Journal of Comparative Economics 48 (2020) 729–748

Contents lists available at ScienceDirect

Journal of Comparative Economics

journal homepage: www.elsevier.com/locate/jce

# Religious rules as a means of strengthening family ties: Theory and evidence from the Amish☆

James P. Choy

*Bureau of Economic Analysis, 4600 Silver Hill Road, Washington, DC 20233, United States*

ARTICLE INFO

*Keywords:*
Cooperation
Altruism
Family relationships
Cultural economics
Religion

*JEL classification:*
D13
D64
Z10
Z12

ABSTRACT

Religious institutions impose many rules on their members. I argue that a central function of these rules in many religious communities is to promote altruistic and cooperative behavior *within families*. My argument contrasts with the dominant view in the literature that the primary economic function of religious rules is to promote altruism and cooperation *between unrelated religious community members*. I develop my argument using a novel dataset on the Amish. My theory can explain selection patterns into the Amish church, Amish cultural persistence, persistent inequality between the Amish and non-Amish people, and high Amish fertility rates.

## 1. Introduction

Religion is one of the most important institutions in nearly all human societies. One way in which religious institutions affect economic outcomes is by imposing rules that affect religious community members' behavior. In this paper, I study the function and effects of these religious rules.

I argue that one of the main functions of many religious rules is to promote altruistic and cooperative behavior *within families*. According to my argument, religious rules complement and strengthen biological altruistic ties, generating levels of cooperation within families that are higher than would be expected on the basis of biological altruism alone. This function of religion is illustrated by the fifth commandment, "Honor thy father and thy mother," which directs children to behave more altruistically towards their parents. It is also illustrated by religious movements such as modern evangelical Protestantism that promote "family values" by imposing rules such as bans on abortion that affect family relationships.

My argument contrasts with the now standard view of religious rules in the literature, that the primary economic function of religious rules is to promote altruism and cooperation *between unrelated religious community members*. The classic statement of this argument in economics is the club goods model of religion, developed by Iannaccone (1992). Iannaccone focuses on seemingly irrational religious rules, such as prohibitions against eating certain foods or working on certain days of the week, and he argues that

☆ I am grateful to Robby Akerlof, Sascha Becker, Donald Kraybill, David McConnell, Andrew Oswald, Steven Nolt, Debraj Ray, and various seminar participants for helpful discussions. I thank David McConnell for data on Amish affiliations. The Centre for Competitive Advantage in the Global Economy (CAGE) at the University of Warwick provided financial support for this research. The views expressed in this paper are those of the author and do not necessarily represent the views of the Bureau of Economic Analysis.

*E-mail address:* james.choy@bea.gov.

https://doi.org/10.1016/j.jce.2019.12.007
Received 10 March 2019; Received in revised form 18 December 2019; Accepted 24 December 2019
Available online 03 January 2020

these seemingly irrational rules promote cooperation between unrelated religious community members.[1] On this view, religious institutions substitute for biological altruism by helping to create relationships between unrelated individuals that mimic biologically altruistic relationships. This function of religion is illustrated by the Christian admonishment, "Love thy neighbor as thyself," which directs people to behave more altruistically towards other unrelated members of the community. It is also illustrated by religious movements that focus on community solidarity and support for the poor, such as Catholic liberation theology, and by cooperative religious communes such as monasteries. It is interesting to note that organizations like monasteries often impose explicit anti-family rules, such as the requirement that members be celibate, in order to increase members' commitment to the community. This observation shows that promoting cooperation within families and promoting cooperation between unrelated community members are distinct and sometimes conflicting goals. However, it is also possible for religious communities to promote both goals simultaneously.

More generally, I argue that religious communities exist on a spectrum, with some communities being more concerned with promoting cooperation within families and other communities being more concerned with promoting cooperation between unrelated community members. Most religious communities are at least partially concerned with promoting both goals to some extent. Thus, I do not argue that club goods model is incorrect as a description of most religious communities. Instead, I argue that the club goods model is incomplete, since it fails to explain the interaction between religious rules and family life. My theory explains a variety of phenomena associated with religion that are not explained by the club goods model.

I develop my argument in the context of one unusual religious community, the Amish. One of the most distinctive features of Amish society is the many seemingly irrational rules that Amish communities impose on their members, such as rules forbidding members from driving cars, using electricity from the grid, or attending school past the age of fourteen. Amish communities also impose a punishment, called shunning, on children who leave the community to avoid following the rules. I argue that a primary purpose of Amish rules is to encourage children to allocate their time according to their parents' preferences by reducing the productivity of other uses of time such as wage work or private leisure. Thus, Amish rules effectively increase the level of children's altruistic and cooperative behavior towards parents. Parents have political control over the community, and they use their control to impose these rules to benefit themselves at the expense of their children. Amish parents then impose punishments on children who leave the community in order to force children to remain and to follow the rules. This could mean, for example, that Amish children allocate more time towards caring for their parents in old age instead of working for wages or enjoying private leisure activities. It could also mean that Amish children allocate more time towards raising their own children (their parents' grandchildren).

I use my model to explain a number of facts about Amish institutions. One set of facts that my model can explain are selection patterns into Amish communities. Using a novel data set covering nearly the entire Amish population of Holmes County, Ohio, over more than 20 years, I establish two results. First, the severity of the punishment of shunning imposed on children who leave the church varies across Amish communities. I show that children are less likely to leave Amish communities that practice stricter forms of shunning, even after controlling for other observable factors, and that the association between shunning and the probability of exit is very strong. This suggests that a central purpose of shunning is to prevent children from leaving the community, as predicted by my theory. Second, I show that children of richer Amish parents are more likely to remain in the church than children of poorer Amish parents. My explanation of this result is that because parents want their children to remain in the church, parents use their resources to influence their children's religious choice, for example by threatening to disinherit children who leave the church. The fact that parents use their resources to affect whether their children join the church illustrates the conflict of interest between parents and children about whether children should join the church.

My model can also explain the persistence of Amish church membership across generations, and persistent inequality between the Amish and their non-Amish neighbors. Amish church membership has been transmitted from parents to children for more than three hundred years. At the same time, Amish families remain persistently poorer than their non-Amish neighbors, in large part because Amish community rules reduce economic productivity. Not only are the Amish poorer than their neighbors, but they appear to have lower overall welfare, as evidenced by relatively high exit rates of Amish children born into communities that do not impose severe punishments on children who leave. In my model, community punishments for children who leave the church are the key to explaining both the transmission of Amish church membership across generations and persistent wealth and welfare inequality between the Amish and the non-Amish. Amish community punishments force children to remain in the church when they would have preferred to leave. When the children become parents themselves, they impose the same punishments on their children, explaining transmission of church membership across multiple generations. Family dynasties that initially belong to the Amish church are persistently poorer than other dynasties even after many generations, because Amish rules reduce children's productivity and welfare. If the costs that Amish rules impose on children outweigh the benefits to parents, the Amish can be persistently disadvantaged relative to the non-Amish in lifetime welfare, even after accounting for the benefits that Amish rules create for Amish parents.

Finally, my model can explain the extremely high Amish fertility rate. The fertility rate among the Amish is nearly eight births per woman. In my model, Amish rules encourage Amish community members to spend more time than they otherwise would on family activities such as child raising, which may help to explain the high Amish fertility rate.

In the main version of my model, I assume that rules encouraging Amish community members to spend more time raising children primarily increase the quantity, rather than the quality, of Amish children. This assumption is consistent with the high Amish fertility rate and the low level of Amish investment in children's schooling. An alternative assumption is that religious rules encourage religious community members to invest more in child quality, rather than child quantity. In this case, members of strict religious communities may be persistently wealthier and better educated than non-members. This possibility may help to explain religious

[1] See also Norenzayan et al. (2016) for more recent evidence supporting this point of view from the anthropological literature.

communities with persistently high levels of human capital such as non-ultra-Orthodox Jews (Botticini and Eckstein, 2012) and German Protestants (Becker and Woessmann, 2009)

It is possible to think of alternative explanations for each of the facts described above, taken individually. However, I believe that it is difficult to find a model of comparable simplicity and generality that explains all of the facts about Amish institutions. I now discuss the main alternative theories, and argue that they each fail to explain some or all of the facts about Amish institutions. I begin with the club goods model.

In the club goods model, the purpose of religious rules is not to encourage members to spend more time with their families, but rather to encourage members to spend more time on activities that benefit the entire community, such as organizing community events or engaging in charitable activities. The club goods model is probably part of the explanation for Amish institutions, as there are in fact high levels of cooperation between unrelated members of Amish communities. However, the club goods cannot be the whole explanation of Amish institutions, because it has difficulty explaining the facts described above. I discuss five ways in which Amish institutions are hard to explain in the club goods model.

First, the club goods model predicts that strict religious communities should impose punishments on people who *join* the community, as a way of screening *out* prospective members who do not share community values and who are less likely to contribute to community club goods. This prediction is developed in more detail by Berman (2000). This prediction contrasts with my finding that Amish communities impose punishments on children who *leave* in order to screen children *into* the community.

Second, Iannaccone argues that benefits of cooperation within strict religious communities, such as mutual insurance, are relatively more valuable to the poor, and that the poor should therefore be more likely to join strict religious communities. This argument is developed by Abramitzky (2008), who shows people with less human capital are more likely to remain members of Israeli kibbutzim. This prediction contrasts with my finding that children of richer Amish parents are more likely to remain in the Amish church.

Third, the club goods model does not explain the persistence of Amish church membership across generations. Since the club goods model has no intergenerational component, taken literally it implies that children of Amish parents should be no more likely to join the Amish church or other similarly strict religious communities than children from secular families with similar levels of wealth and human capital. However, it is clear that Amish children are more likely to join the Amish church than secular children are to join similarly strict religious communities, even after controlling for wealth and human capital endowments.

Fourth, the club goods model has difficulty explaining how religious institutions could persistently reduce their members' welfare. The club goods model implies that the community services provided by strict religious communities make members better off than non-members. This prediction is hard to reconcile with the persistently lower welfare of the Amish, as evidenced by high exit rates from Amish communities that do not impose punishments on children who leave.

Fifth, my model provides a more compelling explanation for high Amish fertility than the extension of the club goods model developed by Berman (2000). Berman argues that ultra-Orthodox Jews have high fertility rates because community rules reduce wages and hence the opportunity cost of time, especially for ultra-Orthodox women. While this is likely part of the explanation for high ultra-Orthodox fertility rates, it seeems incomplete because the ultra-Orthodox fertility rate appears to be much higher than the fertility rate even among poor secular Israelis who face comparable wages. I argue that rules among groups like the Amish and the ultra-Orthodox not only reduce the productivity of wage work, but are also aimed explicitly at increasing the amount of time devoted to family activities such as raising children. This dual effect may increase the fertility rate above the fertility rate in secular populations with comparably low wages.

Other models of religion and culture emphasize the role of religion and culture in determining people's preferences or beliefs, rather than the material rules and regulations imposed by religious institutions. Bisin and Verdier (2000, 2001) argue that parents have a preference that their children join the same religion, and that parents expend resources to instil this preference in their children. Kuran and Sandholm (2008) argue that people tend to adapt their preferences to fit both their own actions and the actions taken by the other members of their communities. Giuliano and Nunn (2019) argue that children acquire their parents' culture by imitating their parents' behavioral rules of thumb. All of these models are designed to explain the persistence of cultural traits such as religious affiliation across generations, and so they can help to explain the persistence of Amish church membership. With suitable modification, these models may be able to explain why children of richer Amish parents are more likely to remain Amish than children of poorer Amish parents. However, none of these models can explain the particular form of Amish rules and institutions, or the purpose of these rules. The only possible explanation for seemingly irrational rules in these models is that the rules were a rational response to environmental conditions at some point in the past, and that the rules were then transmitted culturally across generations even as the environment changed. In contrast, my model explains how seemingly irrational Amish rules in fact serve a rational purpose in the present, thereby explaining why Amish rules take the particular form that they do.

Carvalho and Koyama (2016), and Carvalho et al. (2017), also construct models in which religion is both a social and political institution that imposes rules on its members, and a cultural practice that is transmitted across generations. These models do not emphasize the role of community punishments in maintaining religious communities, although it would be easy to extend the models to include community punishments. Suitably modified, these models may be able to generate empirical predictions that are similar to mine. An advantage of my model relative to these other models is theoretical parsimony. My model explains both the political structure of Amish communities and the transmission of Amish religious affiliation across generations as a result of parents' desire to control their children's behavior. My model does not require the additional assumption that parents have an intrinsic preference that their children remain Amish for "identity" reasons (Akerlof and Kranton, 2000). In contrast, in Carvalho and Koyama (2016) religious rules are chosen by a "religious authority" whose preferences are distinct from the preferences of any other member of the community. Carvalho and Koyama do not explain how an agent with these preferences could be chosen to lead the community, and so

they do not explain why the religious community has the political structure that it does. In Carvalho et al. (2017), the results are driven by the assumption that people have intrinsic "identity" preferences. Again, it is not explained why people have these preferences, in contrast to my model in which parents' reduced form preference that their children remain in the Amish church is given a deeper microfoundation.

The Amish are a small and esoteric community, and it may seem that my model does not have much applicability beyond this unusual context. I argue, however, that many other religious and cultural communities have features similar to the Amish, if perhaps not to the same degree. Many religious and cultural communities impose rules that are meant to regulate family life, such as rules about proper gender roles, courtship and marriage, and child raising practices. Many religious and cultural communities also impose penalties on children who leave. For example, the traditional punishment in Islam for apostasy, conversion from Islam to another religion, is death (see Friedmann, 2003 for discussion of punishments for apostasy in Islam). Similarly, Kymlicka (1995) argues that communal property ownership in some minority cultural communities makes it difficult for children to assimilate into the majority by preventing children from taking property with them when they leave. I believe that my model is useful for understanding these other religious and cultural communities and hence has broader significance. In particular, my model may help to explain both how religious and cultural minority communities can persist, and how these communities can be persistently disadvantaged relative to the majority.

More positively, my theory suggests that religious institutions may help to strengthen family ties, which may increase investment in children's human capital. The declining strength of family ties and the resulting decrease in investment in children's human capital is a serious problem (Heckman, 2011; Lundberg et al., 2016), and promoting religion may help to prevent or reverse this trend. This possible positive effect of religion seems to apply more strongly to religious communities like non-ultra-Orthodox Jews that promote investment in child quality, unlike the Amish who seem to promote investment primarily in child quantity.

The remainder of the paper is organized as follows. Section 2 presents background on the Amish. Section 3 presents my theoretical framework. Section 4 presents my empirical results. Section 5 discusses alternative explanations. Section 6 concludes.

## 2. Background on the Amish

The Amish are one of several Protestant sects, collectively known as Anabaptists, which emerged in Switzerland and Germany in the aftermath of the Protestant Reformation.[2] The Anabaptists were among the most radical Protestant groups, in the sense that they went the furthest in rejecting the traditional authority structures of medieval life. In particular the Anabaptists rejected the traditional church hierarchy, instead holding meetings in private homes with lay preachers, and they also denied the right of the ruler of a territory to determine the religion of the territory's inhabitants. The Anabaptists believed instead that each person should make an informed choice of what religion to belong to on reaching adulthood. One consequence of this belief was that the Anabaptists did not practice infant baptism, hence their name. The Amish denomination was founded in 1693 in Switzerland by Jakob Amman, who split from the existing Anabaptist congregations over the issue of shunning. I discuss shunning in more detail below.

The Amish faced persecution in Europe due to their religious beliefs, and as a result some Amish people emigrated to the United States in the 18th and 19th centuries. The remaining Amish church in Europe did not prosper, and the last European Amish congregation disbanded in 1937. In the Americas, however, the Amish church grew quickly. Table 1 shows Amish population estimates in North America from 1900 to the present. This growth is due almost entirely to the high Amish fertility rate and not to conversion into the Amish church from outside, which is nearly impossible. In fact the Amish population has long been a subject of genetic research due to its extreme genetic isolation.

Amish ethical teachings emphasize humility, simplicity, sublimation of personal desire, and surrender to God's will. One aspect of this philosophy is separation from the outside world, which the Amish regard as materialistic and corrupt. As part of their separation from the outside world, the Amish speak a dialect of German called Pennsylvania German, which is related to the dialect of southern Germany where the Amish faith originated. Amish children are also taught to read standard German, and Amish church services are held in standard German using Martin Luther's 1534 German translation of the Bible. Most Amish people are also fluent in English and speak English with outsiders.

### 2.1. Amish community organization

Contemporary Amish society is organized into settlements, affiliations, and districts. A settlement is a geographical area with many Amish people. The largest settlement, centered in Holmes county, Ohio, and including parts of six surrounding counties, has more than 30,000 members. There are two other settlements of similar size, in Lancaster county, Pennsylvania, and Elkhart and Lagrange counties, Indiana, and many smaller settlements. When these settlements were originally formed they were located in rural areas, and traditionally all Amish men were farmers. However, in recent years urban development has encroached on Amish settlement areas and now many parts of the larger Amish settlements are more suburban or exurban than rural. Urban expansion and rising land prices have led many Amish men to abandon farming for other trades. In the Holmes county settlement, which is the main object of study in this paper, many Amish men work as artisanal craftsmen in areas such as woodworking and masonry, or as small business owners in businesses such as retail or construction. Other Amish men work as unskilled or semi-skilled laborers, including

[2] This overview of Amish history, culture, and institutions is drawn from Hostetler (1993); Hurst and McConnell (2010); Kraybill (2001), and Kraybill et al. (2013).

**Table 1**
North American Amish population estimates.

| Year | Estimated population |
|---|---|
| 1901 | 6300 |
| 1911 | 6600 |
| 1921 | 9960 |
| 1931 | 14,300 |
| 1941 | 21,100 |
| 1951 | 27,675 |
| 1961 | 36,855 |
| 1971 | 50,280 |
| 1981 | 77,955 |
| 1991 | 123,025 |
| 2001 | 189,335 |
| 2012 | 273,710 |

Source: Kraybill et al. (2013), p. 156.

for non-Amish employers. Finally, some Amish men remain in the traditional occupation of farming. Married women rarely work outside the home. Settlements do not have well-defined borders, and within each settlement Amish people and non-Amish people live side by side.

Except in the Lancaster county, Pennsylvania settlement, there are no settlement-wide governance institutions. Governance instead happens at the level of the district. A district is a collection of 20–40 households that meets once every two weeks in a member's house or barn to hold religious services. The size of a district is limited by the number of people who can fit in a house or barn for the service, and when a district becomes too large, it splits. Each district is responsible for making and enforcing its own rules, which I describe in more detail below. Rules are decided in special meetings held twice per year in each district. In theory, all adult members of the district have equal say in these meetings, but in practice most power in each district is exercised by the district leader, known as the bishop, through his ability to set the agenda. Bishops are chosen for life from among the male elders of a district through a complex procedure that is partly an election but that also involves elements of chance. All adult members of the district have a vote in these elections. However, because bishops serve for life, and because only established members of the community are eligible to be nominated as bishops, it seems reasonable to conclude that district rules primarily reflect the preferences of older members of the community.

In addition to enforcing the rules, districts provide many other goods and services for their members. For example, there is a substantial amount of risk sharing and charity within the Amish community, and many of these services are provided at the level of the district. Each district has an officer called the deacon who is responsible for maintaining an alms fund for members in difficulty, and richer district members may be required to provide no-interest loans to poorer members to buy houses or start businesses. The district also provides care for sick and elderly members, and organizes social events. Although it is hard to quantify the amount of redistribution generated by these community services, the overall level of redistribution seems to be higher than in non-Amish society.

While each district is responsible for deciding its own rules, in most cases groups of districts choose similar rules. A group of districts with similar rules is called an affiliation, and two districts in the same affiliation are said to be "in fellowship". Districts in fellowship may exchange guest preachers for Sunday services. In the Holmes county, Ohio settlement, there are four major affiliations, the Old Order, the New Order, the Andy Weaver affiliation, and the Swartzentrubers. The Swartzentrubers are not included in my dataset and hence not considered in the empirical section of this paper. The Old Order are the largest affiliation and represent the main line of the Amish tradition, going back to the Amish settlement of Ohio in the 19th century. The Swartzentrubers split from the Old Order in 1919, the Andy Weaver affiliation split from the Old Order in 1954, and the New Order split from the Old Order in 1966.

Movement across affiliations is sometimes possible, but it is rare. Most families encourage their children to marry within the same affiliation. In addition, in order to join a new district an Amish person must present a letter of recommendation from the bishop of his or her previous district. This discourages Amish people from switching affiliations.

As discussed in more detail below, some Amish people either decide that they do not want to be church members or are involuntarily expelled from the church. Amish people who leave the church have several options. Some join the Mennonites, another Anabaptist group which is theologically very similar to the Amish but which does not impose nearly as many rules on its members. There is a large Mennonite community in Holmes county and so joining this alternative community is not difficult. Other ex-Amish people join other Christian denominations or become atheists, and become "normal" members of the larger American culture. The many non-Amish people with Amish family names living in the Holmes county area provide evidence of this kind of assimilation in the past. It is worth emphasizing that because Amish and non-Amish people live side-by-side in Holmes county, "leaving" the Amish church does not necessarily imply moving geographically. Instead, a person who leaves the church can simply find a new social circle among his or her non-Amish neighbors.

### *2.2. Amish rules*

Amish districts impose many rules on their members. Roughly speaking, these rules fall into four categories. First, there are rules about the use of technology. Second, there are rules about family life and child raising. Third, there are rules about dress and

appearance. Fourth, there are rules about the punishment of district members who break the other rules of the district.

Amish rules about the use of technology are the aspect of Amish culture that is perhaps best known to outsiders. The rules forbid the use of a wide variety of modern technologies. Most notably, all districts prohibit driving cars. Instead Amish people travel in horse-drawn buggies, and these buggies are a common sight on roads in and near Amish settlements. Many districts also have rules about the use of other kinds of transportation and motorized equipment. For example, many districts have rules about the use of tractors. Some districts permit all tractors, some districts permit only tractors with steel wheels, which cannot be driven on the road, and some districts forbid tractors entirely. Some districts also prohibit bicycles. Most Amish districts prohibit the use of electricity from the grid, although some districts permit electric appliances powered by gas generators. Finally, most districts have rules restricting modern information and communication technology such as telephones and computers, although the extent to which these technologies are restricted varies from district to district.

Amish districts have many rules and customs regarding child raising and family life, including rules about proper gender roles in marriage, courtship customs, and child discipline. From the perspective of an economist the most important of these rules is the rule prohibiting children from attending school past the age of fourteen. This rule is followed in all districts, and as a result all Amish people have the same amount of education.[3] The Amish were specifically exempted from state mandatory education laws in the 1972 Supreme Court case *Wisconsin v. Yoder*. Before the age of fourteen, Amish children can either attend public schools along with non-Amish children, or they can attend private Amish schools. Private Amish schools are taught in one-room schoolhouses, typically by unmarried Amish women in their teens or early twenties who themselves have not attended school past the age of fourteen. It seems likely that private Amish schools are less effective than public schools at transmitting skills that are useful in the labor market, although they may be more effective at transmitting Amish cultural values.

Rules about dress and appearance serve to distinguish the Amish visually from non-Amish people. Men and women both dress modestly, and men have beards but not moustaches and distinctive haircuts. There are subtle differences between the clothing required by different affiliations, such as different rules across affiliations concerning the width of brims for men's hats. For someone familiar with the culture it is possible to distinguish members of different affiliations in this way.

Districts enforce the rules using a penalty called shunning. A shunned member is not allowed to interact in certain ways with church members in good standing. For example, it is forbidden to share a meal with or to accept gifts from someone who has been shunned. Shunning may be temporary or permanent. Temporary shunning lasts for two to six weeks and is a response to less serious violations of the rules. A church member who repeatedly breaks the rules, or who refuses to show remorse for past rule-breaking, may be permanently shunned. Permanent shunning lasts indefinitely, but even a member who has been permanently shunned may be accepted back into the community after making a confession and repenting in front of the congregation. Information about shunning is shared across districts in the same affiliation, and if it is discovered that a person has been shunned in his or her home district, members of other districts in the same affiliation will respect the prohibition against interacting with that person. Church members who fail to respect the prohibition against interacting with someone who has been shunned may be shunned themselves.

As I discuss below, in some cases church members may be shunned for voluntarily leaving the church, even if they have not otherwise done anything wrong. Church members who leave but who are not shunned may continue to interact with their friends and family who remain Amish, and in fact such continued interaction is common. Continued interaction is possible because leaving the church does not necessarily imply moving geographically. For example, Amish parents may rely on their children who have left the church to run errands that require access to a car.

Although the official sanctions imposed through the shunning process may seem relatively mild, the Amish describe shunning as a severe and traumatic punishment. Above all, this is due to the shame caused by shunning and the resultant loss of love and respect from friends and family members who remain members in good standing of the Amish church. One Amish person describes the impact of shunning, saying "You suddenly lose all your security, and you become a goat, like a piece of dirt." (Kraybill, 2001,p. 138) Hurst and McConnell (2010) survey excommunicated members of the Amish church and find that two-thirds believe that their families look down on them rather than seeing them as equals. Thus, even though shunning is rarely imposed, the threat of shunning affects every aspect of Amish behavior and social organization.

### *2.3. Variation in rules and child raising practices across affiliations*

The Amish describe the variation in rules and ideology across affiliations as being on a scale from "low" to "high". Lower affiliations are more removed from the outside world, while higher affiliations interact more with non-Amish society. The order of Holmes county affiliations from low to high is the Swartzentrubers, the Andy Weaver affiliation, the Old Order, and the New Order.

Districts in low affiliations typically are more restrictive about technology use than high affiliations. For example, districts in lower affiliations may forbid bicycles, tractors, and gas-powered appliances, while districts in higher affiliations may permit the use of these technologies. There is also variation in child raising practices across affiliations. Parents in lower affiliations are more likely to send their children to private Amish schools instead of secular public schools. Interestingly, parents in higher affiliations appear to devote more effort to teaching and enforcing Amish values in areas other than choice of school type. In particular, parents in higher affiliations devote more time to monitoring their children's behavior than parents in lower affiliations. Regarding parental supervision of children in the New Order, the highest affiliation, Hurst and McConnell (2010, p. 71) write, "The signature accomplishment

[3] Kraybill (2001, p. 81) reports the results of a survey of educational attainment among 812 Amish adults in Lancaster county, Pennsylvania. He finds zero who have more than an eighth grade education.

of the New Order churches, for instance, has been the extent to which they have been able to gain control over the lives of their young people." For example, New Order parents have abolished unsupervised "singings", gatherings of teenagers that start with group hymn singing followed by dancing, drinking, and listening to music. While unchaperoned singings remain popular in the lower affiliations, the New Order require that all singings be attended by parents. Similarly, the New Order has abolished the practice of "bed courtship." Bed courtship is a kind of dating allowed by lower affiliations, in which romantically attached teenagers spend the night together without parental supervision. Theoretically bed courtship does not involve sexual activity, although it is hard to know whether this rule is followed in practice. As a result of this more intensive monitoring, "the New Order did in fact develop a reputation for 'clean living'" (Hurst and McConnell, 2010, p. 71). Moreover, "there is widespread agreement across the settlement that New Order youth are also less likely to drink, smoke, and do drugs" (Hurst and McConnell, 2010, p.74).

Although there are differences in technology use and child raising practices across affiliations, these differences are not the most important differences in rules between affiliations. Instead, the central point of disagreement between affiliations concerns the proper use of shunning. Districts that have similar attitudes towards shunning are likely to be in fellowship even if they have divergent attitudes towards technology or family life. For example, the Lancaster county, Pennsylvania Amish districts are in fellowship with the Andy Weaver affiliation in Ohio, because they share the same policies regarding shunning, even though the Lancaster county Amish permit the use of more modern technologies than the Andy Weaver affiliation. A belief that shunning should be enforced more strictly caused the original split between the Amish church and other Anabaptist groups in 1693, and disagreements about shunning were also at the heart of the conflicts that caused the Swartzentrubers, the Andy Weaver affiliation, and the New Order to split from the Old Order in this century.

Affiliations disagree in particular about the extent to which shunning should be applied to members who have voluntarily left the church but who have not otherwise done anything wrong. The Swartzentrubers and the Andy Weaver affiliation practice strict shunning or *streng meidung*. This means that any member who leaves the church is permanently shunned. Shunning can be ended only if the wayward member returns to a district in the relevant affiliation. In the Old Order, shunning practices vary from district to district and from case to case. Members can often leave to join other Anabaptist denominations such as the Mennonites without incurring a penalty. Members who leave to join other Christian denominations, or who become atheists, are more likely to be shunned by their former Old Order districts. New Order districts do not shun former members as long as those members join a "Bible-believing church", and many New Order Amish believe that the practice of shunning should be abandoned completely. In principle, it is possible for children to leave any affiliation without incurring a penalty if they leave before being baptized, where baptism typically happens in the late teens or early twenties. In practice, however, there is considerable social pressure on children to be baptized. Community members may also "draw back" from interaction with children who leave the church before being baptized, even if the children are not officially shunned. Like the official punishment of shunning, this unofficial punishment is more likely to be applied by lower affiliations.

## 3. Theory

In this section, I develop a simple model to explain Amish institutions. My model is an overlapping generations variation on the club goods model of religion developed by Iannaccone (1992). I present two versions of the model that show two different ways in parents and children may have conflicting preferences over children's time allocation. In both models, parents can increase their welfare by imposing religious rules on their children, thereby encouraging children to allocate their time according to the parents' preferences and benefiting parents at the expense of their children. In the first version of the model, there are two generations, and parents want their children to spend time taking take care of them in their old age, as in Bernheim et al. (1985). In the second version of the model, there are three generations, and grandparents are altruistic towards their children and also directly altruistic towards their grandchildren, as in Saez Marti and Weibull (2005) and Galperti and Strulovici (2017). Because grandparents get utility from their grandchildren both directly and through the grandchildren's parents, grandparents want parents to allocate more time to taking care of grandchildren than parents want to allocate themselves.

After presenting the two versions of the model, I discuss Iannaccone's (1992) club goods model of religion, and explain the empirically testable differences between my model and the club goods model.

### 3.1. Model 1

Consider a religious community that consists of a number of families. In the first version of the model, there are only two generations in each family, so that each family consists of one parent and one (adult) child. Each child is endowed with one unit of time. The child can allocate his time towards either of two activities. The first activity yields only a private benefit to the child. This activity could be working for wages, or spending time on private leisure. I let $y$ denote the quantity of time spent on the private activity. The child can also devote his time towards an activity that benefits both the child and the child's parent. I let $x$ denote the quantity of time spent on the parent-benefiting activity. In this first version of the model, it is helpful to think of $x$ as time used by the child to take care of the parent in old age. The child's budget constraint is

$$x + y \leq 1$$

All parents are members of the religious community, and children can decide whether to join the religious community or not. Let $m \in \{0, 1\}$ be an indicator of whether a child is a member of the community, with $m = 1$ indicating that a child is a member. A child's community membership directly affects the child's utility in three ways. First, there is an intrinsic benefit to remaining in the Amish

church, which I label $a$. This benefit could represent an attachment to Amish identity (Akerlof and Kranton, 2000), or a preference for an Amish lifestyle developed through habit formation. I assume that $a$ varies in the population of children according to some distribution. For some children $a$ may be negative; this indicates that the child has an intrinsic preference for leaving the Amish church independently of church rules. Second, the cultural community may inflict punishments on children who leave. In the Amish case, this is the punishment of shunning. Let $p_c$ denote the punishment inflicted by the community on children who leave, and suppose that $p_c$ cannot be greater than some maximum punishment $P$. The reason that the community can punish children who leave is that children would like to maintain ties with their friends and family members who remain in the community. Thus, the community can punish children who leave by cutting them off from contact with their friends and family. Finally, parents can choose to inflict further individual punishments on their children in addition to any punishments inflicted by the community. For example, parents can threaten to disinherit children who leave the community.[4] Let $P(w_i)$ be the maximum additional individual punishment that can be inflicted by a parent $i$ with wealth $w_i$ on a child who leaves, and let $p_i \in [0, P(w_i)]$ be the punishment that parent $i$ actually chooses to inflict. The maximum individual punishment $P(w_i)$ is a function of the parent's wealth $w_i$. I assume that $\partial P(w_i)/\partial w_i > 0$, that is, that richer parents are able to inflict more severe punishments. For example, the threat of being disinherited by a rich parent is more severe than the threat of being disinherited by a poor parent. Parental punishments could also take less material forms. For example, if children want their parents to approve of their life choices, parents could punish children by withholding approval from children who leave the community. Children may value their parents' approval more if their parents are wealthier and hence viewed as more worthy of respect.

A child who chooses to remain a member of the community must follow religious rules. For example, as discussed above Amish communities require members to refrain from driving cars, using electricity from the grid, and attending school after the age of 14. Let $r \in [0, 1]$ indicate the strictness of religious rules, with larger values of $r$ indicating stricter rules. Strict religious rules reduce the utility productivity of children's private time. For example, Amish rules may reduce children's wages by restricting the occupations they can choose, or they may directly reduce the utility of private leisure time by restricting access to private leisure activities like watching television. Putting all of these components together, a child's utility is

$$U^C = U^C(x, (1 - rm)y, am - (p_c + p_i)(1 - m))$$

I assume that $U^C$ is increasing in all of its arguments. Each child chooses $m$, $x$, and $y$ to maximize his utility, subject to the child's budget constraint.

The main question of interest in the simple framework sketched so far is how children's behavior responds to changes in rules $r$ and punishments $p_c$. An increase in $r$ increases the effective price of private activity time $y$ for children who remain members of the community. This has both a substitution effect, causing the child to shift from the private activity to the parent-benefiting activity, and an income effect, causing children to decrease effective time allocation to both activities (assuming that both activities are normal goods). If the substitution effect is sufficiently strong it may outweigh the income effect, in which case an increase in $r$ can cause an increase in $x$. In other words, strict religious rules may reduce the amount of time that children spend working for wages and on private leisure, and may increase the amount of time that children spend with their families. I assume for the remainder of this section that an increase in $r$ does in fact increase $x$.

Let $\bar{x}_1(r)$ and $\bar{y}_1(r)$ be the optimal values of $x$ and $y$ for a child who is a member of the church given rules $r$, and let $\bar{x}_0$ and $\bar{y}_0$ be the optimal values of $x$ and $y$ for a child who is not a member of the church. A child chooses to leave the church if $U^C(\bar{x}_1(r), (1 - r)\bar{y}_1(r), a) < U^C(\bar{x}_0, \bar{y}_0, -(p_c + p_i(w)))$. An increase in $r$ reduces $U^C(\bar{x}_1(r), (1 - r)\bar{y}_1(r), a)$ without affecting $U^C(\bar{x}_0, \bar{y}_0, -(p_c + p_i(w)))$. Thus, an increase in $r$ increases the probability that a child chooses to leave the church. This tendency could be counteracted by a corresponding increase in the punishment $p_c$, which would reduce each child's propensity to leave. An increase in parents' wealth $w$ also reduces children's propensity to leave the church by increasing the severity of the individual parental punishment $p_i$ for children who leave.

Having outlined children's behavior, I now describe parents' preferences and actions. Parents are altruistic towards their children, but also get utility from the time that children spend caring for them in old age. Assuming that parents' utility is separable in their own consumption and their children's utility, I write a parent's utility function as

$$U^P = u^P(x) + \alpha U^C$$

I assume that $u^P$ is increasing in $x$. Here $\alpha$ measures the strength of a parent's altruism towards her child.

Parents have political control over the community, which allows them to choose the strictness of religious rules $r$ and the strength of community punishments $p_c$. Parents collectively choose $r$ and $p_c$ before children make their decisions. Each parent also chooses whether to inflict an additional individual punishment $p_i$. Parents can commit to the individual punishment $p_i$ for children who leave despite their altruism towards their children.[5]

Now, suppose that $r = 0$, and consider the effect of a marginal increase in the strictness of religious rules $r$. Assuming that $p_c$ is sufficiently large that an increase in $r$ does not induce the child to leave the community, the effect of an increase in $r$ on a parent's utility is

[4] It is interesting to note that while it is impermissible in Amish society to accept a gift from someone who has been shunned, it is permissible to give a gift to someone who as been shunned. So shunning by itself does not prevent parents from making bequests to their children.

[5] Parents may be able to commit to punishing their children because of community pressure to inflict punishments, or because they want to maintain their reputations with their other children (Hao et al., 2008).

$$\frac{\partial U^P}{\partial r} = \frac{\partial u^P}{\partial x}\frac{\partial x}{\partial r} + \alpha\frac{\partial U^C}{\partial r}$$

The first term in this expression is positive and the second term is negative. However, if the second term is sufficiently small, then an increase in $r$ may increase parents' utility. This explains why parents may want to choose $r > 0$, even though parents are altruistic towards their children and an increase in $r$ reduces children's welfare. Moreover, parents may want to choose $p_c > 0$ in order to ensure that $r > 0$ does not cause children to leave the community. This explains why parents may want to choose $p_c > 0$. If strict religious rules $r > 0$ reduce children's welfare by more than they increase parents' welfare, then strict religious rules may reduce overall lifetime utility of Amish community members.

The key assumption in model 1 is that parents get utility from old age care provided by their children as well as altruistically from their children's utility, as in Bernheim et al. (1985). Next I consider a model in which only get utility from children altruistically, and do not care directly about children's actions. Nevertheless, there is still intergenerational conflict about how younger generations allocate their time.

### 3.2. Model 2

In the second version of the model, again there is a religious community that consists of a number of families. However, in this version of the model, each family consists of three generations, one grandparent, one parent, and some number of children. All grandparents are members of the religious community, and it is grandparents who have political control over the community and choose religious community rules $r$ and punishments $p_c$. Each grandparent may also choose to inflict the additional individual punishment $p_i \leq P(w_i)$, where $w_i$ is grandparent $i$'s wealth. Parents allocate time between a private activity $y$ and an activity that benefits other family members $x$, subject to the budget constraint $x + y \leq 1$. However, instead of interpreting $x$ as time spent by parents to care for grandparents in old age, in this version of the model I interpret $x$ as time spent by parents on taking care of children. Parents also choose whether to remain in the church or to leave, and may suffer punishments if they leave. As in the first version of the model, parents get intrinsic benefit $a$ from remaining in the church, and parents who leave suffer the community punishment $p_c$ and the individual punishment $p_i$. Children do not take any actions.

The amount of time that a parent spends taking care of children may affect both the quality and the quantity of the parents' children. Let $n(x)$ be the number of children that a parent has, and let $U^C(x)$ be the quality (that is, the utility) of each child. Parents are altruistic towards their children and so parents get utility from their children's utility. Parents also get utility from time spent on the private activity $y$, but as in model 1, religious rules reduce the productivity of private time. I assume that a parent's utility is separable in her private consumption and her children's utility. Thus I write a parent's utility as:

$$U^P = u^P((1 - rm)y,\ am - (p_c + p_i)(1 - m)) + \alpha n(x)U^C(x)$$

As in model 1, $m$ is an indicator variable that equals 1 if the parent remains in the community and equals zero otherwise, and $\alpha$ measures the strength of parents' altruism towards children. Parents choose $x$, $y$, and $m$ to maximize their utility subject to the budget constraint, given the grandparents' choices of $r$, $p_c$, and $p_i$.

Grandparents are altruistic towards parents, in the same way that parents are altruistic towards children. Following Saez Marti and Weibull (2005) and Galperti and Strulovici (2017), I assume that grandparents are also directly altruistic towards their grandchildren. In contrast to model 1, grandparents do not get any direct utility from consumption other than the altruistic utility that they get from parents' and children's utility. Assuming that grandparents' utility is separable in parents' and children's utility, I write grandparents' utility as

$$U^G = \alpha U^P + \beta n(x)U^C(x) = \alpha u^P((1 - rm)y,\ am - (p_c + p_i)(1 - m)) + (\alpha^2 + \beta)n(x)U^C(x)$$

Here $\beta$ indicates the strength of grandparents' direct altruism towards their grandchildren.

Suppose that $p_c$ is sufficiently large that the parent remains in the religious community. Then the solution to the parent's problem is to choose $x$ and $y$ such that

$$(1 - r)u_1^p((1 - r)y,\ a) = \alpha\left[n(x)\frac{\partial U^C}{\partial x} + \frac{\partial n(x)}{\partial x}U^C(x)\right] \tag{1}$$

Here $u_1^p$ is the partial derivative of $u^P$ with respect to its first argument.

On the other hand, if the grandparent could control the parent's time allocation, then the grandparent would choose $x$ and $y$ to solve

$$(1 - r)u_1^p((1 - r)y,\ a) = \left(\alpha + \frac{\beta}{\alpha}\right)\left[n(x)\frac{\partial U^C}{\partial x} + \frac{\partial n(x)}{\partial x}U^C(x)\right] \tag{2}$$

Comparing (1) and (2) shows that the grandparent wants the parent to allocate more time to child care than the parent wants to allocate herself. From (1), an increase in the strictness of the rules $r$ may increase the amount of time that parents allocate to child care, which increases grandparents' utility by moving parents' time allocation closer to grandparents' preferences. If increasing $r$ does not reduce parents' utility too much, then grandparents may want to choose $r > 0$. If grandparents choose $r > 0$, then they may also want to choose a community punishment $p_c > 0$ to ensure that parents remain in the community and follow the rules.

To provide a concrete example in which an increase in $r$ increases grandparents' utility, suppose that

$u^p = (1-r)y + am - (p_c + p_i)(1-m)$, $n(x)U^C(x) = (1-\epsilon)x$, and $\alpha = 1$, where $\epsilon > 0$ is some small number. Then for $r = 0$, the parent chooses $y = 1$ and $x = 0$, so that the grandparent's utility is 2. On the other hand, if $r > \epsilon$, then the parent chooses $y = 0$ and $x = 1$, so that the grandparent's utility is $(1-\epsilon)(1+\beta)$. If $\epsilon$ is sufficiently small relative to $\beta$, then choosing $r > \epsilon$ increases the grandparent's utility.

As discussed in Section 2, Amish parents do not send their children to school past the age of 14. On the other hand, as will be discussed in more detail below, the Amish fertility rate is very high. This suggests that to a first approximation, Amish rules do not increase parents' investment in child quality, that is, $\partial U^C(x)/\partial x \approx 0$. On the other hand, it seems likely that Amish rules have a significant effect on parents' investment in child quantity, that is, $\partial n(x)/\partial x > > 0$. Given these assumptions, strict religious rules $r > 0$ do not increase the utility of any individual child, even though they increase the total number of children. If strict religious rules $r$ reduce parents' welfare by more than they increase grandparents' welfare, then strict religious rules may reduce total lifetime welfare for Amish community members, as in model 1.

An alternative assumption, that does not seem to apply to the Amish but that may apply to other strict religious communities, is that strict religious rules have relatively little effect on child quantity $\partial n(x)/\partial x \approx 0$ but a large effect on child quality $\partial U^C(x)/\partial x > > 0$. In this case, members of strict religious communities may have higher human capital, wealth, and lifetime welfare than non-members. This result may help to explain religious communities with persistently high levels of human capital relative to non-members, such as non-ultra-Orthodox Jews (Botticini and Eckstein, 2012) and German Protestants (Becker and Woessmann, 2009).

The key assumption in model 2 is that grandparents get utility from their grandchildren both directly and indirectly through parents. If $\beta = 0$, that is, if grandparents are not directly altruistic towards their grandchildren, then (1) and (2) are the same and grandparents and parents have the same preferences over parents' time allocation. In this case, grandparents do not want to impose strict religious rules $r > 0$ on parents. The idea that direct altruism towards grandchildren can generate conflict between grandparents and parents also appears in Saez Marti and Weibull (2005) and Galperti and Strulovici (2017).

### 3.3. Discussion and empirical implications

Both versions of my model imply that older generations may want to impose religious rules on younger generations. These rules cause younger generations to allocate their time in accordance with the preferences of older generations, thereby benefitting older generations at the expense of younger generations. In both versions of the model, older generations ensure that younger generations follow the rules by imposing both community punishments and individual punishments on members of the younger generation who leave. Richer members of the older generation are able to impose more severe individual punishments because they have more resources. There are two difference between the two models. First, the two models differ in the behavior that older generations want to encourage by imposing rules. In the first version of the model, parents want their children to spend more time caring for the parents in old age. In the second version of the model, grandparents want parents to spend more time caring for children. Second, the two models may have different welfare implications.

My model suggests the following empirical implications that can be tested by studying Amish institutions:

1. Some Amish communities may impose punishments on children who leave the church. Children in communities that impose more severe punishments are less likely to leave the church.
2. Children of wealthier Amish parents are less likely to leave the church than children of poorer Amish parents.
3. When members of one generation become parents, they impose the same punishments on the next generation as were imposed on the first generation when they were children. As a result, Amish church membership is transmitted across multiple generations, and family dynasties that belong to the Amish church in one generation are likely to remain in the church even after many generations.
4. If the costs to children of following religious rules are greater than the benefits to parents of having children who follow the rules, then the Amish may have lower lifetime welfare than their non-Amish neighbors, and lower Amish welfare may persist across many generations. Welfare differences between the Amish and the non-Amish imply that there may be high exit rates from Amish communities that do not impose punishments on children who leave.
5. In the second version of the model, Amish rules encourage parents to spend more time on child-raising. Since time spent by the Amish on child-raising seems to increase child quantity more than child quality, the Amish have higher fertility than other populations with similar levels of wealth and human capital. In other strict religious groups, time spent child-raising may increase child quality more than child quantity. This may help explain the phenomenon of strict religious groups whose members, unlike the Amish, have persistently higher human capital than non-members.

#### 3.3.1. Why does the severity of community punishments vary across Amish affiliations?

As discussed in Section 2, there is variation across affiliations in the severity of the punishment of shunning imposed by different affiliations. My model suggests that one factor that may cause variation in the severity of shunning is variation the strength of parental altruism $\alpha$. Variation in parental altruism may also affect child raising practices, which may in turn affect children's intrinsic attachment to the Amish church $a$. Thus, the community punishment $p_c$ and intrinsic attachment to the Amish church $a$ may be jointly determined. This possible correlation between $a$ and $p_c$ is important for the interpretation of my empirical results. In order to explore this possible correlation, I now present an extension of my model in which both variation in community punishments and variation in children's intrinsic attachment to the Amish church depend on variation in parental altruism.

To save space I develop my extension only in the context of the first version of my model, although it would be easy to develop a

similar extension of the second version of the model. Recall that in the first version of the model there are two generations in each family, a parent and a child. The child chooses between spending time on a private activity $y$ and old age care for the parent $x$, subject to the budget constraint $x + y \leq 1$. Parents collectively choose religious rules $r$ and the community punishment $p_c$. For simplicity I assume that there is no individual parental punishment $p_i$. In a given affiliation, all parents have the same level of altruism $\alpha$, although parental altruism may vary across affiliations.

In addition to choosing community rules and punishments, parents make decisions about how to raise their children. Child raising decisions can vary along two dimensions. First, parents can choose how much effort to expend teaching their children Amish values and monitoring their behavior. For example, as discussed in Section 2, parents in different Amish affiliations differ in the amount of time they spend supervising their children's social lives. Let $e$ be a parent's socialization effort and let $c(e)$ be the cost of effort, where $c(e)$ is increasing and convex. Second, parents can choose how much secular education to give their children. As discussed in Section 2, although all Amish children have the same number of years of education, Amish parents can choose whether to send their children to secular public schools or to private Amish schools. Since the relevant form of secular education is public education, I assume that increasing $s$ has no material cost.

As discussed in Section 2, secular public schools are probably better at increasing children's productivity that private Amish schools. I assume that secular education increases the utility productivity of both the child's private activity and old age care for the parent. For example, a better educated child might have higher wages, and might also be better equipped to help the parent navigate the health care system or deal with end-of-life legal issues. Thus, I assume that the child's utility function is

$$U^C = U^C(xs, (1 - rm)ys, am - p_c(1 - m))$$

Assuming that each parents' utility function is separable between her own consumption, her child's consumption, and the cost of child raising effort, I write the parent's utility function as

$$U^P = u^P(xs) - c(e) + \alpha U^C$$

Each child's attachment to the church $a$ is stochastic, and the distribution of $a$ may depend on child raising effort $e$ and secular education $s$ chosen by parents. Let $f_{se}(a)$ be the distribution of $a$. I assume that if $s' > s$ then $f_{se}$ first order stochastically dominates $f_{s'e}$, and that if $e' > e$ then $f_{se'}$ first order stochastically dominates $f_{se}$. That is, parental child raising effort increases intrinsic attachment to the Amish church and secular schooling reduces intrinsic attachment to the Amish church. A child chooses to leave the church if the child's utility of leaving is greater than the utility from staying, where the utility from staying includes any utility from the realized value of $a$. Parents do not know the realization of their child's value of $a$ at the time at which they choose the church's rules, community punishments, child raising effort, and level of secular education, but they do know the distribution of $a$. Let $\pi(s, e, r, p_c)$ be the probability that a child leaves the church, given actions $s$, $e$, $r$, and $p_c$ chosen by parents. It is easy to see that $\pi$ is increasing in $s$ and $r$ and decreasing in $e$ and $p_c$.

Let $EU^C(s, e, r, p_c)$ be a child's expected utility given parents' choices. It is also easy to see that $EU^C(s, e, r, p_c)$ is increasing in $s$ and $e$, and decreasing in $r$ and $p_c$ as long as $\pi > 0$. Let $Eu^P(s, e, r, p_c)$ be expected value of $u^P$, given the parent's and the child's choices. Using this notation, I can write each parent's problem as

$$\max_{s,e,r,p_c} Eu^P(s, e, r, p_c) - c(e) + \alpha EU^C(s, e, r, p_c)$$

subject to the constraint

$$p_c \leq P.$$

I assume that $P$ is sufficiently small that $\pi > 0$ for all feasible values of $s$, $e$, $r$, and $p_c$.

The objective function is supermodular in $\alpha$ and $\{s, e, -r, -p_c\}$. Therefore, by Topkis' theorem, I can state the following results:

1. An increase in $\alpha$ reduces the severity of punishment $p_c$.
2. An increase in $\alpha$ reduces the strictness of religious rules $r$.
3. An increase in $\alpha$ increases child raising effort $e$.
4. An increase in $\alpha$ increases secular education $s$.

Notice that parents may choose both $s > 0$ and $r > 0$, even though increasing $s$ and increasing $r$ have opposite effects on the productivity of the child's private time $y$, because increasing $s$ also increases the productivity of time $x$ spent caring for the parent in old age.

These results are consistent with the empirical facts described in Section 2, namely that parents in higher affiliations impose less severe forms of shunning and less strict rules, are more likely to educate their children in secular schools, and exert more child raising effort. The relationship between affiliation and the average level of children's attachment to the Amish church is ambiguous. On the one hand, parents in higher affiliations are more likely to send their children to secular schools, which may reduce $a$ on average. On the other hand, parents in higher affiliations are more likely to exert effort to monitor children's behavior and guide their social lives, which may increase $a$ on average.

Qualitative accounts of Amish society support the hypothesis that family altruism is a key factor explaining reluctance to impose severe forms of shunning in some affiliations. Altruism may cause people to resist shunning their family members, even at the risk of being shunned themselves if discovered by the community. Kraybill (2001, p. 139) describes this phenomenon as follows:

"The application of shunning varies widely from family to family. Many times it is relaxed in private homes but tightened in public settings if other church members are present, attesting to its ritual character and ceremonial role in the community. Many families treat family members under the ban [that is, family members who have been shunned] with love and care in the privacy of their homes. Despite its theological purposes, shunning is a painful process. One woman said, 'I'm not responsible for being born into a church that practices shunning. I have an uncle and aunt and cousins in the ban, and it may separate us on a social level, but it could never sever the cord of love.'"

Similarly, Hurst and McConnell (2010, p. 92) write:

"In some cases family and church members are less likely to take the hard line on shunning when no one else is watching, as this ex-Swartzentruber male relates: 'I have brothers that'll shun somebody right in front of the whole family, but when they're by themselves, they won't.'"

### 3.4. Comparison with the club goods model

It is worthwhile to compare the empirical predictions of my model to the predictions of Iannaccone's (1992) club goods model of religion. In order to make this comparison, I now review a stripped-down version of the club goods model.

In the club goods model, there are $n$ agents and a religious community. In contrast to my model, there are no family ties between agents and so there is no distinction between parents and children. Each agent has an endowment of one unit of time, which can be spent either on an activity that yields only private benefits such as wage work or private leisure, or an activity that benefits the entire religious community if the agent is a member of the community. For example, this community activity could be organizing community events or providing charitable aid to community members who have suffered negative shocks. Following the notation introduced earlier, let $y$ be the amount of time spent on the private activity and let $x$ be the amount of time spent on the community activity. A community member's budget constraint is

$$x + y \leq 1$$

Let $m^i$ denote whether agent $i$ is a member of the religious community, with $m^i = 1$ indicating that the agent is a member and $m^i = 0$ indicating that the agent is not a member. Let $M = \sum_{i=1}^{n} m^i$ denote the number of members. As in my model, the community can inflict punishments. However, unlike my model, in the club goods model it is not possible for the community to inflict punishments on non-members, since non-members have no ties to the community and so the community has no way to hurt non-members. Thus, the punishment inflicted on agent $i$ is $p^i m^i$, where $p^i$ measures the severity of the punishment and $m^i$ indicates whether agent $i$ is vulnerable to punishment through community membership. Iannaccone suggests examples of punishments such as painful initiation rites inflicted by religious communities on members. Berman (2000) argues that ultra-Orthodox Jews require community members to spend twenty years or more in unproductive Yeshiva study as a form of punishment.

Community members get utility from the private activity, from their own contributions to the community activity, and from the average community contribution of all community members. As in my model, the community can impose rules that reduce the utility productivity of time spent on the private activity. As before, let $r$ the measure the strictness of the rules. Community members may also lose utility from community punishments. Non-members do not benefit from the community contributions of others, but are also not subject to community punishments. Putting everything together, the utility of agent $i$ is

$$U^i = U^i\left(x^i,\ (1 - rm^i)y^i,\ m^i\frac{1}{M}\sum_{j=1}^{n} x^j m^j,\ -p^i m^i\right)$$

Once again, I assume that $U^i$ is increasing in all of its arguments.

As in my model, in the club goods model an increase in $r$ affects religious community members' choices through an income effect and a substitution effect. If the substitution effect is sufficiently strong, an increase in $r$ may cause an increase in $x^i$ for some or all religious community members $i$. In turn, if community members care enough about the average community contribution $(1/M)\sum_{j=1}^{n} x^j m^j$, then an increase in $r$ may be welfare improving for some or all members of the religious community.

Suppose that agent $i$ is deciding whether to join the religious community, and suppose that some subset of other agents $\mathcal{M}$ have already decided to join the community. Let $\hat{x}^{1i}(r)$ and $\hat{y}^{1i}(r)$ be the Nash equilibrium values of $x^i$ and $y^i$ that agent $i$ chooses if he joins the community, conditional on the strictness of the community's rules $r$. Let $\hat{x}^{0i}$ and $\hat{y}^{0i}$ be the values of $x$ and $y$ that the agent chooses if he does not join the community. The agent joins the community if $U^i(\hat{x}^{1i}(r),\ (1 - r)\hat{y}^{1i}(r),\ \frac{1}{1+|\mathcal{M}|}(\hat{y}^{1i}(r) + \sum_{j\in\mathcal{M}} y^j),\ p^i) > U^i(\hat{x}^{0i},\ \hat{y}^{0i},\ 0,\ 0)$. Because community members benefit from the average community contribution of other community members, community members may benefit if prospective members who are likely to contribute less than average to the community can be prevented from joining. Thus, community members may want to choose positive $p^i$ for prospective members $i$ who are likely to contribute less than average. If community members cannot determine which prospective members are likely to contribute less than average, the community may still benefit by imposing a uniform punishment $p$ on all members, if prospective members who are likely to contribute less than average also get lower utility on average from community membership. In this case, a uniform punishment can screen out prospective members who are likely to contribute less than average to the community.

The club goods model yields empirical predictions that contrast with the empirical predictions of my model. I present these predictions again in a numbered list, with each numbered item corresponding to the same numbered prediction from my model.

1. Religious communities do not impose punishments on people who leave the community. However, religious communities may impose punishments on some or all people who join the community, in order to screen out potential members who are likely to contribute less than the community average to the community good. The fraction of potential members who join is lower in communities that inflict more severe punishments.
2. Iannaccone argues that in practice, the club goods provided by religious communities usually benefit the poor more than the rich. For example, one of the main goods provided by many religious communities is mutual insurance, which is less valuable to the rich because it is easier for the rich to self-insure. Thus, the poor are more likely to join strict religious communities than the rich.
3. There is no reason in the club goods model to think that membership in strict religious communities should be especially likely to persist across generations. Children of strict religious community members should be no more likely to join strict religious communities themselves than other children with similar levels of wealth and human capital.
4. Strict religious community members have equal or higher welfare than non-members, after taking into account possible differences in endowments of wealth and human capital between members and non-members.
5. Members of strict religious community members may have higher fertility than non-members, because religious rules reduce wages and hence the opportunity cost of time spent on child-raising, especially for women. However, the club goods model does not predict that members of strict religious communities have higher fertility than non-members with similarly low wages.

In the next section, I argue that the facts about Amish institutions conform better to the predictions of my model than they do to the predictions of the club goods model, at least with respect to the five predictions advanced in this section. Thus, while the club goods model may be part of the explanation of Amish institutions, it is incomplete. I argue that my model must be part of the explanation of Amish institutions as well.

## 4. Evidence

### *4.1. Preliminary evidence*

Before presenting my main evidence, I discuss preliminary evidence from sociological accounts suggesting that my model is an important part of the explanation of Amish institutions. Consistent with my model, the Amish themselves believe that one of the most important functions of their rules is to encourage Amish community members to spend more time with their families. For example, Kraybill and Nolt (2004) argue that a primary purpose of Amish church rules restricting the size of Amish-owned businesses is to encourage business owners to spend more time with their families instead of becoming caught up in work. They write (pp. 125–126):

"The Amish fear large businesses for a number of reasons.... A large business easily encroaches on family life, as entrepreneurs become engulfed in manufacturing, sales, or bookkeeping. Although the community values hard work, it frowns on jobs that completely dominate everything else. 'A lot of Amish are workaholics,' said an Amishman who manufactures storage barns. 'They pretty easily get too caught up in their work. The church doesn't want a business to get too awfully big, or pretty soon you're living for your job. You get too caught up in it. That's not right.'"

Similarly, one of the major justifications for church rules favoring farming and other home-based business is that these businesses allow families to spend more time together. Kraybill and Nolt (2004, p. 94) quote one Amish man saying, "We're family centered, and the farm is the best place to raise a family." They quote another Amish man saying (p.94), "There is always a need for community-oriented shops, but once again, the setting should be such that the father can be at home with the family."

These quotes provide preliminary evidence that one function of Amish rules and institutions is to increase the amount of time that Amish church members spend with their families. Next I discuss more quantitative evidence.

### *4.2. Data*

I study the Amish using data from two sources. The first source is the Amish directory for the Holmes county settlement. The Amish directory is a list of all the households in the settlement, with the exception of households in the Swartzentruber affiliation who do not participate. It is published irregularly every 5–10 years. For each household, it contains the names and birth dates of all members of the household, the occupation of the household head, and the names of all of the children of the household head and the spouse of the household head. For each child the directory states whether the child is a member of the Amish church and whether the child lives in the settlement. The main outcome variable that I am interested in is whether a child leaves the church.

I make use of two editions of the Amish directory, the 1988 edition and the most recent 2010 edition. My sample is all children of Amish parents between the ages of 8 and 16 who are living in the settlement in 1988. I match this sample with the 2010 directory in order to find outcomes for children in 2010, when the children are between the ages of 30 and 38. Since decisions to leave the church are mostly made early in life, it is likely that most of these children have permanently chosen their religious status by age 30.[6] I

[6] I chose my sample in order to ensure that children would be young enough in 1988 that they had not yet married, but old enough in 2010 that they would have had the opportunity to become baptized and officially join the church. Baptism typically occurs in the late teens or early twenties, as discussed in Section 2. The youngest children alive in 1988 are only 21 or 22 in 2010. Many of these children likely have not yet have had the opportunity to become baptized by 2010 and so have not made an affirmative choice to remain in the church or to leave. In contrast, it is safe to assume that 30 year olds who remain in the church have made an affirmative choice to remain. This concern motivates the lower bound on the age

**Table 2**
Summary statistics.

| | All matched observations | Observations with building value > 0 |
|---|---|---|
| **Outcome variable** | | |
| Child leaves church | 0.147 | 0.148 |
| | (0.354) | (0.355) |
| **Child-level variables** | | |
| Child's age | 11.79 | 11.90 |
| | (2.579) | (2.580) |
| Child is male | 0.511 | 0.509 |
| | (0.500) | (0.500) |
| **Household-level variables** | | |
| Father is farmer | 0.513 | 0.493 |
| | (0.500) | (0.500) |
| Father's age | 41.13 | 41.64 |
| | (6.403) | (6.430) |
| Mother's age | 40.08 | 40.59 |
| | (6.344) | (6.339) |
| Building value | 59,893 | 72,726 |
| | (41,107) | (33,443) |
| Completed sibship size | 7.836 | 7.834 |
| | (3.112) | (3.152) |
| Observations | 3627 | 2987 |

Standard deviations in parentheses. Dollar values are 1988 dollars. Child's age is in 1988.

**Table 3**
Summary statistics, by affiliation.

| | All matched observations | | | Observations with building value > 0 | | |
|---|---|---|---|---|---|---|
| | New Order | Old Order | Andy Weaver | New Order | Old Order | Andy Weaver |
| **Outcome variable** | | | | | | |
| Child leaves church | 0.360 | 0.144 | 0.047 | 0.370 | 0.146 | 0.0319 |
| | (0.481) | (0.351) | (0.212) | (0.484) | (0.353) | (0.176) |
| **Child-level variables** | | | | | | |
| Child's age | 11.9 | 11.8 | 11.6 | 11.9 | 11.9 | 11.7 |
| | (2.61) | (2.58) | (2.55) | (2.60) | (2.58) | (2.57) |
| Child is male | 0.523 | 0.509 | 0.514 | 0.507 | 0.510 | 0.509 |
| | (0.500) | (0.500) | (0.500) | (0.501) | (0.500) | (0.500) |
| **Household-level variables** | | | | | | |
| Father is farmer | 0.500 | 0.488 | 0.620 | 0.447 | 0.482 | 0.567 |
| | (0.501) | (0.488) | (0.487) | (0.498) | (0.500) | (0.496) |
| Father's age | 43.2 | 41.1 | 39.9 | 43.8 | 41.6 | 40.5 |
| | (6.61) | (6.33) | (6.30) | (6.65) | (6.33) | (6.42) |
| Building value | 63,120 | 60,370 | 56,301 | 76,010 | 72,076 | 73,718 |
| | (40,432) | (41,052) | (41,507) | (31,409) | (34,176) | (31,148) |
| Completed sibship size | 6.97 | 7.69 | 8.88 | 6.88 | 7.70 | 8.96 |
| | (2.91) | (3.07) | (3.12) | (2.93) | (3.07) | (3.33) |
| Observations | 342 | 2629 | 656 | 284 | 2202 | 501 |

Standard deviations in parentheses. Dollar values are 1988 dollars. Child's age is in 1988.

observe children's outcomes through entries for their parents in the 2010 Amish directory, so I am able to observe outcomes even for children who have left the church or who have migrated, as long as their parents remain members of the Holmes county Amish church. Children who migrate may remain in the church, since children may migrate to other Amish settlements; conversely, children may leave the church without migrating, since there are also many non-Amish people in Holmes county. I do not observe outcomes for children of parents who die between 1988 and 2010, or for children of parents who leave the church or the settlement between 1988 and 2010. I am able to observe outcomes for 93% of the children in my sample.

I match this data with 1988 property tax data collected from the county treasurer's offices of the counties that contain the Holmes county settlement, which I use as a measure of parents' wealth. Not surprisingly the majority of the children in my sample live in Holmes county, but the settlement also includes parts of neighboring Tuscarawas, Coshocton, Wayne, and Stark counties, Ohio. I observe the assessed value of buildings and land owned by each parent in 1988. Table 2 shows summary statistics, for the full sample of children who I am able to match to outcome in 2010, and for the subset of children who live in households that own buildings with value greater than zero in 1988. Table 3 shows summary statistics by affiliation. Table 3 shows that members of the Andy Weaver

(*footnote continued*)
in my sample. However, the lower age bound is somewhat arbitrary. I chose the age bounds when collecting data and before having computed any results; the results presented in the paper use all of the data that I collected.

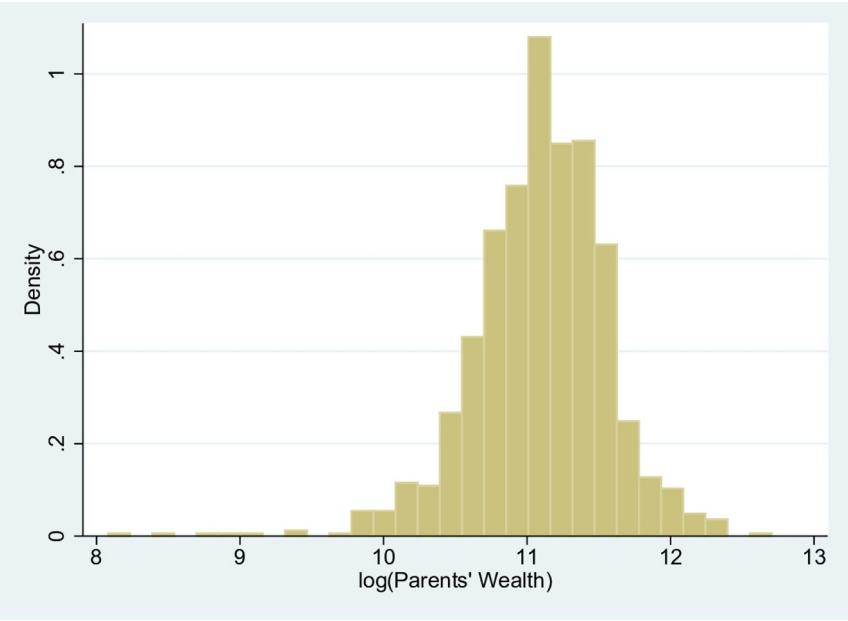


**Fig. 1.** Distribution of log(Parent's Wealth). This figure shows the distribution of the logarithm of the value of buildings held by all parents in the sample whose children I can match to outcomes and for whom the value of buildings owned is greater than zero.

affiliation appear to be somewhat more likely to be farmers and to have somewhat larger families that members of other affiliations.

My preferred measure of parents' wealth is the logarithm of the value of buildings owned by the parents. I think of this as a proxy for the flow value of housing consumed by each household, which in turn is a proxy for each household's permanent income. I prefer the value of buildings, and not the total value of buildings and land, for two reasons. First, for parents who are farmers land is a productive asset rather than a consumption good, and so the value of buildings owned by parents is a better proxy for consumption and permanent income. Second, and more importantly, while the method of assessing building values for tax purposes in Ohio is relatively straightforward, the method of assessing land values is complicated, non-transparent, and highly political. In particular, there are two regulatory regimes for assessing land values. Some parcels of land are assessed at market value, while other parcels are assessed according to their Current Agricultural Use Value (CAUV). CAUV values are lower than market values by up to a factor of ten, and there is no easy way to convert CAUV values to market values. For most households, I observe both the market value and the CAUV value of land. However, for households in Wayne and Tuscarawas counties, I observe only the CAUV value for parcels that fall under the CAUV regime. Thus, values of CAUV parcels in Wayne and Tuscarawas counties are not comparable to parcels in other counties or to parcels in the same county that do not fall under the CAUV regime. For these reasons I use the value of buildings as my primary measure of parents' wealth. In some regressions I use the total value of buildings and land owned by parents as an alternate measure; in these regressions I drop observations from Wayne and Tuscarawas counties. Fig. 1 shows the distribution of the logarithm of the total value of buildings owned for all parents in the sample whose children I can match to outcomes and for whom the value of property owned is greater than zero. Overall the distribution appears to be roughly log-normal. There are some outliers, especially people who own very low values of property. I suspect that these outliers represent cases in which I do not observe all of the property owned by a household. I address this issue by excluding the top and bottom 1% of observations by property value.

### *4.3. Quantitative results*

I begin by comparing children's propensity to leave the church across affiliations. Fig. 2 shows the propensity to leave the church in each affiliation. 36% of New Order children, 14% of Old Order childen, and 5% of Andy Weaver children leave the church. Since the main difference between the different affiliations is the severity of the punishment inflicted on a child who leaves the church, these results are consistent with the first prediction of my model, that Amish communities use punishments to reduce the fraction of children who leave. The results are also inconsistent with the first prediction of the club goods model, that strict religious communities use punishments to reduce the proportion of potential members who join the community.

I confirm the results from Fig. 2 using regression analysis to control for other observable differences between affiliations. I estimate the following linear probability model:

$$Y_{ihd} = \beta \mathbf{X}_{ihd} + \gamma_a + \epsilon_{ihd} \tag{3}$$

Here $Y_{ihd} = 1$ if child $i$ in household $h$ in district $d$ has left the church by 2010, and $Y_{ihd} = 0$ otherwise. $\mathbf{X}_{ihd}$ is a vector of control variables, $\gamma_a$ is a fixed effect for each affiliation, and $\epsilon_{ihd}$ is an error term. I cluster standard errors by district to account for possible correlated shocks within districts.

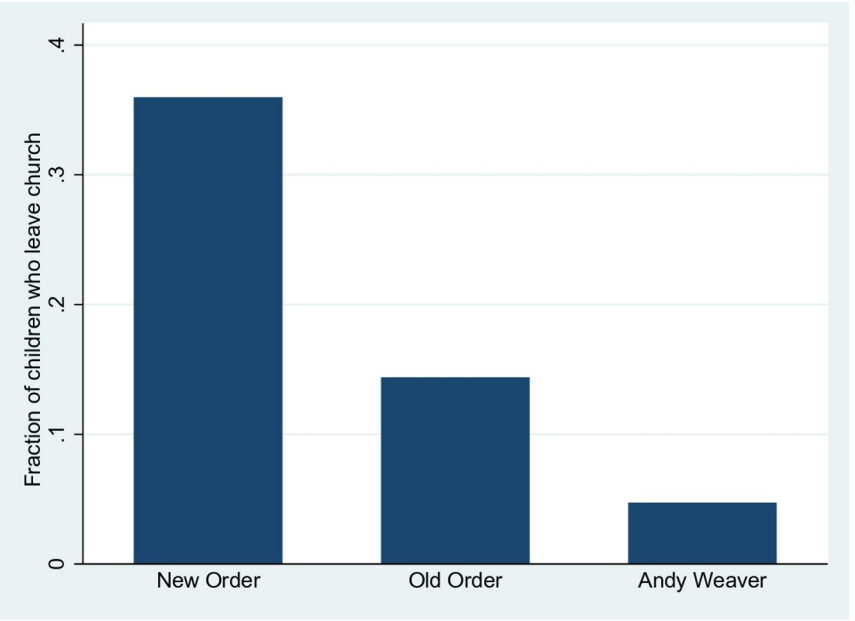


**Fig. 2.** Children's probability of leaving Amish church by affiliation. This chart shows the fraction of children who leave the church in each affiliation, using the entire sample.

One of the control variables in my regression is my measure of household permanent income. As discussed above, my primary measure of household permanent income is the logarithm of the value of buildings owned by the household. A difficulty with this measure is that 18% of children in my sample are members of households that do not own any property. These households are presumably not homeless, but rather are renting housing or, more likely in this context, living with other relatives. Since I think of the value of buildings as a proxy for the flow value of housing consumed by the household, households who own no property should be thought of as households for whom my measure of wealth is unobserved, not households for whom my measure of wealth is zero. For this reason I drop households who own no property when using parents' wealth as a control variable.

Table 4 presents estimates of (3). Column 1 presents results without control variables, and column 2 presents results with controls other than parents' wealth. Even after controlling for observable differences across affiliations, the results imply that New Order children are 21–22 percentage points more likely to leave the church than Old Order children, and that Andy Weaver children are 8–10 percentage points less likely to leave the church than Old Order children. The association between affiliation on children's

**Table 4**
Determinants of children's decision to leave the Amish church.

| | | | | | | |
|---|---|---|---|---|---|---|
| New Order affiliation | 0.216*** | 0.210*** | 0.211*** | 0.206*** | | |
| | (0.0374) | (0.0169) | (0.0384) | (0.0390) | | |
| Andy Weaver affiliation | −0.0965*** | −0.0799*** | −0.100*** | −0.0929*** | | |
| | (0.0169) | (0.0372) | (0.0161) | (0.0162) | | |
| Log(parents' wealth) | | | −0.0573*** | −0.0580*** | −0.0437* | −0.0445** |
| | | | (0.0210) | (0.0207) | (0.0221) | (0.0221) |
| Child is male | | 0.0548*** | 0.0576*** | 0.0576*** | 0.0572*** | 0.0566*** |
| | | (0.0133) | (0.0146) | (0.0146) | (0.0146) | (0.0145) |
| Child's age | | 0.00252 | 0.000875 | 0.000324 | 0.000670 | 0.0000812 |
| | | (0.00226) | (0.00263) | (0.00258) | (0.00277) | (0.00272) |
| Father is farmer | | −0.0939*** | −0.0823*** | −0.0736*** | −0.0861*** | −0.0761*** |
| | | (0.0174) | (0.0188) | (0.0204) | (0.0202) | (0.0209) |
| Father's age | | 0.00308** | 0.00523*** | 0.00582*** | 0.00586*** | 0.00652*** |
| | | (0.00122) | (0.00147) | (0.00139) | (0.00152) | (0.00149) |
| Number of siblings | | | | −0.00618** | | −0.00735** |
| | | | | (0.00281) | | (0.00281) |
| District fixed effects? | NO | NO | NO | NO | YES | YES |
| Observations | 3,627 | 3,627 | 2,936 | 2936 | 2,936 | 2,936 |

The dependent variable is equal to one if the child has left the church in 2010, and zero otherwise. The excluded affiliation is the Old Order affiliation. Parents' wealth is the value of buildings owned by parents. Observations in which parents' wealth is equal to zero, and the top and bottom 1% of observations by parents' wealth for observations in which parents' wealth is greater than zero, are dropped in columns 3 through 6. Columns 5 and 6 do not include affiliation fixed effects since these are collinear with district fixed effects. Standard errors clustered by district in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

propensity to leave the church is highly significant and stronger than the effect of any other observable factor. Columns 3 and 4 present results including parents' wealth as a control variable, dropping observations for which parents own no property, as discussed above. Columns 5 and 6 present results including district fixed effects. These regressions omit affiliation fixed effects, since affiliation and district are collinear.

There may be other factors that vary across affiliations and that may affect children's propensity to leave the church, and that are not fully accounted for by the control variables in Table 4. In particular, parents in affiliations that impose more severe punishments are also less likely to send their children to secular schools, and exert less effort to monitor their children's behavior. As discussed in Section 3.3.1, these differences may lead to differences in the mean level of intrinsic attachment to the Amish church $a$ across affiliations. Thus, the coefficients on the affiliation dummies in Table 4 may not represent the causal effect of shunning on children's propensity to leave the church. Because the average level of $a$ may be lower or higher in affiliations that impose more severe punishments, the causal effect of shunning on children's propensity to leave the church may be lower or higher than the effect indicated by the coefficients on the affiliation dummies in Table 4. It seems unlikely, however, that shunning has no causal effect on children's propensity to leave the church, for two reasons. First, the size of the coefficients in Table 4 seem too large to be explained merely by differences in preferences across affiliations alone. Second, if differences in preferences explained all of the variation in propensity to exit across affiliations, it would be difficult to understand why Amish communities impose the punishment of shunning at all.

Columns 3 through 6 of Table 4 show that children of wealthier Amish parents are less likely to leave the church than children of poorer Amish parents. This result is consistent with the second empirical prediction of my model, and is inconsistent with the second empirical prediction of the club goods model.

The result that children of richer parents are less likely to leave the church could be vulnerable to omitted variable bias. One way in which omitted variables may bias the result is if the value of buildings owned by parents measures parents' permanent income with error, and if this error is correlated with factors affecting children's choice of religious identity. For example, parents who desire more children may spend a larger fraction of their permanent income on housing, and parents who desire more children may also exert more effort to socialize their children with Amish values. In order to partially address this concern, in columns 4 and 6 of Table 4 I control for the completed fertility of each household. Reassuringly, adding this control does not qualitatively change any of the results.

As discussed previously, my preferred measure of parents' wealth is the value of buildings owned by parents. As a robustness check, I run the same regressions from columns 3 through 6 of Table 4 using the total value of buildings and land owned by parents as the measure of parents' wealth. In these regressions I drop observations from Wayne and Tuscarawas counties. Results are presented in Table 5. The results are qualitatively similar to the results from Table 4.

In order to explore the data further, I run separate regressions for each affiliation. The results are reported in Table 6. The main qualitative conclusions hold in each affiliation analyzed separately, although the precision of the estimates decreases. The correlation between parents' wealth and children's propensity to leave the church is insignificant in the Andy Weaver affiliation, suggesting that the individual parental punishments $p_i$ may be less important for retaining children in the church when the community punishment $p_c$ is very severe.

**Table 5**
Determinants of children's decision to leave the Amish church using alternate measure of parents' wealth.

| | | | | |
|---|---|---|---|---|
| New Order affiliation | 0.183*** | 0.178*** | | |
| | (0.0490) | (0.0492) | | |
| Andy Weaver affiliation | −0.0791*** | −0.0729*** | | |
| | (0.0181) | (0.0185) | | |
| Log(parents' wealth) | −0.0649** | −0.0652** | −0.0595** | −0.0597** |
| | (0.0276) | (0.0207) | (0.0275) | (0.0274) |
| Child is male | 0.0536*** | 0.0536*** | 0.0556*** | 0.0552*** |
| | (0.0150) | (0.0149) | (0.0149) | (0.0147) |
| Child's age | 0.000703 | 0.000392 | 0.000324 | 0.000219 |
| | (0.00312) | (0.00309) | (0.00331) | (0.00325) |
| Father is farmer | −0.0896*** | −0.0824*** | −0.0843*** | −0.0750** |
| | (0.0235) | (0.0256) | (0.0258) | (0.0258) |
| Father's age | 0.00565*** | 0.00608*** | 0.00607*** | 0.00681*** |
| | (0.00167) | (0.00157) | (0.00180) | (0.00177) |
| Number of siblings | | −0.00504 | | −0.00777** |
| | | (0.00364) | | (0.00345) |
| District fixed effects? | NO | NO | YES | YES |
| Observations | 2105 | 2105 | 2105 | 2105 |

The dependent variable is equal to one if the child has left the church in 2010, and zero otherwise. The excluded affiliation is the Old Order affiliation. Parents' wealth is the value of buildings and land owned by parents. Observations from Wayne and Tuscarawas counties are dropped. Top and bottom 1% of observations by parents' wealth are dropped. Standard errors clustered by district in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

**Table 6**
Determinants of children's decision to leave the Amish church, by affiliation.

| | New Order | | Old Order | | Andy Weaver | |
|---|---|---|---|---|---|---|
| Log(parents' wealth) | −0.229*** | −0.231*** | −0.0503** | −0.0513** | −0.00912 | −0.00968 |
| | (0.0754) | (0.0747) | (0.0251) | (0.0244) | (0.0300) | (0.0302) |
| Child is male | 0.114* | 0.115* | 0.0537*** | 0.0542*** | 0.0338** | 0.0336** |
| | (0.0572) | (0.0570) | (0.0181) | (0.0179) | (0.0139) | (0.0140) |
| Child's age | 0.0114 | 0.0114 | −0.000346 | −0.00124 | −0.000415 | −0.000591 |
| | (0.0140) | (0.0139) | (0.00294) | (0.00288) | (0.00327) | (0.00333) |
| Father is farmer | −0.0803 | −0.0888 | −0.0897*** | −0.0776*** | −0.0337 | −0.0307 |
| | (0.0921) | (0.0945) | (0.0224) | (0.0245) | (0.0229) | (0.0218) |
| Father's age | 0.0118** | 0.0115** | 0.00559*** | 0.00661*** | 0.000154 | 0.000250 |
| | (0.00463) | (0.00453) | (0.00183) | (0.00171) | (0.00101) | (0.000971) |
| Number of siblings | | 0.00674 | | −0.00910** | | −0.00167 |
| | | (0.0102) | | (0.00365) | | (0.00156) |
| Observations | 281 | 281 | 2160 | 2160 | 495 | 495 |

The dependent variable is equal to one if the child has left the church in 2010, and zero otherwise. Parents' wealth is the value of buildings owned by parents. Top and bottom 1% of observations by parents' wealth are dropped. Standard errors clustered by district in parentheses. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

### 4.4. Econometric issues

A potential problem with the regressions reported in Tables 4–6 is that they do not take into account the possibility that the error terms could be correlated across districts within a given affiliation. In order to take into account this possibility, it would be necessary to cluster standard errors by affiliation. However, I am not aware of any unbiased way to calculate clustered standard errors when each cluster corresponds to a "treatment" (where in this case each affiliation is a separate treatment) and there is only one cluster per treatment (see Mackinnon and Webb, 2018). Thus it does not seem to be possible to quantify the uncertainty in my estimates when error terms are correlated across districts within affiliations.

Although I cannot definitively rule out the possibility that the error term is correlated within affiliations, it is at least not obvious why there would be such correlations. Different affiliations are not geographically segregated within the Holmes county settlement, and so there are unlikely to be geographical shocks that affect members of some affiliations but not others. Districts within an affiliation do not coordinate changes to their rules, so there are no common policy shocks that affect all members of the same affiliation. All Amish people participate in the larger market economy along with non-Amish people, so there are unlikely to be wage or price shocks that affect members of some affiliations but not others.

Despite these arguments, there may other factors which generate common shocks across districts within affiliations. To the extent to which this is the case, the error terms reported in Tables 4–6 do not fully capture the uncertainty associated with my estimates. This possibility should be kept in mind when evaluating my results.

### 4.5. Further evidence

The results so far support the first two predictions of my model, namely that Amish community punishments reduce children's exit rates, and that children of richer Amish parents are less likely to leave the church than children of poorer Amish parents. I now discuss the evidence supporting the remaining three predictions of my model.

The third prediction of my model is that Amish church membership is transmitted across multiple generations, as children are induced to remain in the church by community punishments, and then grow up into parents who impose the same punishments on their children. In fact, as discussed in Section 2 and as shown in Table 1, Amish church membership has been transmitted within families for more than 300 years. The persistence of Amish church membership across generations is not explained by the club goods model of religion, which does not suggest any reason why children of Amish parents should be more likely to join the Amish church or other similar religious groups than children of secular parents with similar levels of wealth and human capital.

The fourth prediction of my model is that family dynasties that belong to the Amish church may have persistently lower welfare than family dynasties that do not belong to the Amish church. The exit rate from the New Order, which does not impose punishments on children who leave, supports this claim. As shown in Fig. 2, 36% of New Order children leave the Amish church, while it is implausible to believe that anywhere close to 36% of children from secular families with levels of wealth or human capital similar to the Amish join religions that are as strict as the Amish. This result suggests that the Amish have lower welfare than the non-Amish. My model explains how this situation can persist. In contrast, the club goods model implies that members of strict religious groups must have equal or higher welfare than non-members after controlling for endowments.

The fifth prediction of my model is that members of the Amish church may have very high fertility. In fact, the average number of children in households in my sample in 2010 is 7.8. This is not the correct measure of the Amish fertility rate, since my sample only includes parents who had children in 1988 and so it undersamples people with lower fertility. The result nevertheless provides some indication of the very high number of children in Amish households. A more careful estimate of the Amish fertility rate using a larger data set comes from Greksa (2002), who estimates that the completed fertility rate in the Geauga county, Ohio settlement is 7.7 births per woman. Thus, the Amish seem to have a much higher fertility rate even than secular populations with comparably low wages. As

discussed in Section 3, in the club goods model low wages among members of strict religious groups may cause high fertility, but the club goods model does not explain why members of strict religious groups have higher fertility than secular populations with comparable wages. My model suggests that, because religious rules are specifically designed to increase the amount of time that religious community members spend with their families, members of religious communities may have higher fertility than secular populations with similar wages.

## 5. Alternative explanations

I have argued that one function of Amish rules and institutions is to promote altruistic and cooperative behavior within families. It is worth emphasizing that this is not the only function of Amish rules and institutions. It seems likely that Amish rules and institutions also serve to promote cooperation between unrelated community members, as in the club goods model. Thus, my argument is not that the club goods model is wrong when applied to the Amish, but merely that it is incomplete. In this section, I discuss some additional possible alternative explanations of my results.

### *5.1. What is the purpose of shunning?*

Affiliations that have more severe forms of shunning are interpreted in my model as communities with higher values of the community punishment $p_c$, which reduces the utility of children who leave the community. It might be argued alternatively that the purpose of shunning is not to reduce the utility of children who leave the community, but rather to change the behavior of people who remain in the community by making it harder for them to interact with outsiders. If it is harder for community members to interact with outsiders, then they might spend more time contributing to activities that benefit the community. In other words, it might be argued that shunning should be interpreted not as increasing the community punishment $p_c$, but rather as increasing the strictness of religious rules $r$ in a community described by the club goods model.

Contrary to this hypothesis, qualitative evidence suggests that the Amish view shunning as a form of punishment and not simply as a way of preventing interaction with outsiders. Shunning-as-punishment is illustrated by the use of *temporary* shunning as response to minor infractions of the rules, as discussed in Section 2. Shunning-as-punishment is also illustrated by the observation that one of the goals of shunning is to encourage members who have left to rejoin the community. This function of shunning is explained in the Dordrecht Confession, one of the canonical statements of Amish beliefs. The Dordrecht Confession says that members should not consider those who have been shunned "as enemies, but admonish them as brethren in order to bring them to knowledge, repentance, and sorrow for their sins so that they may be reconciled to God and His church and consequently be received and taken in again" (Kraybill et al., 2013, p. 168). This suggests that the purpose of shunning is to affect the utility of people who have been shunned, rather than simply to prevent contact between community members and outsiders.

Despite this evidence, it may be the case that part of the effect of shunning is to encourage contributions to club goods by making it harder for community members to spend time interacting with outsiders. My argument is only that this is unlikely to be the complete explanation of the practice of shunning, just as the club goods model is unlikely to be the complete explanation of Amish institutions more generally.

### *5.2. Are there other factors that vary across Amish affiliations?*

In Section 3.3.1, I argued that the average level of children's intrinsic attachment to the Amish church $a$ may vary across affiliations, and that this may account for some of the observed differences across affiliations in children's propensity to leave the church. There may also be other factors that vary across affiliations and that affect children's propensity to leave the church. One such factor is variation in social norms across affiliations that may make it more difficult in some affiliations to discuss plans to leave the church. In lower affiliations, merely discussing the possibility of leaving may trigger social sanctions. This may make it harder for children in some affiliations to coordinate leaving decisions, or it may make it more difficult for children in some affiliations to acquire knowledge about non-Amish lifestyles that would be necessary to make an informed decision about the costs and benefits of leaving. This kind of variation may account for some of the differences across affiliations in propensity to exit the Amish church observed in Table 4, and suggests an additional reason why the estimates in Table 4 should not be interpreted as the causal effect of the punishment of shunning on propensity to exit the church.

Although differences in social norms governing open discussion of leaving the church may account for some of the variation in exit propensity across affiliations, this kind of variation does not help to explain the other empirical facts explained by my model, including the fact that children of richer parents are less likely to leave the church and the high Amish fertility rate. It also seems unlikely that this kind of variation could by itself account for the large size of the difference in propensity to leave the church across affiliations.

## 6. Conclusion

In this paper, I have argued that a central function of Amish rules and institutions is to increase the amount of time that Amish church members spend with their families. Amish parents impose these rules to benefit themselves at the expense of their children (although possibly to the benefit of their grandchildren). My theory helps to explain selection patterns into the Amish church, Amish cultural persistence and persistent Amish disadvantage relative to the non-Amish, and high Amish fertility rates.

Many other religious communities impose rules that affect family relationships. My theory may help to explain why these religious communities persist, and also why some of these religious communities may be persistently disadvantaged relative to secular populations. More positively, my theory suggests that some religious institutions may be valuable for helping to strengthen family ties and increasing investment in children's human capital. This result applies primarily to religious groups like non-ultra-Orthodox Jews which seem to emphasize investment in child quality over investment in child quantity, unlike the Amish. Further study of how religious institutions affect family relationships is an important topic for future research.

## References

Abramitzky, R., 2008. The limits of equality: insights from the israeli kibbutz. Q. J. Econ. 123, 1111–1159.
Akerlof, G.A., Kranton, R., 2000. Economics and identity. Q. J. Econ. 115, 715–753.
Becker, S.O., Woessmann, L., 2009. Was weber wrong? a human capital theory of protestant economic history. Q. J. Econ. 124, 531–596.
Berman, E., 2000. Sect, subsidy and sacrifice: an economist's view of ultra-orthodox jews. Q. J. Econ. 115, 905–953.
Bernheim, B.D., Shleifer, A., Summers, L.H., 1985. The strategic bequest motive. J. Polit. Econ. 93, 1045–1076.
Bisin, A., Verdier, T., 2000. Beyond the melting pot: cultural transmission, marriage, and the evolution of ethnic and religious traits. Q. J. Econ. 115, 955–988.
Bisin, A., Verdier, T., 2001. The economics of cultural transmission and the dynamics of preferences. J. Econ. Theory 97, 298–319.
Botticini, M., Eckstein, Z., 2012. The Chosen Few: How Education Shaped Jewish History 70–1492. Princeton University Press, Princeton, NJ.
Carvalho, J.-P., Koyama, M., 2016. Jewish Emancipation and Schism: Economic Development and Religious Change. J. Comp. Econ 44, 562–584.
Carvalho, J.-P., Koyama, M., Sacks, M., 2017. Education, Identity and Community: Lessons from Jewish Emancipation. Public Choice 171, 119–143.
Friedmann, Y., 2003. Tolerance and Conversion in Islam: Interfaith Relations in the Muslim Tradition. Cambridge University Press, Cambridge, UK.
Galperti, S., Strulovici, B., 2017. A theory of intergenerational altruism. Econometrica 85, 1175–1218.
Giuliano, P., Nunn, N., 2019. Understanding Cultural Persistence and Change. Working paper
Greksa, L., 2002. Population growth and fertility patterns in an old order amish settlement. Ann. Hum. Biol. 29, 192–201.
Hao, L., Hotz, J.V., Jin, G.Z., 2008. Games parents and adolescents play: risky behaviour, parental reputation and strategic transfers. Econ. J. 118, 515–555.
Heckman, J.J., 2011. The american family in black & white: a post-racial strategy for improving skills to promote equality. Daedalus 140, 70–89.
Hostetler, J.A., 1993. Amish Society. Johns Hopkins University Press, Baltimore, MD.
Hurst, C.E., McConnell, D.L., 2010. An Amish Paradox. Johns Hopkins University Press, Baltimore, MD.
Iannaccone, L., 1992. Sacrifice and stigma: reducing free-riding in cults, communes, and other collectives. J. Polit. Econ. 100, 271–291.
Kraybill, D.B., 2001. The Riddle of Amish Culture. Johns Hopkins University Press, Baltimore, MD.
Kraybill, D.B., Johnson-Weiner, K.M., Nolt, S.M., 2013. The Amish. Johns Hopkins University Press, Baltimore.
Kraybill, D.B., Nolt, S.M., 2004. Amish Enterprise. Johns Hopkins University Press, Baltimore, MD.
Kuran, T., Sandholm, W.H., 2008. Cultural Integration and Its Discontents. Rev. Econ. Stud. 75, 201–228.
Kymlicka, W., 1995. Multicultural Citizenship. Oxford University Press, New York, NY.
Lundberg, S., Pollak, R.A., Stearns, J., 2016. Family inequality: diverging patterns in marriage, cohabitation, and childbearing. J. Econ. Perspect. 30, 79–102.
Mackinnon, J.G., Webb, M.D., 2018. The wild bootstrap for few (treated) clusters. Econom. J. 21, 114–135.
Norenzayan, A., Shariff, A.F., Gervais, W.M., Willard, A.K., McNamara, R.A., Slingerland, E., Henrich, J., 2016. Parochial prosocial religions: historical and contemporary evidence for a cultural evolutionary process. Behav. Brain Sci. 39, 43–65.
Saez Marti, M., Weibull, J., 2005. Discounting and altruism to future decision-makers. J. Econ. Theory 122, 254–266.