exclude from the analysis any individual whose ancestry is not European or Canadian.⁴

To measure ancestral densities by area as precisely as possible, I use the STF 3A files from the census. These files provide, for each county, the demographic characteristics for the full sample of individuals who filled out the census long form. I use the counties in the STF 3A files to build up the MSAs in the census, and the PSUs in the GSS. For each MSA and PSU, I compute the share of the area that is of each ancestral group, and the share of the area that is of each age/sex group. To measure the religious distribution of each ancestral group, I use the GSS data to compute the percentage of each ancestry that is of each of the seven religious categories used.⁵

Empirical Strategy - Outcomes

While the GSS data are ideal for examining effects on religious participation, they are not ultimately useful for modeling the impact of religiosity on outcomes in this market structure approach. The sample is sufficiently small that the estimates from doing so are very imprecise.

⁴I divide non-hispanic white GSS respondents into 25 ancestral groupings: Austrian; Belgian; French Canadian; Other Canadian; Czech; Danish; Dutch; English; Finnish; French; German; Greek; Hungarian; Irish; Italian; Lithuanian; Norwegian; Polish; Portugese; Romanian; Russian; Scottish; Swedish; Swiss; and Yugoslavian. This is the largest set of ancestries that can be jointly defined in the GSS and the IPUMS files.

⁵It is possible that it is not just the religious preferences of ancestral groups, but also their degree of religious participation that matters. I have created a measure that accounts for this alternative, by weighting each ancestral group by their frequency of religious attendance in making PRDEN. This has no effects on the results.

Therefore, to model outcomes, I turn to the IPUMS data. Of course, in IPUMS, there is no information on religious preference. But, there is data on ancestry, akin to the data provided in the GSS. And the instrument in the GSS, PRDEN, does not ultimately rely on religious preference, but rather only on ancestry and area. Thus, I use the data the religious preferences of each ancestry from the GSS, and on ancestral densities by MSA from the STF 3A data, to compute PRDEN for each ancestry/MSA combination in the Census. That is, one benefit of this identification strategy is that I need not have actual religious preference data; having information on ancestry is sufficient to identify the model. The Census regression is therefore just the reduced form version of the GSS instrumental variables regression. These regressions model outcomes (e.g. income) as a function of whether individuals live near other ancestries that share their religious preference.

The relevant level of variation for the IPUMS data is at the ancestry/MSA level. Thus, I collapse these data into ancestry/MSA cells for the analysis, and weight the regression by the number of observations in the cell.

In the context of the outcomes analysis, there are two plausible interpretations of the coefficient β . The first is that an increase in market density of one's religion in one's area changes one's level of religious participation. The second is that those who are more religious move to areas where there is a higher density of their religion. In terms of finding a valid instrument for the impacts of religious participation on outcomes, the exact causal mechanism is not relevant; what is critical is simply that both mechanisms are otherwise independent of outcomes. But this may not be true with the selection mechanism. In particular, it may be that those who move to areas where there is a higher density of their religion (or those who don't

move out of such areas) are different, on average, than those who do not. This could arise, for example, if everyone desires to live with those of their religion, but moving is costly, so only higher income individuals can undertake such a move on the margin. In this case, there would be a systematic bias towards finding a positive correlation between religious market density and income.

Of course, the opposite story is also plausible, with the lowest income individuals being the ones who move to be with others of their religion. Indeed, based on the past literature on ethnic agglomeration, this opposite story seems the natural alternative. As reviewed in Edin et al. (2003), past studies tend to find that those who live with others of their same ethnic group are worse off than those who are not “ghetto-ized” in this way. Edin et al. propose a novel solution to this problem, using the random assignment of refugee groups in Sweden, to show that there is a true positive causal effect of ethnic agglomeration on outcomes. But, in a cross-section, the past literature would suggest a negative selection bias to the estimates. Thus, the effect of potential selection bias on the estimates is not clear. I discuss below an approach to addressing the sign of this bias using individuals who do and do not move from their state of birth.

Means

The means from the GSS and IPUMS data sets, at the PSU or MSA level, are presented in Table 1. In the GSS, the typical respondent lives in an area where roughly one-quarter of the population shares his general religious preference, although the variance is large. The instrument, PRDEN, has a slightly lower mean of 0.195, with a much smaller standard

deviation.⁶ Roughly one in six persons in the typical persons area shares their same ancestry, and fewer than one in twenty shares the same age/sex grouping. The mean of the attendance index (on a scale from 0, no attendance, to 8, attend more than once per week) is 3.83. In the IPUMS data, the mean and variance of PRDEN is similar to the GSS, as are ancestral densities; the age/sex density is somewhat higher.

Part III: Results

GSS Results

Table 2 presents the results for the impacts of market density on religious attendance in the GSS. The Table shows only the coefficients of interest, those on RELIGDENS, ANCESDENS, and ASDENS, from regressions that include all the controls shown in equation (1) above. The first column shows the OLS relationship, which is strongly positive: increases in the density of one's religion in one's area increases the frequency of religious attendance. The point estimate suggests that each ten percentage point increase in religious density is associated with a 0.013 unit increment in the attendance index. Recall that the attendance index is roughly a log scale, so that one unit represents a doubling in frequency of attendance. This implies that a 10 percentage point rise in religious density raises religious attendance by 1.3%.

The second column shows the impact of instrumenting religious density using PRDEN the prediction for religious density based on area ancestral mix. Doing so substantially increases the coefficient to 0.845; that is, each 10 percentage point increase in density raises attendance by

⁶The mean of PRDEN is lower than the mean of religious density since own ancestral group is excluded in computing PRDEN.

0.085 index point, or 8.5%.

The next two rows show the impact of area ancestral density and age/sex density on religious attendance. There is a very modest and statistically insignificant effect of ancestral density; having more individuals around of your ancestry group does not much impact attendance. This is striking, given the strong IV results in the second column, which are based on the density of other ancestral groups in the area. That is, living with others of your own ethnic group has little effect on religious attendance, but living with others of complementary ethnic groups that share your religion has a large positive effect.

This finding suggests that there are particular effects of having more individuals around of your own ancestry, as opposed to “complementary” ancestries that share your religion, that offset the positive effects of religious density. For example, for Italians, having more Poles and more Italians in the area increases the density of Catholics and thereby increases participation. But the additional Italians also reduce participation, perhaps because there are other cultural functions that are served outside of the religious setting. This is consistent with the census outcomes results which show negative effects of higher own ethnic density in the area. There is a much stronger effect of age/sex density; having more individuals in your age/sex group significantly increases the rate of religious participation.

One concern with the results presented thus far is that some individuals in the GSS report their religion as “none”. While about one-third of this group does report attending religious services, it is not clear if they should be included in the regression. I exclude them in the bottom panel, and the results are very similar to those reported above.

An additional concern is that the measure of religious participation is a crude one. In

Table 3, I therefore explore two alternative measures. First, I replace the index values with their time-weighted counterparts.⁷ As the first row of Table 3 shows, the effect remains significant with this alternative formulation, implying that each 10 percentage point increase in density leads to 0.95 more days of church attendance. Second, I estimate separate IV (linear probability) models for each of the individual attendance category dummies. While few of these coefficients are significant, there is a clear and consistent pattern of reductions in the low attendance categories, and increases in the high attendance categories. There is a particularly strong reduction in the odds of never attending, and a particularly strong increase in the odds of attending at least weekly.

IPUMS Results

Table 4 presents the results for outcomes using the IPUMS data, which is, as noted, collapsed to the ancestry/MSA level. Each row shows the coefficients on PRDEN (religious density for own ancestry, predicted based on own ancestries religious distribution and the religious distribution of other ancestries in the area), as well as the density of own ancestry and own age/sex group in the MSA.

The results show a very strong positive correlation between predicted own religious density and positive economic outcomes. The coefficient shows that each ten percent increase in predicted own religious density leads to a 0.91 percent increase in household income.

Normalized by the poverty scale, a ten percent increase in predicted own density leads to a 0.55

⁷That is, I convert the values as follows: no attendance = 0; less than annual attendance = 1; once or twice a year = 2; several times a year = 4; monthly = 12; two or three times per month = 30; almost every week = 40; weekly = 52; two or three times per week = 130

percent increase in income as a fraction of poverty. The lower result in the second row implies that density is positively associated with family size, as we will see to be the case below.

The results from the GSS analysis can be used here to form an implicit IV estimate of the effects of religious attendance on outcomes. The coefficient on PRDEN in a reduced form model of attendance in the GSS is 1.077 (0.351); that is moving from a PRDEN of 0 to a PRDEN of 1 doubles the frequency of attendance (a one unit increase in the attendance index). Putting this together with the census results, it implies that doubling the frequency of attendance leads to a 9.1% increase in household income, or a rise of 5.5% as a fraction of the poverty scale.

This implicit IV estimate is an overstatement, however, to the extent that there are other aspects of religiosity that are affected by religious density, and affect outcomes. For example, more religious density may lead not only to more attendance, but also more fervent religious beliefs, which themselves impact outcomes. Indeed, in the GSS, there are strong positive effects of density on the intensity of religious beliefs (IV coefficient of 0.133 (0.054)), and on belief in the afterlife (IV coefficient 0.128 (0.076)). These effects may be an outcome of increased participation, or they may be a direct impact of being in an environment where more individuals share your religion. To the extent that the latter is the case, then the implicit IV calculations presented here are overstating the effects of attendance per se (since they also build in the effects of religious beliefs as well).

The coefficient on ancestry density in these regressions is negative, suggesting that a higher density of own ancestry leads to lower, not higher, incomes. This is consistent with the literature on ethnic concentration reviewed earlier, which reported that, in standard observational models, ethnic segregation is associated with worse outcomes. Once again, this makes it even

more striking that I find positive effects of PRDEN: while being near others of your own ancestral group is negatively associated with outcomes, being near others of ancestral groups that share your religion is positively associated with outcomes. Being near others of the same age/sex group is also negatively associated with outcomes.

The next five rows explore the impact of predicted religious density on education. I first examine total years of education, and find a highly significant positive effect, whereby each 10 percentage point increase in predicted religious density is associated with 0.05 more years of education. Thus, once again using the result from the GSS, this finding suggests that doubling religious participation raises education by roughly 0.5 years. This result arises from large reductions in the odds of being a high school dropout or only having a high school degree, and increases in some college attendance or (particularly) college graduation.

The next two rows examine effects on own and spousal employment. There is no noticeable effect on the odds that respondents are employed, although spousal employment is significantly higher; this effect likely arises through the effects on marriage noted below (since spousal employment is zero either if spouses do not work or individuals are unmarried). The odds of the family receiving welfare income are significantly reduced; this effect is quite large, implying (given the GSS findings) that doubling the rate of attendance is associated with a 0.4% reduction in the odds of receiving welfare, or 16% of baseline receipt rates. There is a negative, but only marginally significant, reduction in the odds of being disabled.

The next two rows examine marital status. There is a significant positive association between predicted religious density and the odds of being married, and a significant negative association with the odds of being divorced. The former implies that doubling the rate of

attendance leads to a 4.4% increase in the odds of being married; the latter implies that this same doubling leads to a 4.0% reduction in the odds of being divorced. Finally, there is a strong positive association with the number of children, which is consistent with the rise in the odds of marriage.

Overall, these IPUMS results are very consistent with a positive effect of religious participation on economic measures of well-being. Higher religious density, which is associated with more religious attendance, is also associated with higher incomes, more education, higher family labor supply, lower rates of welfare receipt, a higher odds of marriage, and a lower odds of divorce.

Part IV: Alternative Interpretations

The results thus far have shown that there is a very strong positive correlation between religious market density, as proxied by ancestral densities in an area, and both religious attendance and positive census outcomes. Given the striking nature of these findings, I turn now a host of specification tests designed to explore the robustness of my finding to alternative explanations.

Before doing so, it is important to recall the wide set of alternative interpretations that are already ruled out by the results to this point. The models estimated above include area fixed effects, so that any explanation based on general differences across areas is controlled for in the analysis. That is, these findings do not simply indicate that more religious areas have better outcomes, but rather that outcomes are better for groups that have more individuals of their religion (through complementary ancestries) in their area, relative to other groups in that same

area with fewer individuals of their religion.

The model includes own ethnic density (which is also excluded from the computation of PRDEN), so that any explanation based on within-ethnic group spillovers is controlled for in the analysis. Moreover, the coefficient on own ethnic density is generally opposite-signed from the coefficient on PRDEN. The only alternatives which can explain these results, therefore, are ones where the density of *other* ancestral groups that share one's religion, but not the density of one's *own* ethnic group, cause positive outcomes.

Spillovers from "Nearby" Ancestral Groups

One such alternative is that, even though I control for own religious density, there are other spillover effects from "nearby" ancestry groups. That is, even controlling for the number of Swedish in an area, the number of Norwegians in the area may affect behavior in manners beyond religious spillovers, e.g. through other sources of Scandinavian group interaction.

I address this in two ways. First, just as I control for own group density in the model, I can also control for density of groups in your region of the world. I divide the ancestral groups into four for this exercise: Commonwealth/English-speaking nations; Northern Europe; Western Europe; and Eastern Europe. The results of including this control in the model are shown in the third column of Table 2 for the GSS. Adding this control has little effect on the estimated coefficient on religious density. The control itself is negative, but insignificant. The effect of adding this control in the IPUMS is shown in the second and third columns of Table 5; the first column provides the base case results from Table 4 for comparison. It is clear that adding this control has little substantive effect on the results; if anything, they become somewhat stronger.

The area density itself is typically opposite signed from PRDEN, which is similar to the own ancestry coefficient.

A more rigorous approach is to define religious density exclusive of all ancestries in one's region of the world. That is, I recompute density based only on that set of ancestries that are not in a given ancestry's region (akin to excluding one's own ancestry from the original computation of PRDEN). This provides a measure of density which is completely independent of not only of a group's own ancestry, but also of the ancestries of all (geographically) nearby groups.

The results of this exercise are shown in the final column of Table 2 (for GSS), and the fourth column of Table 5 (for IPUMS). Once again, even this more rigorous approach to dealing with potential spillovers has little impact on the findings in either data set.

Do Results Reflect Religious Participation Only?

The results thus far indicate fairly strongly that living near complementary ancestral groups is associated with a higher rate of religious participation, and with better economic outcomes. These findings, however, do not necessarily establish that the mechanism translating higher complementary ancestral density to outcomes is through religious participation. It may be, for example, that higher complementary ancestral density leads to more civic participation in general, not only through religious participation.

This alternative seems unlikely, since the key regressor is based on groups that share a group's religious makeup, not other factors that drive civic participation. Nevertheless, I can explore this alternative in the GSS, by using the available data on other measures of "social

capital”. The GSS collects data on whether the respondent is a member of one of a set of organizations. I use as a dependent variable both a dummy variable for belonging to *any* organization, and the count of the number of organizations to which the individual belongs. If PRDEN is simply capturing some unobserved source of social capital in an area, then we would expect these coefficients to be positive, as we see for attendance.

In fact, as Table 6 shows, both coefficients are negative, although neither is significant. I also show in this table the means and number of observations for these variables. These estimates are fairly imprecise, relative to the sample means. But there is certainly no evidence of increases in other social capital. Indeed, if I estimate separate models for every group membership, the only two coefficients that are positive are for membership in a professional organization and membership in a church group (although neither is significant); the other 12 membership coefficients are all negative.

Are Results Driven by Selection?

As noted earlier, there are two possible explanations for the findings presented thus far: that the evolution of density in one’s area changes religious participation and outcomes; and that those who are more religious move to areas with a higher density of their religion. And, as noted, in terms of finding a valid instrument for the impacts of religious participation on outcomes, the exact causal mechanism is not relevant; what is critical is simply that both mechanisms are otherwise independent of outcomes. But there is reason to believe that the selection mechanism is not independent of outcomes. As noted, the bias is not obvious *ex ante*

It is difficult to definitively rule out selection as an explanation for these results, but there

is one quite suggestive test: examining the results for movers. An implication of the selection explanation for these results is that they should be found more strongly among those who are moving into an area than for those already living in the area.

An extreme thought experiment is helpful here. Suppose that individuals in the U.S. were assigned randomly to a place to live many generations ago, and were not allowed to move. Suddenly, for the current generation interviewed in the 1990 Census, individuals were allowed one move at age 25. In this case, there would be a clean test for selection bias. Among movers, any measured effect on outcomes at age 26 (immediately after moving is allowed) of an increase in religious density, representing an improved match, would be pure selection. That is, a finding that those who move to places with a better religious match have better outcomes could not reflect the causal effect of the better match, since outcomes at age 26 are determined solely by place of birth. Rather, such a finding would reflect the fact that higher quality individuals are moving to where the match is better. In other words, if the effect of religious density is truly causal, then any change in religious density for movers should be irrelevant to their outcomes, since their outcomes at age 26 should be fully determined by their place of birth. If an increase in religious density is associated with better outcomes, it occurs through selection.

Of course, we do not have this type of clean experiment. In this section, I approximate such a test by examining the correlation between outcomes and the change in religious density for those who move from their state of birth to some new state. Such an approximation has two offsetting biases relative to the thought experiment. First, if all selection has occurred in previous generations (something ruled out by assumption in the thought experiment), then it will not be reflected in a test on movers in the current generation, biasing the test statistic towards

zero. So long as some selection is ongoing with the current generation, however, the test will still provide the correct sign, if not magnitude. Second, since movers have moved at some point since birth, if there is a true positive causal effect of religious density, then the improved match may be reflected in their ex-post outcomes even absent selection. This will bias the test towards finding selection, and we cannot fix this problem since we don't know the age that individuals moved.

To implement this test, I measure PRDEN at the state, rather than the MSA level. I then reestimate the census models just for individuals who have moved from their state of birth, modeling their outcomes as a function of the *change* in PRDEN between their state of birth and their current state. The positive selection story would predict a strong positive coefficient on this change: those who move to improved matches (higher PRDEN) would be those with the best outcomes. I include in this model the change in ancestry and age/sex density between state and birth and state of residence, as well as a full set of fixed effects for both state of birth and state of residence.

The result of these regressions is shown in the first column of Table 7. Strikingly, the signs for virtually all of the regressions are now the opposite of those from the earlier regressions: those who move from their birth state to another state with a higher PRDEN have *worse* outcomes along most dimensions. There are significant negative coefficients for income, education and employment, and significant positive coefficients for disability and welfare receipt. The results for marriage and divorce are also opposite signed (negative effects on marriage, positive on divorce), although not statistically significant. This suggests that there is actually negative selection among movers: those who move to better matches in terms of religious density

are actually worse on average than those who do not.

To confirm this point, the second column in Table 7 shows the results (at the state level) for those who remain in their state of birth. Here, the results are generally positive and larger than those found earlier. The exceptions is marriage, where the effect is not significant, divorce, where the effect is significant but weaker than in Table 4, and number of children, where the effect is now negative. Thus, consistent with the literature on ethnic segregation, there appears to be negative selection at work among movers that leads us to understate the overall positive effect of religious density on outcomes.

Part IV: Conclusions

Religion remains an important aspect of life in the U.S. Yet we know very little about the impacts the religious participation has on economic outcomes. I have attempted to remedy this shortcoming by instrumenting for religious density with the extent to which others in an area share one's religious heritage, and to study the impact of religious density on outcomes in the Census IPUMS data.

The findings of this analysis are striking: a higher density of your religion in your area, as proxied by the ancestral mix of area residents, leads to significantly more religious participation, and to better outcomes along a variety of dimensions, such as education, income, and marital status. These effects are sizeable, and are robust to a variety of specification checks. And they do not appear to be driven by selection of higher ability individuals into areas where there is more density of their religion.

This paper has not at all investigated the mechanisms behind these findings. There are

several candidate explanations that could be explored in future work. One is social capital effects: being near those of one's religion could increase the total number of social interactions, and this can have positive effects of the kind emphasized by the social capital literature (e.g. Glaeser, Laibson and Sacerdote, 2000; Putnam, 2000). Given the negative effects of being near one's own ethnic group, however, this must be a particular channel of social capital that works through interactions in religious settings. A second and related channel could be insurance provision. Religious institutions may operate as both financial and emotional insurers during negative shocks, improving overall welfare through filling missing markets for capital or other forms of support. This interpretation is supported in a developing country context by Chen (2003), who finds that those most affected by the Indonesian financial crisis of the 1990s were most likely to increase their degree of religious intensity, but that this effect was mitigated by the availability of other forms of credit.

A third channel could be through religious education: a higher religious density could lead to more attendance at religious schools, and this could have a positive effect on long run outcomes. Neal (1997) shows that Catholic children are more likely to attend Catholic schools if there are more Catholics nearby. Whether this leads to better outcomes is a source of some contention in the literature, as Altonji, Elder and Taber (2002) review.

Finally, it is possible that more religious faith improves well-being directly. For example, those with more faith may be less "stressed out" about daily problems that impede success in the labor market and the marriage market, and are therefore more successful. Tests which could separate out this and the other channels would be quite useful for explaining the results in this paper.

References

- Altonji, Joseph, Todd Elder and Christopher Taber (2002). "An Evaluation of Instrumental Variables Strategies for Estimating the Effects of Catholic Schools," NBER Working Paper #9358.
- Bachman, Jerald, Patrick O'Malley, John Schulenberg, Lloyd Johnston, Alison Bryant, and Alicia Merline (2002). *The Decline in Substance Use in Young Adulthood: Changes in Social Activities, Roles, and Beliefs*. Mahway, New Jersey: Lawrence Erlbaum Associates.
- Chaves, Mark and Philip Gorski (2001). "Religious Pluralism and Religious Participation," *Annual Review of Sociology*, 27, 261-281.
- Chen, Daniel (2003). "Economic Distress and Religious Intensity: Evidence from Islamic Resurgence During the Indonesian Financial Crisis," mimeo, MIT.
- Edin, Per-Anders, Peter Fredriksson, and Olof Aslund (2003). "Ethnic Enclaves and the Economic Success of Immigrants - Evidence from a Natural Experiment," *Quarterly Journal of Economics*, 329-357
- Ellison, Christopher (1991). "Religious Involvement and Subjective Well-Being," *Journal of Health and Social Behavior*, 32, 80-99.
- Ellison, Christopher, J.A. Burr, and P.L. McCall, "Religious Homogeneity and Metropolitan Suicide Rates," *Social Forces*, 76, 273-299.
- Evans, T. David, Francis Cullen, R. Gregory Dunaway, and Velmer Burton (1995). "Religion and Crime Reexamined: The Impact of Religion, Secular Controls and Social Ecology on Adult Criminality," *Criminology*, 33, 195-215.
- Finke, Roger and Rodney Stark (1988). "Religious Economics and Sacred Canopies," *American Sociological Review*, 53, 41-49.
- Finke, Roger and Rodney Stark (1992). *The Churching of America, 1776-1990: Winners and Losers in Our Religious Economy*. New Brunswick, NJ: Rutgers University Press.
- Freeman, Richard (1986). "Who Escapes? The Relation of Churchgoing and Other Background Factors to the Socioeconomic Performance of Black Male Youths from Inner-City Tracts," in Richard Freeman and Harry Holzer, eds. *The Black Youth Employment Crisis*. Chicago: University of Chicago Press, 353-76.
- Glaeser, Edward, David Laibson and Bruce Sacerdote (2000). "The Economic Approach to

- Social Capital”, NBER Working Paper #7728.
- Gruber, Jonathan (2003). “Pay or Pray: The Impact of Charitable Subsidies on Religious Participation,” mimeo, MIT.
- Gruber, Jonathan and Sendhil Mullainathan (2002). “Do Cigarette Taxes Make Smokers Happier?,” NBER Working Paper #8872, April 2002.
- Hout, Michael and Andrew M. Greeley (2003). “Religion and Happiness,” paper prepared for the annual meeting of the American Sociological Association.
- Hull, Brooks, and Frederick Bold (1995). “Preaching Matters - Replication and Extension,” *Journal of Economic Behavior and Organizations*, 27, 143-149.
- Hummer, Robert, Richard Rogers, Charles Nam, and Christopher Ellison (1999). “Religious Involvement and U.S. Adult Mortality,” *Demography*, 35, 273-285.
- Iannaccone, Laurence (1991). “The Consequences of Religious Market Structure: Adam Smith and the Economics of Religion,” *Rationality and Society*, 3, 156-177.
- Iannaccone, Laurence, Roger Finke and Rodney Stark (1997). “Deregulating Religion: The Economics of Church and State,” *Economic Inquiry*, 35, 350-364.
- Iannaccone, Laurence (1998). “Introduction to the Economics of Religion,” *Journal of Economic Literature*, 36, 1465-1496.
- Lehrer, Evelyn and Carmel Chiswick (1993). “Religion as a Determinant of Marital Stability,” *Demography*, 30, 385-404.
- Levin, Jeffrey S. (1994). “Religion and Health: Is There an Association, Is it Valid, and Is It Causal?,” *Social Science and Medicine*, 38, 1475-1482.
- Levin, Jeffrey S. and Harold Y. Vanderpool (1987). “Is Frequent Religious Attendance *Really* Conducive to Better Health? Toward an Epidemiology of Religion,” *Social Science and Medicine*, 24, 589-600.
- Lipford, Jody, Robert McCormick, and Robert Tollison (1993). “Preaching Matters,” *Journal of Economic Behavior and Organizations*, 21, 235-250.
- Olson, Daniel (1998). “The Influence of Religious Pluralism on Close Social Ties and Religious Involvement”. Paper presented at the Annual Meeting of the Society for the Scientific Study of Religion.

- Perl, Paul, and Daniel Olson (2000). "Religious Market Share and Intensity of Church Involvement in Five Denominations," *Journal for the Scientific Study of Religion*, 39, 12-31.
- Phillips, Rick (1998). "Religious Market Share and Mormon Church Activity," *Sociology of Religion*, 59, 117-130.
- Putnam, Robert (2000). *Bowling Alone: The Collapse and Revival of American Community*. New York: Simon and Schuster.
- National Center on Addiction and Substance Abuse (2001). *So Help Me God: Substance Abuse, Religion and Spirituality*. New York: Columbia University.
- Neal, Derek (1997). "The Effects of Catholic Secondary Schooling on Educational Attainment," *Journal of Labor Economics*, 15, 98-123.
- Roof, Wade Clark, and William McKinney (1987). *American Mainline Religion: Its Changing Shape and Future*. New Brunswick, NJ: Rutgers University Press
- Tomes, Nigel (1985). "Religion and the Earnings Function," *American Economic Review*, 75, 245-250.
- Voas, David, Daniel Olson and Alasdair Crockett (2002). "Religious Pluralism and Participation: Why Previous Research is Wrong," *American Sociological Review*, 67, 212-230.
- Wallace, John, and David Williams (1997). "Religion and Adolescent Health-Compromising Behavior," in J. Schulenberg, J.L. Maggs, and K. Hurrelman, eds., Health Risks and Developmental Transitions During Adolescence. Cambridge, UK: Cambridge University Press, 444-468.
- Zaleski, Peter and Charles Zech (1995). "The Effect of Religious Market Competition on Church Giving," *Review of Social Economy*, 53, 350-367.

| Table 1: Means | | |
|----------------------------|--------|---------------|
| | Mean | Standard Dev. |
| GSS | | |
| Religious Density | 0.266 | 0.164 |
| PRDEN | 0.195 | 0.029 |
| Ancenstral Dens. | 0.187 | 0.155 |
| Age/Sex Dens. | 0.030 | 0.012 |
| Area Dens. | 0.361 | 0.171 |
| Attend Index | 3.83 | 2.73 |
| Number Obs | 14,976 | |
| Census | | |
| PRDEN | 0.197 | 0.029 |
| Ancestral Dens. | 0.190 | 0.145 |
| Age/Sex Dens. | 0.077 | 0.013 |
| Area Dens. | 0.332 | 0.141 |
| HH Income | 50,560 | 10,470 |
| Percentage of Poverty Line | 3.578 | 0.354 |
| Years of Education | 13.3 | 0.721 |
| HS Dropout | 0.132 | 0.060 |
| HS Graduate | 0.323 | 0.071 |
| Some College | 0.294 | 0.059 |
| College Graduate | 0.248 | 0.093 |
| Employed | 0.652 | 0.066 |
| Spouse Employed | 0.411 | 0.053 |
| Receive Welfare | 0.024 | 0.013 |
| Disabled | 0.117 | 0.029 |

| | | |
|----------------|-------|-------|
| | 32 | |
| Married | 0.631 | 0.050 |
| Divorced | 0.103 | 0.026 |
| Number of Kids | 0.656 | 0.136 |
| Number Obs | 5023 | |

Notes: Table shows means and standard deviations from GSS and IPUMS data sets described in text.

| Table 2: Attendance Results in GSS | | | | |
|--|------------------|-------------------|-------------------|-------------------|
| | OLS | IV | IV, Area Control | IV, Exclude Area |
| All Observations | | | | |
| Religious Density | 0.125 (0.039) | 0.845 (0.278) | 0.911 (0.281) | 0.951 (0.288) |
| Ancestry Density | 0.020 (0.031) | -0.054 (0.042) | -0.016 (0.050) | -0.065 (0.043) |
| Age/Sex Density | 0.282 (0.123) | 0.243 (0.125) | 0.243 (0.126) | 0.236 (0.126) |
| Area Density | | | -0.101 (0.067) | |
| Number of Obs | 14,976 | 14,976 | 14,976 | 14,976 |
| Exclude Those with No Religious Preference | | | | |
| PRDEN | 0.101 (0.043) | 0.853 (0.282) | 0.928 (0.283) | 0.983 (0.290) |
| Ancestry Density | 0.023 (0.033) | -0.055 (0.045) | -0.010 (0.054) | -0.069 (0.045) |
| Age/Sex Density | 0.264 (0.132) | 0.258 (0.134) | 0.262 (0.134) | 0.256 (0.134) |
| Area Density | | | -0.119 (0.072) | |
| Number of Obs | 13,606 | 13,606 | 13,606 | 13,606 |

Notes: Tables shows coefficients of interest from regressions such as equation (1) in text; in addition to covariates shown, models include fixed effects for age/sex, ancestry, religious preference, PSU, and year. Each column in each panel shows coefficients from a single regression. Standard errors in parentheses. Top panel shows regressions including all observations; bottom panel excludes those who report no religious preference. First column is estimated by OLS; remaining columns by instrumental variables, where instrument is PRDEN. Third column adds control for % of individuals in PSU that are from individual's area of the world; fourth column recomputes PRDEN to exclude those individuals.

Table 3: Attendance Categories and Time Weighted Measure

| Dependent Variable | IV | IV, Exclude Area |
|---------------------------|-------------------|-------------------|
| Time-Weighted Attendance | 9.46 (3.73) | 11.40 (3.85) |
| Never Attend | -0.069 (0.038) | -0.060 (0.039) |
| Less than Once Per Yr. | -0.026 (0.031) | -0.024 (0.032) |
| Once/ Twice Per Year | -0.002 (0.038) | -0.026 (0.039) |
| Several Times Per Year | -0.026 (0.037) | -0.040 (0.038) |
| Once Per Month | -0.032 (0.028) | -0.025 (0.029) |
| 2 or 3 Times Per Month | 0.005 (0.03) | -0.008 (0.031) |
| Nearly Every Week | 0.018 (0.026) | 0.036 (0.026) |
| Every Week | 0.101 (0.045) | 0.108 (0.046) |
| Several Times Per Week | 0.029 (0.028) | 0.039 (0.029) |

Notes: Tables shows coefficients of interest from regressions such as equation (1) in text; in addition to religious density, models includes ancestral density, age/sex density, and fixed effects for age/sex, ancestry, religious preference, PSU, and year. Each cell shows coefficient from a separate regression. Standard errors in parentheses. First row shows regression where attendance index is replaced by time weighted values. Remaining rows show individual linear probability models of attendance frequency categories. First column shows basic PRDEN results; second column shows results when PRDEN excludes individuals from that individual's part of the world.

Table 4: Census Outcome Results

| | PRDEN | Ancestry Density | Age/Sex Density |
|--------------------------------------|---------------------|--------------------|--------------------|
| Log of Household Income | 0.091 (0.012) | -0.032 (0.001) | -0.000 (0.014) |
| Log of Percentage of Poverty Line | 0.055 (0.007) | -0.013 (0.001) | -0.027 (0.008) |
| Years of Education | 0.499 (0.064) | -0.300 (0.008) | -0.302 (0.075) |
| HS Dropout | -0.039 (0.006) | 0.022 (0.001) | 0.011 (0.007) |
| HS Graduate | -0.075 (0.007) | 0.021 (0.001) | 0.031 (0.008) |
| Some College | 0.025 (0.006) | -0.009 (0.001) | 0.002 (0.006) |
| College Graduate | 0.088 (0.008) | -0.035 (0.001) | -0.043 (0.010) |
| Employed | 0.003 (0.005) | -0.003 (0.001) | -0.012 (0.005) |
| Spouse Employed | 0.032 (0.005) | -0.000 (0.001) | 0.014 (0.006) |
| Receive Welfare | -0.0044 (0.0018) | 0.0007 (0.0002) | 0.0056 (0.0021) |
| Disabled | -0.007 (0.004) | 0.002 (0.0005) | -0.012 (0.004) |
| Married | 0.044 (0.006) | 0.014 (0.0006) | 0.044 (0.006) |
| Divorced | -0.030 (0.004) | -0.002 (0.0005) | -0.015 (0.004) |
| Number of Kids | 0.091 (0.011) | 0.007 (0.001) | 0.434 (0.013) |

Notes: Tables shows coefficients of interest from regressions such as equation (1) in text; in addition to covariates shown, models include fixed effects for age/sex, ancestry, religious

preference, PSU, and year. Each row shows coefficients from a separate regression, where dependent variable is outcome listed in first column. Standard errors in parentheses.

Table 5: Census Specification Checks

| | Base Case | Control for Area Density | | Exclude Own Area |
|-----------------------------------|---------------------|--------------------------|--------------------|---------------------|
| | PRDEN | PRDEN | Area Density | PRDEN |
| Log of Household Income | 0.091 (0.012) | 0.101 (0.012) | -0.019 (0.002) | 0.094 (0.011) |
| Log of Percentage of Poverty Line | 0.055 (0.007) | 0.061 (0.007) | -0.011 (0.001) | 0.057 (0.006) |
| Years of Education | 0.499 (0.064) | 0.556 (0.065) | -0.103 (0.013) | 0.525 (0.060) |
| HS Dropout | -0.039 (0.006) | -0.044 (0.006) | 0.010 (0.001) | -0.035 (0.006) |
| HS Graduate | -0.075 (0.007) | -0.075 (0.007) | 0.0004 (0.0014) | -0.069 (0.006) |
| Some College | 0.025 (0.006) | 0.025 (0.006) | -0.000 (0.001) | 0.013 (0.005) |
| College Graduate | 0.088 (0.008) | 0.095 (0.008) | -0.011 (0.002) | 0.092 (0.008) |
| Employed | 0.003 (0.005) | 0.002 (0.005) | 0.001 (0.001) | 0.001 (0.004) |
| Spouse Employed | 0.032 (0.005) | 0.031 (0.005) | 0.001 (0.001) | 0.026 (0.005) |
| Receive Welfare | -0.0044 (0.0018) | -0.0050 (0.0018) | 0.0010 (0.0004) | -0.0044 (0.0017) |
| Disabled | -0.007 (0.004) | -0.008 (0.004) | 0.002 (0.0008) | -0.007 (0.004) |
| Married | 0.044 (0.006) | 0.042 (0.006) | 0.003 (0.001) | 0.042 (0.005) |
| Divorced | -0.030 (0.004) | -0.032 (0.004) | 0.002 (0.001) | -0.028 (0.003) |
| Number of Kids | 0.091 (0.011) | 0.089 (0.012) | 0.003 (0.002) | 0.086 (0.011) |

Notes: Tables shows coefficients of interest from regressions such as equation (1) in text; in addition to covariates shown, models include ancestral density, age/sex density, and fixed effects for age/sex, ancestry, religious preference, PSU, and year. Standard errors in parentheses. First column shows PRDEN coefficients from Table 4, for coefficient of interest; each row is a separate regression. Second and third columns show PRDEN and area density coefficients from regressions that adds control for % of individuals in PSU that are from individual's area of the world; fourth column recomputes PRDEN to exclude those individuals.

Table 6: GSS Results for Other Participation

| Measure of Participation | IV Coefficient |
|--------------------------|---------------------------------------|
| Member of Any Group? | -0.056 (0.089) [0.756] 6,112 |
| Number of Groups | -0.629 (0.503) [1.98] 6,112 |

Notes: Tables shows coefficients of interest from regressions such as equation (1) in text; in addition to religious density, models includes ancestral density, age/sex density, and fixed effects for age/sex, ancestry, religious preference, PSU, and year. Each cell shows coefficient from a separate regression, where dependent variable is listed in first column. Standard errors in parentheses, variable mean in square brackets, and number of observations at bottom.

Table 7: Examining Selection

| | Movers out of Birth State | State of Residence = State of Birth |
|-------------------------|---------------------------|--|
| Log of Household Income | -0.050 (0.011) | 0.219 (0.030) |
| Log % of Poverty Line | -0.010 (0.007) | 0.110 (0.017) |
| Years of Education | -0.354 (0.041) | 1.206 (0.130) |
| HS Dropout | 0.034 (0.004) | -0.101 (0.014) |
| HS Graduate | -0.001 (0.006) | -0.070 (0.010) |
| Some College | -0.005 (0.005) | 0.021 (0.008) |
| College Graduate | -0.029 (0.006) | 0.150 (0.015) |
| Employed | -0.011 (0.005) | 0.016 (0.007) |
| Spouse Employed | -0.010 (0.005) | 0.014 (0.007) |
| Receive Welfare | 0.0040 (0.0020) | -0.0055 (0.0027) |
| Disabled | 0.008 (0.003) | -0.019 (0.006) |
| Married | -0.008 (0.005) | 0.008 (0.007) |
| Divorced | 0.005 (0.003) | -0.026 (0.005) |
| Number of Kids | -0.017 (0.012) | -0.009 (0.016) |

Notes: Tables shows coefficients of interest from regressions such as equation (1) in text; in addition to covariates shown, models include ancestral density, age/sex density, and fixed effects

for age/sex, ancestry, religious preference, PSU, and year. Standard errors in parentheses. First column shows coefficient on change in PRDEN (state of residence minus state of birth PRDEN), and second column shows coefficient on PRDEN for those still in state of birth; each row is a separate regression.

Appendix: Illustrating PRDEN Computation

This appendix provides an illustrative example of the computation of the PRDEN measure that is the centerpiece of the empirical strategy used in this paper. I consider a very simplified example of two ancestries (Italian and English), two religions (Catholic and Protestant), and two areas.

As the first row of Appendix Table 1 shows, 80% of Italians are Catholic, but only 10% of English are. The first row in the second panel of the table shows, likewise, that 20% of Italians are Protestant, and 90% of the English are. In this example, we consider two fictional areas: 60% of the population in Area 1 is Italian, and 40% is English; in Area 2, 20% of the population is Italian and 80% is English.

This Table computes the terms in the last part of expression (2), by multiplying the ancestral religious distribution times ancestral population distributions in an area to obtain a predicted religious density. Thus, from these figures, we predict that 52% of Area 1 will be Catholic, with 48% of the population Protestant; 24% of Area 2 is predicted to be Catholic, and 76% is predicted to be Protestant. That is, since Area 1 is slightly more Italian, it is predicted to be slightly more Catholic; since Area 2 is much more English, it is predicted to be much more Protestant.

Appendix Table 2 then takes this information, and uses it to compute PRDEN, the predicted religious density of each area for each ancestry. In Area 1, for example, 52% of the population is predicted to be Catholic, and 80% of Italians are Catholic; 48% of the population is predicted to be Protestant, and 20% of Italians are Protestant. As a result, for Italians in Area 1, on average 52% of the population will be of their same religion. For the English in Area 1, on average 48% of the population will be of their same religion. In Area 2, only 34% of the population will be of the same religion as the Italians in that area, but 70% of the population will be of the same religion as the English in that area.

These results illustrate the sources of variation in PRDEN. The more highly segregated areas are into ancestral groups that are themselves religiously segregated, the larger will be the variance in PRDEN. So, there is much more variance in PRDEN in Area 2, since that location is more heavily segregated, and the dominant group (English) are strongly segregated into Protestantism.

Appendix Table 1

| Catholics | | | |
|----------------------------|----------|---------|-----------------------------|
| | Italians | English | Total Predicted Catholics |
| Ancestry Religion % | 80% | 10% | |
| Ancestry Density in Area 1 | 60% | 40% | 52% |
| Ancestry Density in Area 2 | 20% | 80% | 24% |
| Protestants | | | |
| | Italians | English | Total Predicted Protestants |
| Ancestry Religion % | 20% | 90% | |
| Ancestry Density in Area 1 | 60% | 40% | 48% |
| Ancestry Density in Area 2 | 20% | 80% | 76% |

Appendix Table 2

| | Italians | | | English | | |
|------------|------------------------|------------------|---------|------------------------|------------------|---------|
| | Predicted % in Area | % in Ancestry | Product | Predicted % in Area | % in Ancestry | Product |
| Area 1 | | | | | | |
| Catholic | 52% | 80% | 42% | 52% | 10% | 5% |
| Protestant | 48% | 20% | 10% | 48% | 90% | 43% |
| PRDEN | | | 52% | | | 48% |
| Area 2 | | | | | | |
| Catholic | 24% | 80% | 19% | 24% | 10% | 2% |
| Protestant | 76% | 20% | 15% | 76% | 90% | 68% |
| PRDEN | | | 34% | | | 70% |