JOURNAL FOR THE SCIENTIFIC STUDY OF RELIGION

Dynamics of Religious Group Growth and Survival

TONGZHOU CHEN
School of Mathematics
Georgia Institute of Technology

MICHAEL MCBRIDE
Department of Economics
University of California, Irvine

MARTIN B. SHORT
School of Mathematics
Georgia Institute of Technology

We model and analyze the dynamics of religious group membership and size. A group is distinguished by its strictness, which determines how much time group members are expected to spend contributing to the group. Individuals differ in their rate of return for time spent outside of their religious group. We construct a utility function that individuals attempt to maximize, then find a Nash equilibrium for religious group participation with a heterogeneous population. We then model dynamics of group size by including birth, death, and switching of individuals between groups. Group switching depends on the strictness preferences of individuals and their probability of encountering members of other groups. We show that in the case of only two groups—one with finite strictness and the other with zero—there is a parameter combination that determines whether the nonzero strictness group can survive over time, which is more difficult at higher strictness levels. We also show that a high birth rate can allow even the strictest groups to survive. Finally, we consider cases of several groups, gaining insight into strategic choices of strictness values and displaying the rich behavior of the model.

Keywords: *religious affiliation, population dynamics, religious growth, competition, pluralism.*

INTRODUCTION

Religion is a rich and dynamic phenomenon. Rates of religiosity rise and fall over time, perhaps to rise and fall again, and some religious groups grow while others shrink.[1] These patterns are influenced by societal-level factors including the promotion or suppression of certain or all forms of religious practice by government (Fox and Tabory 2008; Iannaccone 1991; Pew Research Center 2017b) and by generational differences in religiosity (Pew Research Center 2018). Yet, at the heart of this vibrancy are the choices made by thousands, millions, or even billions of individual persons deciding whether and how much to participate within a religious group, whether to remain in one group, switch to another group, or abandon religious participation altogether, and more. For example, the United States has experienced a decline in overall rates of religious participation during the last several decades (Voas and Chaves 2016), but during that time there has also been tremendous churn among American religious groups as individuals affiliate, disaffiliate, and switch affiliations (Putnam and Campbell 2012). While some groups

Acknowledgments: This work was funded by ARO MURI Grant W911NF-11-1-0332 and AFOSR MURI Grant FA9550-10-1-0569. For constructive comments, we thank colleagues, participants at the 2018 Association for the Study of Religion, Economics, and Culture Conference, and anonymous reviewers of our article.

Correspondence should be addressed to Michael McBride, 3151 Social Science Plaza, Irvine, CA 92697-5100. E-mail: mcbride@uci.edu

[1]There are many trends to consider. The Pew Research Center, for example, publishes various reports and statistics about trends in religiosity worldwide (Pew Research Center 2017a).

win in this competition for adherents, others lose, and the outcome is a religious marketplace with a diversity of forms of religious practice.

A body of research during the last few decades has drawn inspiration from economic models of markets and group production to link individual religious decisions with larger patterns of religiosity, yet this work provides only an incomplete understanding of the dynamic processes in religious markets. Two central thrusts of research are most relevant. The first identifies the locus of religious activity in the religious group, with the group serving as a collective-production entity that is susceptible to free-rider problems. The theory demonstrates that strict religious groups better confront free-rider problems than their less-strict counterparts, thereby enabling the strict groups to more successfully provide religious goods and services (Carvalho, Iyer, and Rubin forthcoming; Iannaccone 1992; McBride 2007). This insight helps to explain why strict churches have grown faster than less-strict churches during the last several decades (Iannaccone 1994). The second is that wide diversity of religious preferences can sustain a wide range of religious groups and practices when religious suppliers are allowed to enter and compete (McBride 2008, 2010; Stark and Finke 2000). As in the markets for other goods, there are differences in tastes for different types and styles of religion, and a diversity of forms of religious practice are needed to satisfy the diverse tastes. When entrepreneurs are allowed by government to supply this diversity, high religious pluralism can result as religious consumers with different tastes make their optimal affiliation decisions.

This article constructs and examines a dynamic model of religious competition that combines these two theories into a single framework. In so doing, we are able to reconcile what may at first appear to be a contradiction between the two views. While the latter theory recognizes the viability of all types of religious practice styles, the former implies that strict religious groups should outperform and possibly drive their less-strict competitors out of the market. We propose that the theories do work together but that additional factors are also relevant to a broader understanding of religious competition. In particular, we suppose that the relative success of different religious groups will depend not just on strictness but also on several other factors mentioned in the literature but not yet examined in a formal dynamic framework. Among these other factors are the strength of the cultural transmission of religious preferences across generations, the likelihood of exposure to other groups, the underlying distribution of preferences for nonreligious goods, and birth rates.

The incorporation of these features into our model draws inspiration from two other literatures. One literature establishes the vital role of demographic factors in the growth and decline of religious groups (Hout, Greeley, and Wilde 2001; Scheitle, Kane, and Hook 2011). The second literature uses dynamic models of cultural transmission to understand the spread of religious practices within and across generations (Bisin and Verdier 2000; Carvalho 2013). A primary contribution of our article is in combining several ideas in different literatures into one rich, dynamic model that yields new results. Specifically, we combine key conceptual elements of the club model of religious production, spatial models of religious competition, demographic models of religious growth and decline, and the dynamic models of cultural transmission.

Our mathematical and computational analysis reveals that the dynamics of a religious market with these many features are rich with notable emergent patterns. One finding is that very strict groups will die out unless they have sufficiently high birth rates and retention. This finding has been predicted in prior work (McBride 2015), and our analysis reveals that it is robust to several additional complexities in the market. A second finding is that moderate groups can survive if their strictnesses advantageously places them near the mean of the underlying distribution for nonreligious goods. That is, if most religious consumers have moderate opportunity costs of religion, then religious groups that only ask for a moderate commitment will appeal to a large fraction of the population, thus enabling those moderate groups to survive despite not producing religious goods as intensely as the much stricter groups. That moderate groups can survive and thrive has been noted before in explicit dynamic studies (Makowsky 2011; Montgomery 1999), but we demonstrate how several other factors not mentioned in those studies can also contribute to the persistent success of moderate groups.

A few prior attempts to explicitly model religious market dynamics are most similar to ours. Montgomery (1999) examined an environment with three strictness levels and the ability for religious groups to adjust strictness levels as their membership compositions changed over time. He found that the low-strictness groups do shrink and die out as predicted under some parametric configurations, but that they also survive and thrive under others. Makowsky (2011) allowed for a wider range of possible strictnesses to show why the less-strict groups might thrive. Lighter membership requirements allow for larger in-group heterogeneity in more moderate rates of free-rider mitigation, thus allowing for a degree of success in the market. Carvalho and Sacks (2018) provide a dynamic, theoretical model of a minority group competing with a rival mainstream culture. As in our setting, they allow for heterogeneous religious preferences, exit options, and competition, but we consider group survival and demographic change with more than two groups, while they examine the rise and fall of extremism in a two-group setting. Finally, Scheitle, Kane, and Hook (2011) simulated the growth of a hypothetical American religious group under different assumptions about in-group fertility and religious switching. They show that both fertility and switching play key roles, and that switching plays a particularly important role in the long run. Our model differs from these prior studies in its formal synthesis of the several factors mentioned earlier, i.e., cultural transmission across generations, differential rates of interaction among individuals of different groups, and variation in birth rates. Ultimately, our work demonstrates how these many factors contribute to the variety of outcomes possible in a religious market.

Ours is a theoretical study. The formalism of the mathematics enforces coherence and specificity in assumptions about the process of religious decision making. Such assumptions are meant to capture the essence of factors believed to be important, but we acknowledge that they are, by their nature, abstractions that emphasize some psychological and social forces at the expense of others. We pursue this approach because we believe that exploring the logical relationships between assumptions about religious choices and emergent patterns is itself a worthy endeavor. At the same time, we recognize that the value of purely theoretical study is further enhanced and, ultimately, assessed by its ability to explain and predict empirical patterns. Although explicitly testing our model empirically is beyond the scope of our article, we do discuss how our findings add to our current understanding and provide direction for future work. We also refer to specific real-life examples that we believe demonstrate key ideas in the article.

Single-Group Model

Individual Utility Function

Each individual must decide what portion of his or her time will be devoted to in-group activities, with the remaining portion devoted to out-group activities.[2] Without loss of generality, we assume that each person has a total time of 1, and the amount that individual i then devotes to in-group activities is denoted by t_i. It is assumed that in-group time is spent communally by the members of the group in production of "goods" that are distributed among the group members evenly, regardless of their individual contribution.

[2] We here describe the decisionmaker as acting instrumentally in pursuit of her goals as represented by a utility, whatever those goals may be. We find this interpretation to be the natural one in our setting because it fits the process of religious switching described later in the article. Using this language also eases the exposition. We note, however, that a utility representation does not have to be interpreted as instrumental choice. Utility representations assume consistency in choice, not instrumentality (Gaus 2007).

We define the in-group utility function U_{in} of person i as

$$U_{in} = c \left(\frac{\sum_j t_j}{N_g} \right)^{1/2}, \tag{1}$$

where $c > 0$ is a constant, N_g is the total population of the religious group g to which individual i belongs, and $\sum_j$ is taken over all members of group g, including member i.

As is standard in economic models, the utility function represents how the individual ranks different possible alternatives that may arise in the course of social interaction. A technical condition is that, as all $t_j \leq 1$, the sublinearity of the in-group utility in terms of the mean in-group time contribution causes the utility to be greater than the mean, reflecting the efficiency of group work. Of course, raising the mean in-group time contribution to any positive power less than 1 would do the same; we choose the power to be $1/2$ for simplicity. We also assume that all groups share the same factor c, such that no groups are inherently better at producing group utility than others. Finally, implicit in Equation (1) is the assumption that $N_g > 1$, otherwise the "group" would merely be a single individual. If $N_g = 1$, then $U_{in} = 0$.

Out-group activities yield a utility that is linear in the time spent outside the group, $1 - t_i$, such that

$$U_{out} = r_i(1 - t_i). \tag{2}$$

The factor $r_i \geq 0$ could reflect something like an hourly wage that can be earned at a job away from the group, but more generally reflects how much an individual personally values her time away from the group, during which she can engage in whatever activities she prefers. We will assume throughout that the values of r_i for the various individuals are, when not determined by inheritance (detailed in subsection "Birth, Death, and Inheritance"), drawn from a probability density $R(r)$.

The religious group is subject to potential free-rider problems. The amount that individual i earns from in-group activities may be dominated by the various t_j of the other group members, while the out-group utility is determined solely by the actions of individual i in such a way that time spent in-group returns a smaller U_{out}. Hence, many individuals may maximize their utility by simply choosing to contribute $t_i = 0$, which will maximize U_{out} while in many circumstances leaving U_{in} relatively unchanged. To combat such behavior, we allow the group to administer a punishment such that those members contributing less than what the group deems a minimal acceptable level will have their utilities reduced by an amount

$$Pun = \beta_g(\lambda_g - t_i)_+.$$

Here, $\beta_g \geq 0$ sets the overall scale for punishment within group g, while $\lambda_g \in [0, 1]$ is defined to be the "group strictness," which is the main trait that will serve to differentiate groups within our model, and $(\cdot)_+$ denotes the positive part of $(\cdot)$. The larger λ_g is, the stricter the group and the more time the group demands of its members. However, a member is only punished if she fails to contribute at least λ_g to the in-group activity. The punishment conceptualized here may be reflected in many ways: actual withholding of some of the group-produced goods from the individual, social pressures that may lead to ostracizing, or something else. Stricter groups have the means to enforce in-group norms, including norms related to in-group contributions.

The overall utility function U of person i in group g is equal to the sum of in-group production U_{in} and out-group production U_{out} minus the punishment Pun. Without loss of generality, we will scale all utilities by the common factor c and redefine r_i and β_g in terms of this standard scale, such that our final individual utility function is

$$\begin{aligned} U_i &= U_{in} + U_{out} - Pun \\ &= \left(\frac{\sum_j t_j}{N_g} \right)^{1/2} + r_i(1 - t_i) - \beta_g(\lambda_g - t_i)_+. \end{aligned} \tag{3}$$

Single-Group Nash Equilibrium

We now consider the case of a single group with parameters $\lambda_g = \lambda$ and $\beta_g = \beta$ and with a fixed set of members, such that the population size $N_g = N$ and the set of r values present within the group are unchanging. Then, to be determined for each individual in the group is what value of t_i she should choose. It is assumed that every individual is attempting to maximize her own personal U_i through this choice, but note that each person's U_i is also partially determined by the decisions of every other group member through the U_{in} term. Then, this is a classical game-theoretic problem, where the standard solution concept is the Nash equilibrium. In this case, a Nash equilibrium would be a set of in-group times of each member $\vec{t} = \{t_1, t_2, \ldots, t_i, \ldots, t_N\}$ with corresponding member utilities $\vec{U} = \{U_1, U_2, \ldots, U_i, \ldots, U_N\}$ such that there does not exist any alternative $\vec{t'} = \{t_1, t_2, \ldots, t_i', \ldots, t_N\}$ in which only member i has changed her choice such that the corresponding $\vec{U'} = \{U_1', U_2', \ldots, U_i', \ldots, U_N'\}$ would satisfy $U_i' > U_i$, for any i. In other words, in a Nash equilibrium, no individual i can increase her utility by unilaterally changing to a different t_i.

In principle there are five options for t_i that could possibly maximize U_i for an individual, given that all other t_j are fixed: 0, 1, λ, and two potential critical points we might call $\lambda < t_a < 1$ and $0 < t_b < \lambda$ located at

$$t_a = \frac{1}{4Nr_i^2} - T_i \equiv a_i + (1 - a_i)\lambda, \quad 0 < a_i < 1, \tag{4}$$

$$t_b = \frac{1}{4N(r_i - \beta)^2} - T_i \equiv b_i\lambda, \quad 0 < b_i < 1, \tag{5}$$

where $T_i = \sum_{j \neq i} t_j$. Note that due to constraints on the intervals where they may be located, t_a is only available to individuals whose r_i satisfies

$$\frac{1}{2\sqrt{N(1 + T_i)}} < r_i < \frac{1}{2\sqrt{N(\lambda + T_i)}}, \tag{6}$$

and t_b is only a valid critical point for individuals with

$$\frac{1}{2\sqrt{N(\lambda + T_i)}} + \beta < r_i < \frac{1}{2\sqrt{NT_i}} + \beta. \tag{7}$$

For fixed T_i , $U_i(1) = U_i(\lambda)$ at an r_i value within the range of values in Equation (6), with r_i values higher than this causing $U_i(1) < U_i(\lambda)$. Similarly, $U(\lambda) = U(0)$ at an r_i value in the interval in Equation (7), with r_i values higher than this causing $U_i(0) > U_i(\lambda)$. Also note that for fixed T_i and $t_i = 0$, 1, or λ, U_i is trivially nondecreasing in r_i. For t_a, which is a function of r_i, we have

$$U_i(t_a) = \frac{1}{4Nr_i} + r_i(1 + T_i), \tag{8}$$

Figure 1
An illustration of the existence and determination of the Nash equilibrium for the single-group case
[Color figure can be viewed at wileyonlinelibrary.com]

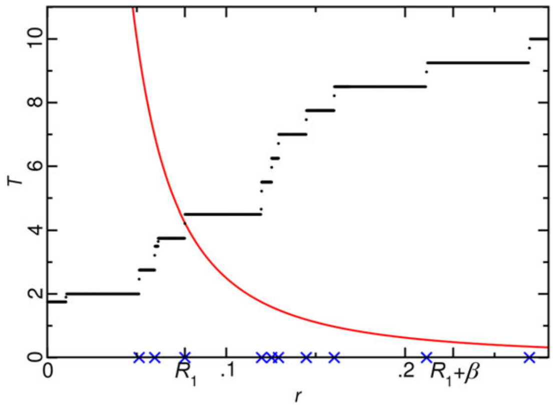

Notes: Here, the r_i of the group members is shown as blue X marks on the horizontal axis, the solid red curve is Equation (A2), and the discontinuous black curve is Equation (A1).

which is also increasing on the region of r_i values for which t_a is available, and ranges over values between $U_i(1)$ and $U_i(\lambda)$. Similarly,

$$U_i(t_b) = \frac{1}{4N(r_i - \beta)} + r_i(1 + T_i) - \beta(\lambda + T_i), \tag{9}$$

which is also increasing in r_i for all values for which it is available and ranges between the values $U_i(\lambda)$ and $U_i(0)$. Hence, for fixed T_i, the maximal value of U_i is a nondecreasing function of r_i, and the optimal t_i continuously transitions from $1 \to t_a(r_i) \to \lambda \to t_b(r_i) \to 0$ as r_i ranges from $0 \to \infty$. These results motivate the following Nash equilibrium for a single group.

Theorem 1. *Let $\vec{r}$ denote the list of r_i values for the N members of the group, sorted from least to greatest. There exists a number $R_1 > 0$ that is a function of $\vec{r}$, N, λ, and β such that, if all individuals with $r_i < R_1$ choose $t_i = 1$, all with $r_i = R_1$ choose t_a with a potentially specific value of a, all with $R_1 < r_i < R_1 + \beta$ choose $t_i = \lambda$, all with $r_i = R_1 + \beta$ choose t_b with a potentially specific value of b, and all with $r_i > R_1 + \beta$ choose $t_i = 0$, the system is in a Nash equilibrium.*

The proof (and all other proofs) are found in the Appendix. Intuitively, the amount of time that each group member devotes to in-group activities in the equilibrium is decreasing in that person's opportunity cost of time r_i. Those with very low r_i will devote all or most of their time to the group while those with higher r_i will devote less and less as r_i increases. An illustration of the Nash equilibrium is shown in Figure 1. Here, 10 individuals with r_i values shown as blue X marks on the horizontal axis are members of a group with $\lambda = .25$ and $\beta = .15$. The solid red curve is Equation (A2), while the discontinuous black curve is Equation (A1), with $a \approx .62$ and $b = .6$ (though the value of b is unimportant in this particular case, as no individuals play t_b). Note the single point of intersection of these two curves, guaranteeing that the Nash equilibrium exists, which in this case occurs at the r value of one of the group members, and is labeled as R_1. Then the two individuals with $r_i < R_1$ will choose $t_i = 1$, the single individual at $r_i = R_1$ will choose $t_i = a + (1 - a)\lambda$ with $a \approx .62$, the six individuals with $R_1 < r_i < R_1 + \beta$ will choose $t_i = \lambda$, and the one remaining individual will choose $t_i = 0$. Consistent with the intuition, group members' time contributions decrease in r_i values.

Ideal Strictness and Punishment Levels

The previous section considers how a variety of individuals already within a group and with varying r_i will determine their t_i given the group strictness λ and punishment factor β. Here, we study a somewhat different problem, focusing on one value of r_i at a time and asking what are the ideal strictness and punishment factor for individuals with that particular r_i value. To do so, we assume that N people with identical parameters $r_i = r$ are originally unaffiliated, meaning they are not currently a member of any group and receive only U_{out}. They would like to form a group together of strictness λ to get a higher payoff than being unaffiliated. We assume for now that, since all individuals have the same r_i, they will all choose the same t_i; we will prove later that this can be made to be so. If this is the case, then they will all choose $t_i = \lambda$ of the group they have formed; choosing $t_i = 0$ leaves them no better off than they currently are being unaffiliated, and choosing $t_i = 1,\ t_a$, or t_b would be equivalent to forming a group with λ corresponding to that specific choice. Then, each individual will receive payoff

$$U = \sqrt{\lambda} + r(1 - \lambda). \tag{10}$$

This payoff is maximized for ideal strictness $\lambda = 1/4r^2$. Note, though, that since $\lambda \leq 1$ by definition, if $r < 1/2$, the ideal strictness is simply $\lambda = 1$. For this reason, in the remainder of the article we generally assume that all $r_i \geq 1/2$. If the group adopts the ideal strictness level, it will end up with a maximized utility of $U_{max} = r + 1/4r$.

However, we must now determine whether the above situation is a Nash equilibrium as discussed above. Specifically, with all N individuals having the same $r_i = r$, in a group of strictness $\lambda = 1/4r^2$, we require all individuals playing $t_i = \lambda$ to result in a threshold R_1 such that $R_1 < r < R_1 + \beta$, which is the condition needed for the Nash equilibrium. In this case, $T = \lambda N$, so that $R_1 = r/N < r$ from Equation (A2). But, this is only a Nash equilibrium if $r < R_1 + \beta$, so that we need $\beta > r(1 - 1/N)$. To allow for a group of any potential size, then, we could simply use $\beta = r$. With this being the case, the total punishment for a person were she to choose $t_i = 0$ instead of $t_i = \lambda$ would be $P = \beta\lambda = \sqrt{\lambda}/2 \leq 1/2$, which is less than the U_{in} received from being in the group. Hence, a minimum punishment level is necessary to guarantee that no individuals in this group will be tempted to switch from $t_i = \lambda$ to $t_i = 0$, but this punishment level is bounded and need not completely remove the benefits of being in the group (U_{in}) to be entirely effective. This is a classical example of the "free-rider" problem, which in this case can be solved with sufficient punishment for free riding.

It is possible to set a bounded punishment level that dissuades any of the $r_i = r$ members of the group from deviating from the choice $t_i = \lambda$ without completely removing that member's U_{in}. But can the same be done to dissuade outsiders with differing $r_i > r$ from joining the group and playing $t_i = 0$? Imagine another individual with $r_i > r$ joining the existing group, so that N increases by one, but λ and β are as indicated above. Any such individual can only decrease the value of R_1, but never so much that $R_1 + \beta < r$ given our β value, so all the original individuals will always continue to play $t_i = \lambda$. However, the added individual will only free ride if her utility from doing so is greater than her utility from choosing $t_i = \lambda$. That happens if

$$\sqrt{\lambda} + r_i(1 - \lambda) < \sqrt{\lambda N/(N+1)} + r_i - \beta\lambda, \tag{11}$$

which would only necessarily be the case in arbitrarily sized groups if $\beta < r_i$. At the same time, though, this new individual will only join the group to free ride if the utility of doing so is greater than the utility of simply being unaffiliated, which only happens if

$$\sqrt{\lambda N/(N+1)} + r_i - \beta\lambda > r_i. \tag{12}$$

So, by choosing $\beta = \sqrt{N/\lambda(N+1)}$, the group can prevent all possible free riding. Of course, in this case the punishment for free riding is $Pun = \beta\lambda \approx \sqrt{\lambda}$, so that the punishment is to simply remove the entirety of U_{in}.

One can also consider what may happen if an individual with $r_i < r$ joins the group. Again, note that adding such individuals can only decrease R_1, but never so much so that $R_1 + \beta < r$, so any such added individuals will never cause the existing group members to become free riders. The newly added individual can therefore only play $t_i = 1$, λ, or a linear combination of the two (t_a). Upon adding such an individual, the intersection between Equations (A1) and (A2) (see Appendix) can only occur at one of three places: at an $R_1 > r_i$ where Equation (A1) yields $\lambda N + 1$ and $t_i = 1$, at $R_1 = r_i$ where Equation (A1) yields $\lambda(N+1) + a(1-\lambda)$ and t_i is a linear combination of 1 and λ, or at an $R_1 < r_i$ where Equation (A1) yields $\lambda(N+1)$ and $t_i = \lambda$. But in the first case, $R_1 = 1/2\sqrt{(N+1)(1+\lambda N)} < 1/2$ so r_i cannot satisfy $r_i < R_1$, as we have already constrained all $r_i \geq 1/2$, so the first case is ruled out by prior assumptions, and the new individual does not play $t_i = 1$. The third case gives $R_1 = r/(N+1)$, so any $r_i > r/(N+1)$ will cause the new individual to choose $t_i = \lambda$, which is quite likely in a very large group. Finally, any $1/2 \leq r_i \leq r/(N+1)$ will cause this individual to play a linear combination of 1 and λ. In fact, this argument is easily extended to a situation in which there are many r_i values present, all $\leq r$, in which case at most the individual(s) with the smallest r_i may play a linear combination of 1 and λ, but all others will play λ.

Multiple-Group Model

Dynamics of Group Membership

We now turn to a more dynamic situation in which there are potentially several groups to choose from, and individuals may be changing their affiliations over time. The overall goal of the model will be to describe how the sizes of religious groups vary over time given the distribution of r values in the population, the strictness values of the various groups, and other considerations discussed below. This variation is of course directly determined by the rate at which individuals enter a group versus the rate at which they leave a group, and these rates are themselves determined by two mechanisms that we will consider: (1) birth and death of group members and (2) individuals switching group affiliation. Both factors are important for understanding the trajectory of religious group membership (Scheitle, Kane, and Hook 2011). We cover each effect separately below, and summarize the model in Figure 2.

Before describing these effects in detail, we define a few more aspects of the model. First, we assume that the values r_i for all of the individuals within the entire society, encompassing all existing groups, are (at least initially) derived from a probability density $R(r)$. Second, we will at times wish to consider a special group known as the "unaffiliated group" whose strictness is by definition 0 and for whose "members" there is no U_{in}. As the name implies, this group really encompasses all those individuals who are not affiliated with any standard group with $\lambda > 0$; as such, these individuals are not partaking in any in-group activities whatsoever and all choose $t_i = 0$.

Given the results above regarding ideal strictness and punishment levels, we make the following simplification moving forward. Specifically, we will assume that all groups will select a β that dissuades any possible free riding, and that therefore all of the members of any group will simply play $t_i = \lambda$. The only approximation involved in this assumption is that we are ignoring the possibility of the member(s) with the smallest r_i values playing a linear combination of 1 and λ, but this is a very borderline case that should not affect the remainder of the results.

Figure 2
A flowchart detailing the various pieces of the multigroup model

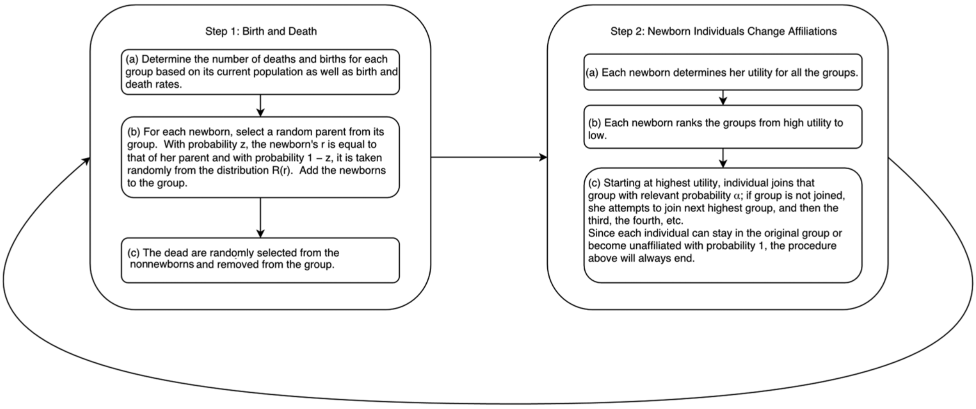

Birth, Death, and Inheritance

We assume that each group has a per capita birth rate b_g, which could potentially be group dependent, but that each group has the same per capita death rate d. When individuals die, they are simply removed from the population, thus decreasing N_g by 1 for the group to which they belonged. When individuals are born, it is assumed that their initial affiliation is the same as that of their parent(s), so they will increase N_g by 1 for the group that they are born into. In addition, whenever a new individual is born, with probability z her r_i is equal to that of her parent and with probability $1 - z$ her r_i is taken randomly from the distribution $R(r)$. Parameter z thus captures the degree of in-group cultural transmission from parent to child. If all groups exhibit the same birth rate, this mechanism will cause the expected distribution of r values within the population to be stationary in time, and equal to the distribution $R(r)$.

Changing Affiliation

We assume that every individual has one chance in life to change group affiliation. This opportunity is given to each individual effectively directly after his or her birth; for simplicity, though, this is meant to capture the possibility of switching groups once an individual becomes an independent adult.

Group switching is conceptualized in the following way for individual i who is currently a member of group g and values out-group activities at rate r_i. First, given the set of M groups and their corresponding strictness values $\lambda_{g'}$, individual i could in principle associate with any of the groups g' and thereby obtain utility

$$U_{ig'} = \sqrt{\lambda_{g'}} + r_i(1 - \lambda_{g'}). \tag{13}$$

The only exception to the above formula is for the unaffiliated group, which we will label as $g' = 0$, whereby the utility is simply $U_{i0} = r_i$.

In a system with perfect and complete information, each individual would simply determine which g' provides the maximum utility and choose that group. However, an important aspect of switching religious groups is exposure to the group: if one is exposed to members of a group frequently, the chance of switching to that group should be higher than that of switching to a group whose members you have never met, all else being equal. This motivates us to define what we will call the exposure probability of an individual currently in group g to members of group g'

$$\alpha_{gg'} = 1 - \left[1 - \frac{(1-\lambda_{g'})N_{g'}}{\sum_k (1-\lambda_k)N_k}\right]^{s(1-\lambda_g)}, \tag{14}$$

where $s > 0$ is a model parameter and N_g is the number of members in group g.

Underlying this exposure probability is the assumption that during out-group time, all members of society are well mixed. Then, for any given out-group chance encounter of an individual, the probability that the person met is in group g' is simply proportional to the number of people from g' who are spending time out of group at that moment, which is $(1-\lambda_{g'})N_{g'}$. Of course, the number of out-group chance encounters that an individual in g experiences throughout her life up until the moment she may select to switch groups is proportional to the amount of time she spends in general outside the group, represented by the $s(1-\lambda_g)$ term. Then, $\alpha_{gg'}$ is the probability that our individual has had at least one encounter with a person from group g' by the time she may choose to switch groups. The exceptions to Equation (14) are for the case $g' = g$, in which case $\alpha_{gg} = 1$ since everyone has had encounters with members of their own group for certain, and the case $g' = 0$, for which we also assume $\alpha_{g0} = 1$ since one need not encounter unaffiliated individuals in order to "join" the unaffiliated "group."

Given then the utilities $U_{ig'}$ and exposure probabilities $\alpha_{gg'}$, switching for individual i currently in g occurs in the following way. First, we sort the groups such that g' is in front of g'' if

1. $U_{ig'} > U_{ig''}$.
2. $U_{ig'} = U_{ig''}$ and $\alpha_{gg'} > \alpha_{gg''}$.
3. $U_{ig'} = U_{ig''}$, $\alpha_{gg'} = \alpha_{gg''}$, and $\lambda_{g'} < \lambda_{g''}$.

This leaves us with a permutation of groups, denoted by $\sigma(j)$, where $j = 0, 1, \ldots, M-1$. Then we simply march down the permutation starting with $j = 0$, at each point determining whether i chooses group $\sigma(j)$ via the exposure probability $\alpha_{g\sigma(j)}$ until she probabilistically joins a group. This procedure is attempting to assign every individual to her highest utility group, but only does so if there was sufficient exposure to that group, else the next highest utility group is attempted, etc. Note that this procedure will always end with i joining some group, because α_{gg} and α_{g0} are both 1, so in the extreme case she can always stay in her current group or become unaffiliated. Because of this, an individual will never switch to a group with lower utility than her current group (though it may be of equal utility).[3]

Also note that if person i spends all her time in in-group activities, which should only occur if $\lambda_g = 1$, then all the $\alpha_{gg'}$ are zero except for α_{g0} and α_{gg}. Such an individual has no opportunity

[3]A potential criticism of this formulation of religious switching is that it applies only to very specific settings, such as Christian America. This criticism has a degree of merit but perhaps less merit than might be first thought. Its merit consists of the reality that political, legal, and social institutions and regulations have inhibited the free movement of individuals in and out of religious groups for much of human history. The type of switching we assume requires the existence of multiple groups, but the presence of alternative religious groups has traditionally been suppressed in many countries and continues to be suppressed in many countries today (Fox and Tabory 2008). However, it could also be argued that, without such external legal and social constraints, the movement of individuals in and out of religious groups would be quite natural. Moreover, the model could be adjusted to include high social and legal costs to religious switching. Such costs can be easily added to the utility function. We do not include them for two reasons. First, the inclusion of these costs would result in very obvious changes to the model, such as a lower rate of switching out of the legally privileged group and the shrinkage of alternative groups. Second, we are specifically interested in the dynamics of a setting with free flows as they serve as an important benchmark for comparison. As an aside, we do note that switching is actually not unique to today's United States, but there have been historical cases of relatively free religious markets (Brekke 2016).

to switch to any but the unaffiliated group, or simply remain in her current group. Furthermore, if any group g has $\lambda = 1$, then for any other group $g' \neq g$ we have $\alpha_{g'g} = 0$. Therefore, the size of a group with strictness 1 will never grow due to new members joining from the outside, and can only drop if members choose to become unaffiliated.

Differential Equation Model for Group Size

Given the dynamics specified above, one could implement a discrete, agent-based model immediately to observe how the system evolves. Here, we instead cast the problem in terms of ordinary differential equations, so as to achieve a greater ability to understand the model analytically. The assumption here is that the overall population size is very large, so that taking an expectation of the stochastic dynamics may yield a good approximation to the discrete case.

We assume going forward that no two groups share the same strictness level: $\lambda_g \neq \lambda_{g'}$ for all $g \neq g'$. Then, given the number of groups M and their various strictness levels, each potential r value from the distribution $R(r)$ can be classified by its permutation $\sigma_r(j)$ of the groups strictly in terms of the utility of the groups to a person with parameter $r_i = r$. As such, we can divide the total population into a finite number S of subpopulations, each of which is labeled by the permutation of groups σ that all members of that subpopulation have in common. Then our model need only track the number of individuals in group g that are members of subpopulation σ over time, labeled as $n_{g\sigma}(t)$. Note that $\sum_{\sigma}^{J} n_{g\sigma}(t) = N_g(t)$. We define the fraction of the distribution $R(r)$ that encompasses subpopulation σ to be f_σ. Then the differential equation governing the expected value of $n_{g\sigma}(t)$ is

$$\frac{dn_{g\sigma}}{dt} = -n_{g\sigma} + \sum_{g'=0}^{M-1} b_{g'} \left[z n_{g'\sigma} + f_\sigma (1-z) N_{g'} \right] p_{g'g\sigma}. \tag{15}$$

Here, we have scaled time by the common death rate d, so that b_g is now the relative (to death) birth rate of group g. The new term $p_{g'g\sigma}$ is simply the probability that when a person currently in g' is given the opportunity to switch groups, she switches to group g, conditional on being a member of subpopulation σ. If group g takes position J in ordering σ, then

$$p_{g'g\sigma} = \alpha_{g'g} \prod_{j=0}^{J-1} \left(1 - \alpha_{g'\sigma(j)}\right). \tag{16}$$

That is, in order to choose g given preference σ, one needs to not choose any of the groups $\sigma(j)$ with $j < J$ that are higher in the ordering, and then needs to choose to join g, with all of the probabilities dictated by the various $\alpha_{g'\sigma(j)}$.

In general, it is more convenient to consider the size of a given population relative to the total population size N, so we now recast Equation (15) in terms of new variables $\tilde{n}_{g\sigma} = n_{g\sigma}/N$ and $\tilde{N}_g = N_g/N$. Given that the differential equation for N in time units scaled by the common death rate d is

$$\frac{dN}{dt} = -N + \sum_{g'=0}^{M-1} b_{g'} N_{g'},$$

we obtain the differential equation

$$\frac{d\tilde{n}_{g\sigma}}{dt} = \sum_{g'=0}^{M-1} b_{g'} \left(\left[z\tilde{n}_{g'\sigma} + f_\sigma (1-z) \tilde{N}_{g'} \right] p_{g'g\sigma} - \tilde{N}_{g'} \tilde{n}_{g\sigma} \right), \tag{17}$$

where the p values are the same as above, and the α values still follow Equation (14) but with N_j replaced with $\tilde{N}_j$. In general, we will use Equation (17) from now on with all tildes dropped, and all references to sizes of populations will be scaled by total population size, which may or may not be constant.

Two Groups

In this section, we present some analytical results for the simplest nontrivial case, that of two groups. The groups here are the unaffiliated group labeled 0 and an affiliated group labeled 1 with some strictness value $\lambda_1 = \lambda > 0$. Since there are only two groups, we only have $S = 2$ subpopulations with different ordering preferences σ, $\{0, 1\}$ and $\{1, 0\}$, which we will refer to as simply $\sigma_0 = 0$ and $\sigma_1 = 1$, respectively. Then, let $f_1 = f$ so that $f_0 = 1 - f_1 = 1 - f$. Finally, note that $N_0 + N_1 = 1$.

According to the rules of switching:

1. People can always stay in the original group if they prefer that group. Thus, $\alpha_{00} = \alpha_{11} = 1$. Then $p_{000} = p_{111} = 1$ and $p_{010} = p_{101} = 0$.
2. People can always switch to the unaffiliated group if they prefer it. Thus, $\alpha_{10} = 1$. Then $p_{100} = 1$ and $p_{110} = 0$.
2. People who are originally in group 0 and prefer group 1 can switch to group 1 with probability

$$p_{011} \equiv p = \alpha_{01} = 1 - \left[1 - \frac{(1-\lambda)N_1}{1-\lambda N_1}\right]^s. \tag{18}$$

4. People who are originally in group 0 and prefer group 1 will nonetheless stay in group 0 with probability $p_{001} = 1 - p$.

First, consider the case in which all birth rates have the same value, which we set to unity. Then Equation (17) becomes

$$\begin{cases} \frac{dn_{00}}{dt} &= -n_{00} + z(n_{00} + n_{10}) + (1-f)(1-z) \\ \frac{dn_{01}}{dt} &= -n_{01} + [zn_{01} + f(1-z)N_0)]\,(1-p) \\ \frac{dn_{10}}{dt} &= -n_{10} \\ \frac{dn_{11}}{dt} &= -n_{11} + [zn_{01} + f(1-z)N_0]\,p + [zn_{11} + f(1-z)N_1]\,. \end{cases} \tag{19}$$

At equilibrium, then, we clearly have $n_{10} = 0$ and $n_{00} = 1 - f$. Given that the total population size adds to unity, we can recast the remaining two equations in terms of a single variable, which we will choose to be $n_{11} = N_1 = n$. For notational simplicity, let $K = z + f(1 - z)$ (so $f \leq K \leq 1$). Then at equilibrium, we have

$$\frac{dn}{dt} = (f - Kn)p(n) + Kn - n \equiv g(n) = 0, \tag{20}$$

where

$$p(n) = 1 - \left[\frac{1-n}{1-\lambda n}\right]^s. \tag{21}$$

Equation (20) is significant because it sets the equilibrium size n of group 1 in the two-group case, given the parameters. With extremely high λ there will be very little switching into the very strict group as would be expected because of its very high strictness. In the extreme case $\lambda = 1$,

Figure 3
The size of group 1 at equilibrium is plotted as a function of its strictness λ for varying values of z, with $s = .75$ fixed and the distribution $R(r)$ chosen at each λ such that $f = .5$ [Color figure can be viewed at wileyonlinelibrary.com]

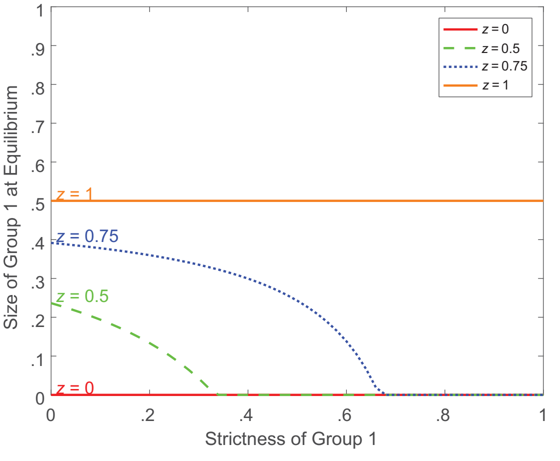

we have $p = 0$ so that at equilibrium $n = 0$ unless $K = 1$, which can only happen if $f = 1$ and/or $z = 1$. For cases $\lambda < 1$, we have the following formal result that matches our intuition.

Theorem 2. *If* $g'(0) = fs(1 - \lambda) + K - 1 \leq 0$, *then the equation* $\frac{dn}{dt} = g(n)$ *has only the trivial equilibrium point* $n = 0$ *and it is stable. Otherwise, the trivial equilibrium point becomes unstable and the equation has another stable equilibrium at a point* n_0 *in* $(0, f)$.

We gain several insights from this result. One is that a group with strictness $\lambda > 0$ will only survive as a finite fraction of the population at equilibrium if inequality

$$fs(1 - \lambda) + f(1 - z) > 1 - z \tag{22}$$

is satisfied. It could be imagined that the inheritance rate z and the parameter s are not under the control of any of the groups, but λ, and thereby f, are. Note that f is given by the fraction of $R(r)$ for which being in group 1 is preferable to being in group 0, and is given by

$$f = \int_{1/2}^{1/\sqrt{\lambda}} R(r)dr.$$

Hence, the left-hand side of the inequality (22) is decreasing in λ, so that more strict groups are more apt to die out over time than less strict groups. In this case, we should not see ultra-strict groups because they would die out. For any given R, z, and s, inequality (22) implies a maximal strictness that the group can adopt and still continue to survive in the long run.

So why then do we see relatively ultra-strict groups such as the Amish and Hasidic Jews survive and grow? One way to combat the threat of dying out is to raise the probability of inheritance z (see Figure 3). Inheriting the r value from their parents means that the children also have the same preference as their parents. Hence, if somebody is already in her optimal group, her descendants who inherit her r value are not going to make any switch, causing the group to maintain its size from internal birth more so than in cases where inheritance is low. In an extreme case that $z = 1$ and everybody can take the r value from her ancestors, at the equilibrium, everybody will stay in her favorite group. The Amish are noted for the very high rate of retaining

Figure 4
The size of group 1 at equilibrium is plotted as a function of its strictness λ for varying values of s, with $z = .5$ fixed and the distribution $R(r)$ chosen at each λ such that $f = .5$ [Color figure can be viewed at wileyonlinelibrary.com]

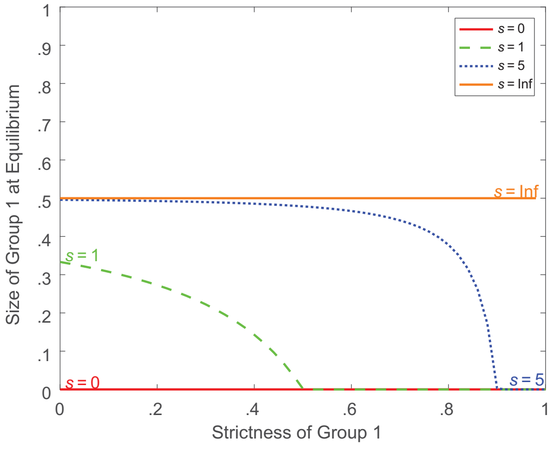

their youth within the group; one estimate puts retention at about 90 percent (Kraybill, Nolt, and Weaver-Zercher 2012).

Increasing the model parameter s has a similar effect (see Figure 4). That is because the larger s is, the greater the $\alpha_{gg'}$s are, so there is a higher probability for people to switch to their optimal group. When s goes to ∞, α_{01} goes to 1 unless $\lambda_1 = 1$. In this extreme case, everybody will end up in their optimal group if $\lambda_1 \neq 1$. On the other hand, when s goes to 0, α_{01} goes to 0. So people in group 0 preferring group 1 are never capable to make the switch. However, people in group 1 can always become unaffiliated if this is a better choice. Then group 1 keeps losing people and will eventually die out.

The results above only apply to the case in which the two groups have a common birth rate, and show that in some circumstances the strict group will die out. In reality, we often see that stricter religious groups have higher birth rates relative to less strict groups (Frejka and Westoff 2008; Scheitle, Kane, and Hook 2011). Amish families, for example, typically have about seven children (Kraybill, Nolt, and Weaver-Zercher 2012). A high birth rate in the ultra-strict group can counteract the relatively low rates of conversion into the group that result from the relatively low exposure that outsiders receive to the group and possibly an inherently smaller fraction of the population for whom such a strict group is ideal. This in turn could potentially allow a stricter group to continue to survive by increasing its internal growth rate. Therefore, we now consider the case in which group 0 retains birth rate 1, but group 1 has birth rate $b \geq 1$. Then the differential equations governing the fractional populations are

$$\begin{cases} \frac{dn_{00}}{dt} &= z(n_{00} + bn_{10}) + (1-f)(1-z)(N_0 + bN_1)(1 - n_{00}) \\ \frac{dn_{01}}{dt} &= [zn_{01} + f(1-z)N_0](1-p) - (N_0 + bN_1)n_{01} \\ \frac{dn_{10}}{dt} &= -(N_0 + bN_1)n_{10} \\ \frac{dn_{11}}{dt} &= [zn_{01} + f(1-z)N_0]\,p + b\,[zn_{11} + f(1-z)N_1] - (N_0 + bN_1)n_{11}. \end{cases} \tag{23}$$

As in the case above, we again find that at equilibrium $n_{10} = 0$, so we can cast the equilibrium equations in terms of $n_{11} = N_1 = n$, with $N_0 = 1 - n$ still. After some algebraic manipulations, we find that the population n at equilibrium satisfies

$$\left\{-K(b-1)n^2 + n\,[(b-1)(Kz + f(1-z)) - K(1-z)] + f(1-z)\right\} p(n)$$

Figure 5
The size of group 1 at equilibrium is plotted as a function of its strictness λ for varying values of b, with $s = .75$ and $z = .5$ fixed and the distribution $R(r)$ chosen at each λ such that $f = .5$ [Color figure can be viewed at wileyonlinelibrary.com]

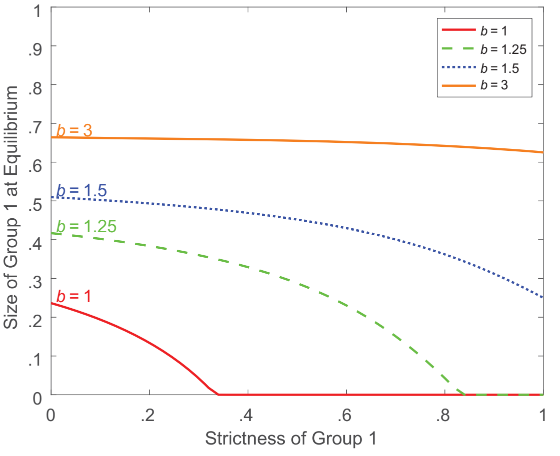

$$+\left\{-(b-1)^2n^2 + n(b-1)(bK+z-2) + (1-z)(bK-1)\right\} n = g_b(n) = 0. \qquad (24)$$

Then the following result holds.

Theorem 3. *For any* $0 < \lambda \leq 1$, *there exists a minimal birthrate* b_{min} *that allows for survival of the stricter group at equilibrium.*

This result indicates that for very strict groups, a higher than average birth rate may be necessary for long-term survival, but also guarantees that this is always possible. Indeed, because b_{min} is increasing in λ, the stricter a group wishes to be, the greater the birth rate necessary for survival, assuming the group could not survive at $b = 1$. Moreover, because the largest possible value of b_{min}, occurring at $\lambda = 1$, could not be greater than $1/z$, any group with a birth rate higher than this is guaranteed to survive regardless of its λ. Figure 5 illustrates that if the strict group has a higher birth rate, while all the other parameters are fixed, it can still survive.

Generating a high birth rate and high rate of inheritance appears to be the strategy that ultra-strict groups like the Amish and ultra-orthodox Jews follow to thrive despite their very high strictness. Other strict (but not ultra-strict) groups like the Mormons also maintain fertility rates higher than the national average. But their higher rates of participation in regular culture also imply that they should experience higher rates of religious switching both in and out of the group. Indeed, in any given year, about one-third of new membership in the Mormon Church comes from births, while two-thirds comes from converts (The Church of Jesus Christ of Latter-Day Saints 2018).

Three Groups

In the previous section, we determined the conditions under which a single group with positive strictness level may survive at equilibrium alongside the unaffiliated group. Of course, in the real religious marketplace, many groups simultaneously coexist, so one would ideally want to analyze multigroup cases within the context of our model. Unfortunately, the model's complexity increases very rapidly with the number of groups due to two main factors: the possibility of inheritance of r values and the rapid growth in the number of group preference orderings σ with number of groups M.

Figure 6
Equilibrium sizes of three groups as the strictnesses of the two affiliated groups vary
[Color figure can be viewed at wileyonlinelibrary.com]

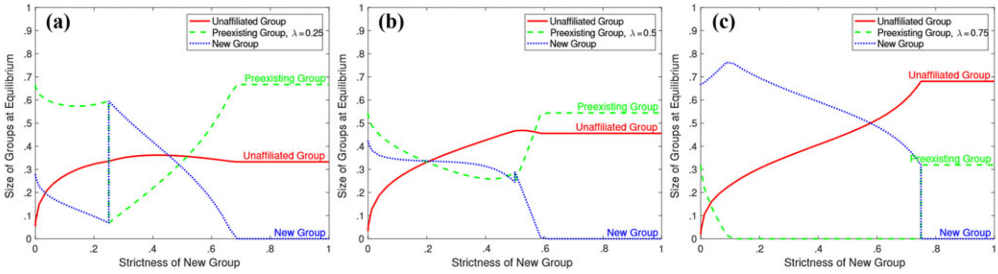

Notes: In all cases, $s = 3$, $z = .5$, birth rates are uniform, and the distribution R is lognormal with parameters given in the text. (Left) The preexisting group has strictness .25; (Center) the preexisting group has strictness .50; (Right) the preexisting group has strictness .75.

For example, consider now a scenario where $M = 3$. Then there are four different σ orderings of the groups that can occur: $\sigma_0 = \{0, 1, 2\}$, $\sigma_1 = \{1, 0, 2\}$, $\sigma_2 = \{1, 2, 0\}$, and $\sigma_3 = \{2, 1, 0\}$. Given an inheritance level $z \neq 0$, we must keep track of the number of individuals of each ordering within each of the three groups, leading to a 12-dimensional system, which can be reduced by one dimension down to 11 because the total population size is 1. Due to the dynamics of group switching, some of the subpopulations will simply exponentially decay, namely, n_{10}, n_{20}, and n_{21}, leaving us with effectively an eight-dimensional system for the case of only three groups. This unfortunately makes analytical work even for this small number of groups quite difficult. We therefore proceed using numerical simulations, a method that has been used before to study religious markets and the dynamics of religious group growth (Iannaccone and Makowsky 2007; Makowsky 2011; Montgomery 1999).

Consider first the results presented in Figure 6, where we explore the equilibrium sizes of each of three groups as the strictnesses of the two affiliated groups vary, given the dynamics of Equation (17) and an initial condition in which all groups are equally sized. In each figure, we fix the strictness value of one of the groups, which we refer to as the "preexisting" group, and plot the equilibrium group sizes as a function of the strictness of the third group, which we refer to as the "new" group; this choice of terminology will be explained below. In all cases, we have chosen parameter values $s = 3$ and $z = .5$, use uniform birth rates, and use a lognormal distribution for R so that $r - 1/2 \sim \text{Lognormal}(\mu, v^2)$, where $\mu = -1/2$ and $v = 2$. The last assumption reflects the fact that income distributions are typically understood to be lognormal (Liberati 2015), and our R distribution can be interpreted as capturing the value of outside-group activities including work for pay. In the two-group case, these parameters and distribution would allow a single group with strictness up to approximately .83 to survive alongside the unaffiliated group, without the need to increase its birth rate beyond the baseline value. We will refer to this strictness value as the absolute maximal strictness in our discussions below.

Some immediate observations stand out from Figure 6. First, if the strictness of the new group is too high, then it cannot sustain its population and eventually dies out as expected given our earlier analysis of the two-group case. Furthermore, the maximal strictness value that the new group can adopt and still survive is always less than the absolute maximal strictness of .83, again as expected. Perhaps less obvious, though, is the fact that the maximal strictness the new group can adopt is not monotonic in the strictness of the preexisting group. When the preexisting group has rather low strictness, the new group may adopt relatively high strictness values and still survive, and as the strictness of the preexisting group increases toward approximately .5 in this case, the maximal strictness of the new group is reduced. But, as the strictness of the preexisting group rises above .5, the maximal strictness of the new group also rises.

Another observation is that, as the strictness of the preexisting group increases, its maximal possible size at equilibrium decreases, as expected; with higher strictness fewer people rank the group highly in their group ordering, and it is less probable for those who do to join the group given the probabilities α. But, more interestingly, the minimal possible size—overstrictnesses of the new group—of the preexisting group is not monotonic in the preexisting group's strictness. When the preexisting group's strictness is very low, a new group with only a slightly higher strictness value will steal most of the members from the preexisting group, making the lowest size for the preexisting group quite small. Similarly, if the preexisting strictness is quite high, any sufficiently low strictness for the new group will completely eliminate the preexisting group. On the other hand, when the strictness of the preexisting group is more moderate—say near .5 in this case—its minimal size is still a relatively large fraction of the overall population.

These two observations become quite important when we imagine groups choosing their strictness levels in a strategic way. Consider a scenario in which only a single, preexisting group exists alongside the unaffiliated group. We might imagine that this group is free to choose whatever strictness level it would like for itself, but should do so in a way that will optimize some objective function. Suppose that the group's main concern is that it have a high membership. Then, if this preexisting group were to ignore the possibility of any new groups forming or breaking away from it, it ought to choose an arbitrarily low strictness level, and thereby recruit almost everyone. However, this choice would leave the preexisting group very vulnerable should a third group form, since, as observed above, the new group could easily steal away almost all of the preexisting group's members by choosing its own strictness carefully. To guard against this, then, the preexisting group should instead choose a somewhat moderate strictness value, such that a new group entering would (a) have fewer possible strictness values to choose from in order to survive and (b) have a minimized possible impact on the size of the preexisting group.

Note that this finding is similar to those found in prior studies (McBride 2008, 2010) but with the added twist that the new group must avoid being too strict to prevent eventually dying out due to loss of members. We thus see strong market pressures toward religious groups that are not so strict, and that can only be countered with sufficiently high birth rates in the strictest groups. This logic can help explain why some of the most successful religious groups that have formed in the last two centuries in the United States, such as the Mormons, Seventh-Day Adventists, and Jehovah's Witnesses, have succeeded while being relatively strict but not ultra strict. These relatively strict groups demand high levels of conformity from their members but do not separate themselves from the rest of mainstream society to the same degree as the ultra-strict groups. Because they are not completely separated, they can engage in active missionary efforts to encourage switching into the group to supplement their relatively high (but not extremely high as, say, the Amish) birth rates.

More insights into the behavior of the three-group case can be seen by examining Figure 7. Here, we display the rates at which individuals transition between the three groups—given by the numbers displayed above the corresponding arrows—and at which they are retained from the births within the group—given by the numbers on the loops starting and ending on the same group—once the system has reached equilibrium. Initial conditions are that every $n_{g\sigma}$ has an equal size, and the equilibrium group sizes are $N_0 \approx 0.344$, $N_1 \approx 0.429$, and $N_2 \approx 0.228$. In this case, we have employed the same lognormal $R(r)$ distribution used to construct Figure 6, have chosen $s = 5$ and $z = .5$ with constant birth rates, and chosen strictness levels for the two affiliated groups such that the fraction of people who rank each group at the top of their ordering is equal for all three groups. Because of this, no group in this case has an inherent advantage merely due to the number of people who might prefer that group above all others, which causes the resulting dynamics to be more dominated by the probabilities of switching directly.

Figure 7
Scaled flow rates of newly born individuals between groups or remaining within a given group at equilibrium

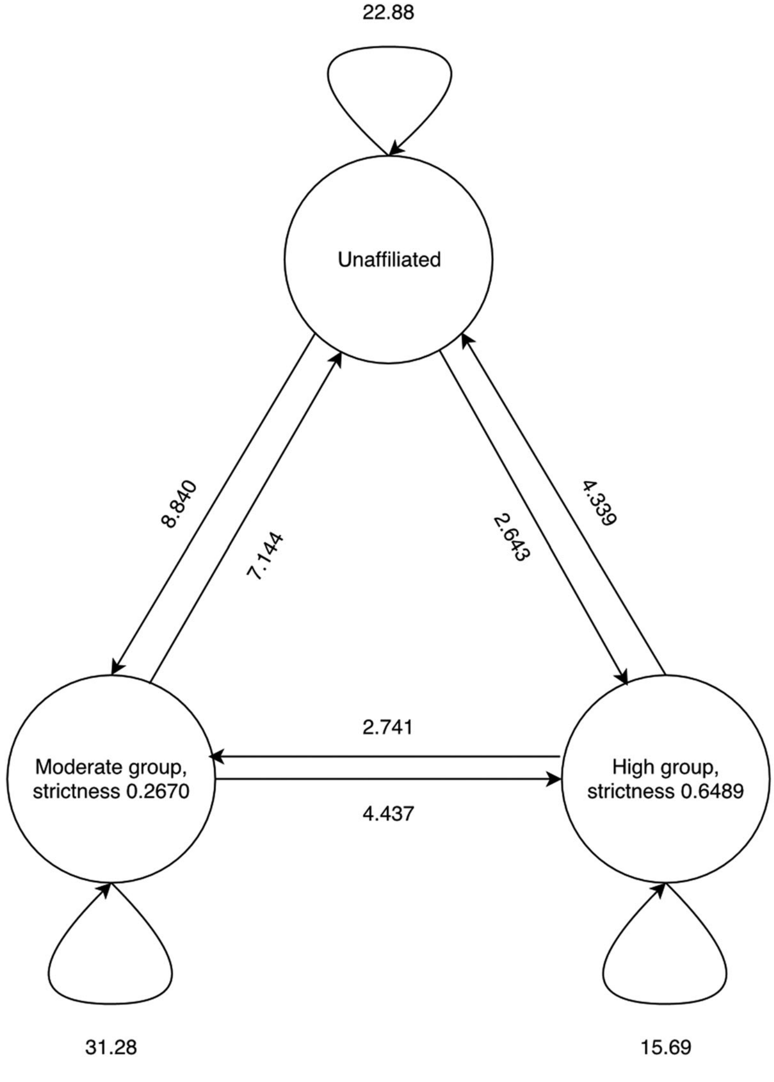

Notes: See text for parameters used in this simulation.

Figure 7 reveals an interesting behavior not seen in prior models but that has been found empirically. Note that there is a larger flow from the high strictness group to unaffiliated than from high strictness to moderate strictness. Prior models based purely on ideal strictness levels (McBride 2008, 2010; Stark and Finke 2000) would generally predict the opposite, as individuals from the highest strictness group would tend to choose the next lowest strictness group when switching groups instead of choosing to not affiliate with any group. Yet, it has been found that religious switchers often became unaffiliated or switch to a group with very different characteristics (Hungerman 2013, 2014). As Hungerman points out, his finding is puzzling in part because the existing models predicted that a religious switcher should move to a relatively similar group rather than become unaffiliated or switch to a very different group.

Figure 8
Normalized group sizes as the birth rate of the strictest group varies
[Color figure can be viewed at wileyonlinelibrary.com]

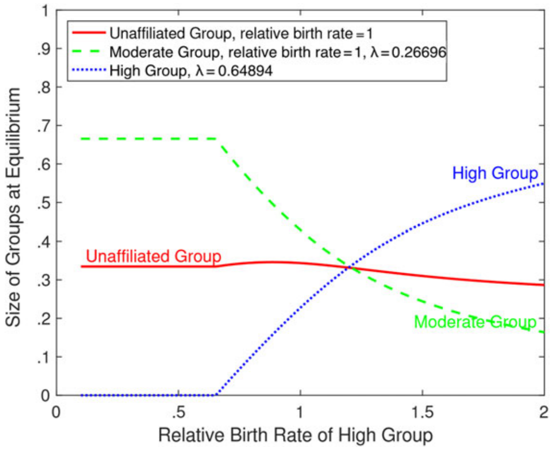

Notes: Parameters are $s = 5$, $z = .5$, R lognormal so that $r - 1/2 \sim \text{Lognormal}(-1/2, 4)$ and strictnesses chosen so that all groups have an equal fraction of the population that most prefers that group.

In our model, however, the transition probabilities that drive switching from one group to another depend on several factors related to cultural transmission. It is much more likely for a member of the high-strictness group to transition to the unaffiliated group than the moderate group, all else being equal, because the individuals who are dissatisfied in the high-strictness group do not get enough exposure to the moderate group to make switching to that moderate group likely. The dissatisfied individuals become unaffiliated because that option is the only alternative that is assumed to not require prior exposure. The empirical result documented by Hungerman can thus potentially be explained by the multifaceted dynamics of cultural transmission. Parents and religious groups are the prime drivers of vertical transmission, but opportunities for exposure to other religious groups drive horizontal transmission. Some stay in the group, others leave after they find something better, and others become unaffiliated when dissatisfied but not yet exposed to something better. The larger lesson can be thus stated: it is not just the existence of religious substitutes but the opportunity to be exposed to them that matter for religious switching.

Another matter of interest is how variation in birth rates affects the market outcomes. Stricter groups tend to have higher fertility than less-strict groups, and these differences in birth rates should contribute to differences in growth rates across groups (see Scheitle, Kane, and Hook 2011). To explore this possibility in the three-group case, we examine the dynamics of our model when the unaffiliated and moderate groups have their birth rates fixed at 1 but the strictest group's birth rate is allowed to vary from less than 1 to above 1. Having the strict group with $b > 1$ matches the real world, where stricter groups have higher fertility, while the $b < 1$ case is for reference.

Figure 8 shows how a change in the strict group's birth rate b affects the long-run proportion of the population in each of the three groups. Given that any $b \neq 1$ will lead to a nonconstant total population size, the sizes in this plot are all normalized by total population, thus allowing us to see the change in each group's relative share of the market. As the strict group's birth rate increases, its long-run share increases, as is expected. The higher birth rate implies a larger number of retained youth, who in turn will have more children that lead to even more retained youth, and so on. Yet, observe that the increase in the strict group's share comes primarily at the expense of the moderate group, whose share decreases dramatically as the strict group's birth rate increases. At the same time, the unaffiliated group's share is only slightly affected, due to

Figure 9
Equilibrium sizes for a system with eight groups, as parameters s or z varies
[Color figure can be viewed at wileyonlinelibrary.com]

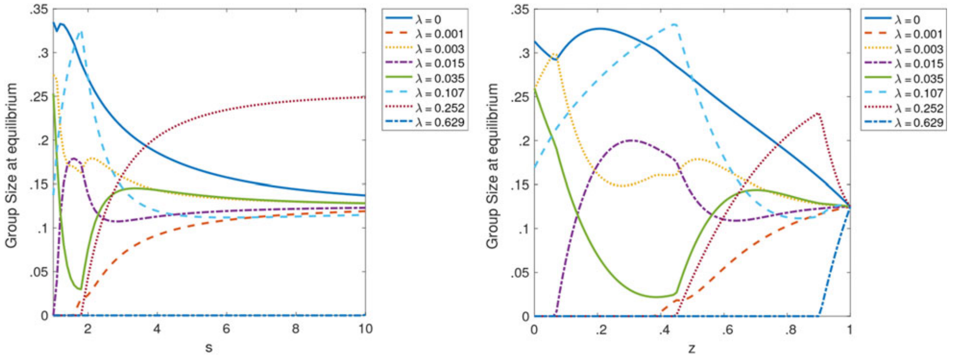

Notes: When not varying, $s = 2$ and $z = 0.5$. The distribution $R(r)$ is lognormal, and strictnesses are chosen such that every group is ranked most highly by an equal fraction of the population. Initial conditions set each subpopulation to an equal size

the underlying switching patterns similar to those in Figure 7. The larger switching from strict to unaffiliated than from strict to moderate means that the strict group's high birth rate does not lead to much of a decline in the share of unaffiliated persons even though the share of moderate group members drops significantly. This finding lends a new potential insight into our understanding of a recent trend in American religion where strict groups are growing faster than less-strict groups, but the unaffiliated share is growing, too. The high fertility of stricter groups, combined with the dynamics of religious inheritance and switching, may actually contribute to the decline of less-strict religious groups (such as the mainline Protestant groups) relative to the unaffiliated group.

Many Groups

As the number of groups continues to increase past three, the system becomes ever more complex, and even simple numerical experiments become unwieldy as there are too many parameters to vary. Nonetheless, we do provide here, as an example, some simulated results in the case of eight groups. Unlike in the three-group case, we set each group's strictness level and hold it fixed for the duration of the system's evolution as we vary s and z.

This exercise requires some decisions to be made about the exact distribution of $R(r)$ and the strictness levels of the group. As before, we use a lognormal $R(r)$ for the distribution of the opportunity cost of time outside the group. For the groups' strictnesses, we initialize the system in a way intended to give each group what might be considered an *a priori* fair chance to survive in the long run. Specifically, the strictnesses are chosen so that each of the eight groups has an equal fraction of the population that ranks that group most highly, and we initialize the system so that each group begins with those individuals in its group who most prefer being in that group. This initialization does not imply that group sizes will be equal in equilibrium because switching rates in and out of groups will not be uniform due to variation in the amounts of time individuals devote to their groups and due to the other dynamics of the system. It will, however, allow us to monitor how the system evolves from what might be considered, in spirit, an equalized starting position.

Figure 9 plots the equilibrium sizes of the eight groups as functions of s (with z fixed at .5) and z (with s fixed at 2). As is clear from the figure, the relative sizes of the groups at equilibrium vary significantly with s and z. Some interesting patterns are evident. For example, some of the low-strictness groups are among the smallest in size, despite the fact that they are equally as preferred as other groups and generally more probable to join. Specifically, the group with strictness .001 has a smaller population than the group with strictness .003 under all parameters tested. Similarly, the .015 strictness group has a smaller size than the .035 group under many parameters combinations. This may be related to the phenomenon observed in the three-group case, whereby a group could steal away many members from a low-strictness group by having a slightly higher strictness value. At the higher end of the strictness scale, we find that the general trend is that the highest-strictness group does quite poorly, while the next two highest groups can do very well. The origin of this effect is a bit clearer. For smaller s or z values, the .107 strictness group evidently picks up all of the individuals who prefer it or the two highest groups; in these parameter regimes, the lower α values for switching to the highest two groups cause them to die out. But, as s or z is increased, the probability for joining the strictness .252 group increases enough so that it now is able to recruit many of those who prefer it, as well as those who prefer the highest-strictness group, all at the expense of the strictness .107 group, whose size drops accordingly.

Conclusion

In this article, we have constructed a dynamic model for the sizes of religious groups, based on a unidimensional categorization of groups by their strictness level, interpreted as the amount of time they expect their members to spend within the group contributing to the common good. This model is similar to previous such models in the way it accounts for how the strictness of a group interacts with the preferences of the members of the overall population, who are effectively described by some distribution over strictness preferences based on the individual's utility function for out of group activities. Based on an individual's rate of utility for these out-group activities r, all existing religious groups can be ranked from highest to lowest utility, based on the group strictness levels. But, our model adds to the existing literature by including a probabilistic component to group switching, such that an individual may not necessarily be able to switch into her most preferred group and have to settle for one of lesser utility. Crucially, the probability of an individual being able to join a group is directly related to the probability of having encountered members of that group during time when both the individual and the group members were engaged in out of group time. Hence, it is more probable to join larger groups, as one is more likely to have encountered their members by sheer number, and to join lower-strictness groups, as those individuals spend more time out of group during which they might be encountered. At the same time, members of high-strictness groups may find it difficult to switch to another group, as they will have spent little time out of group themselves. All of these effects, including possible inheritance from parents to offspring of religious preferences and possibly varying birth rates of the various groups, are summarized by a system of ordinary differential equations for the various population sizes in time.

Analysis of our model has confirmed several phenomenon seen in prior models. For example, we have shown that when the only options are a single group with some finite strictness and another "group" with zero strictness (capturing the ability of people to be unaffiliated with any group), the size of the affiliated group decreases with strictness, such that the group may not be able to survive at equilibrium if its strictness is too high. This effect is not merely due to a reduced fraction of the population that would thrive with such high strictnesses, and is intimately tied to the decreased probability of individuals joining such a high-strictness group. High rates of inheritance can mitigate this effect, as can group switching probabilities that require fewer encounters with members before one can readily join a group. Importantly, we have shown that a group of any strictness can survive if the birth rate of its members is high enough in relation to the birth rate of nonmembers.

Going beyond the two-group case, we also examined in some detail the case of three groups: two affiliations and one unaffiliated group. This scenario is especially relevant when considering the effects of new groups entering the religious market, and our numerical experiments here seem to indicate that moderate-strictness groups are the most robust in terms of how they handle new groups entering the marketplace, whereas very low or very high strictness groups may be effectively driven out of the market in such situations unless they maintain very high birth rates. We plan to consider this specific issue in much greater detail in future work, examining how groups might arrange themselves with regard to strictness to best maximize the goals of the group, be they simply gaining the largest following possible, maximizing the utility of their members, or something else. A related future problem is modeling how religious groups may splinter into offshoots with differing strictness values. Given that our model explicitly tracks the number of each strictness preference subpopulation within each group, it may be plausible to construct a splintering mechanism whereby when a critical number of members of a group have a strictness preference that is substantially different from the group's current strictness value, they splinter off to form a new group.

Finally, we briefly examined from a numerical viewpoint a scenario with several (eight) groups. Our results highlight the inherent complexity of the system, given that the eventual equilibrium varies significantly with parameters, and general trends are somewhat difficult to discern. Further exploration of a setting with several groups will be of interest to social scientists trying to understand the rich dynamics of religious markets, including the forces that drive some groups to thrive and others to die out. Future work using our framework could also explore larger trends in religious markets, such as the recent rise of nonaffiliated persons in the United States. Indeed, a key feature of our model is the explicit characterization of both the demand and supply sides of the religious market, both of which are relevant to understand trends in religious market outcomes. This work may also be of interest to a more general mathematical audience, who might find in it a rich source of interesting mathematical problems.

References

Bisin, Alberto and Thierry Verdier. 2000. "Beyond the melting pot": Cultural transmission, marriage, and the evolution of ethnic and religious traits. *Quarterly Journal of Economics* 115(3):955–88.

Brekke, Torkel. 2016. *Faithonomics: Religion and the free market*. New York: Oxford University Press.

Carvalho, Jean-Paul. 2013. Veiling. *Quarterly Journal of Economics* 128(1):337–70.

Carvalho, Jean-Paul, Sriya Iyer, and Jared Rubin. Forthcoming. Demographic imperatives and religious markets: Considering the individual and interactive roles of fertility and switching in group growth. In *Advances in the economics of religion*. London: Palgrave, Macmillan.

Carvalho, Jean-Paul and Michael Sacks. 2018. *The economics of religious communities: Social integration, discrimination and radicalization*. Working Paper.

Fox, Jonathan and Ephraim Tabory. 2008. Contemporary evidence regarding the impact of state regulation of religion on religious participation and belief. *Sociology of Religion* 69(3):245–71.

Frejka, Tomas and Charles Westoff. 2008. Religion, religiousness and fertility in the US and in Europe. *European Journal of Population/Revue européenne de Démographie* 24(1):5–31.

Gaus, Gerald. 2007. *On philosophy, politics, and economics*. Belmont, CA: Nelson Education.

Hout, Michael, Andrew Greeley, and Melissa Wilde. 2001. The demographic imperative in religious change in the United States. *American Journal of Sociology* 107(2):468–500.

Hungerman, Daniel M. 2013. Substitution and stigma: Evidence on religious markets from the Catholic sex abuse scandal. *American Economic Journal: Economic Policy* 5(3):227–53.

———. 2014. Do religious proscriptions matter? Evidence from a theory-based test. *Journal of Human Resources* 49(4):1053–93.

Iannaccone, Laurence R. 1991. The consequences of religious market structure: Adam Smith and the economics of religion. *Rationality and Society* 3(2):156–77.

———. 1992. Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives. *Journal of Political Economy* 100(2):271–91.

———. 1994. Why strict churches are strong. *American Journal of Sociology* 99(5):1180–1211.

Iannaccone, Laurence R. and Michael D. Makowsky. 2007. Accidental atheists? Agent-based explanations for the persistence of religious regionalism. *Journal for the Scientific Study of Religion* 46(1):1–16.
Kraybill, Donald B., Steven M. Nolt, and David L. Weaver-Zercher. 2012. *The Amish way: Patient faith in a perilous world*. Hoboken, NJ: John Wiley & Sons.
Liberati, Paolo. 2015. The world distribution of income and its inequality, 1970–2009. *Review of Income and Wealth* 61(2): 248–73.
Makowsky, Michael D. 2011. A theory of liberal churches. *Mathematical Social Sciences* 61(1):41–51.
McBride, Michael 2007. Club Mormon: Free-riders, monitoring, and exclusion in the LDS church. *Rationality and Society* 19(4):395–424.
———. 2008. Religious pluralism and religious participation: A game theoretic analysis. *American Journal of Sociology* 114(1):77–106.
———. 2010. Religious market competition in a richer world. *Economica* 77(305):148–71.
———. 2015. Why churches need free-riders: Religious capital formation and religious group survival. *Journal of Behavioral and Experimental Economics* 58:77–87.
Montgomery, James D. 1999. Adverse selection and employment cycles. *Journal of Labor Economics* 17(2):281–97.
Pew Research Center. 2017a. *The changing global religious landscape*. Available at http://www.pewforum.org/2017/04/05/the-changing-global-religious-landscape/.
———. 2017b. *Many countries favor specific religions, officially or unofficially*. Available at http://www.pewforum.org/2017/10/03/many-countries-favor-specific-religions-officially-or-unofficially/.
———. 2018. *The age gap in religion around the world*. Available at http://www.pewforum.org/2018/06/13/the-age-gap-in-religion-around-the-world/.
Putnam, Robert D. and David E. Campbell. 2012. *American grace: How religion divides and unites us*. Riverside: Simon and Schuster.
Scheitle, Christopher P., Jennifer B. Kane, and Jennifer Van Hook. 2011. Demographic imperatives and religious markets: Considering the individual and interactive roles of fertility and switching in group growth. *Journal for the Scientific Study of Religion* 50(3):470–82.
Stark, Rodney and Roger Finke. 2000. *Acts of faith: Explaining the human side of religion*. Berkeley, CA: University of California Press.
The Church of Jesus Christ of Latter-Day Saints. 2018. 2017 statistical report for 2018 April conference. *Ensign* 48(5):119.
Voas, David and Mark Chaves. 2016. Is the United States a counterexample to the secularization thesis? *American Journal of Sociology* 121(5):1517–56.

APPENDIX

Proof of Theorem 1

Consider a set of nonnegative integer values N_1, N_a, N_λ, N_b, and N_0 such that the first N_1 members of $\vec{r}$ choose $t_i = 1$, the next N_a choose $t_i = t_a = a + (1-a)\lambda$ with $0 \leq a \leq 1$, the next N_λ choose $t_i = \lambda$, the next N_b choose $t_i = t_b = b\lambda$ with $0 \leq b \leq 1$, and the final N_0 choose $t_i = 0$. For any given individual, $T_i = T - t_i$, where $T = \sum_j t_j$ for the state we are examining. Then, for all those individuals choosing $t_i = 1$, $T_i = T - 1$ and the lower bound in Equation (6) becomes

$$\frac{1}{2\sqrt{NT}} \equiv R_1.$$

Hence, so long as $r_i < R_1$, for $i \leq N_1$, all of the N_1 individuals will be making their optimal choice and not want to unilaterally switch. Similarly, for the individuals choosing $t_i = \lambda$, $T_i = T - \lambda$, the upper bound of Equation (6) becomes R_1 while the lower bound of Equation (7) becomes $1/2\sqrt{NT} + \beta = R_1 + \beta$. So, as long as $R_1 < r_i < R_1 + \beta$ for all $N_1 + N_a < i \leq N_1 + N_a + N_\lambda$, all of the N_λ individuals will be making their optimal choice and not want to unilaterally switch. In this same way, so long as $r_i > R_1 + \beta$ for all $i > N_1 + N_a + N_\lambda + N_b$ all the N_0 individuals will also be making their optimal choice. For those choosing $t_i = t_a$, $T_i = T - t_a$, so that Equation (4) will be satisfied regardless of t_a so long as $r_i = 1/2\sqrt{NT} = R_1$, which must be the case for these individuals, so they are also playing their optimal choice. Finally, for those choosing $t_i = t_b$, Equation (5) will be satisfied regardless of t_b so long as

$r_i = 1/2\sqrt{NT} + \beta = R_1 + \beta$, which must be the case for these individuals, so that they are also playing their optimal choice. Since all r_i values are accounted for and no individual can increase utility by unilaterally switching strategy, this is a Nash equilibrium. Note, though, that as of yet there is no guarantee that this Nash equilibrium will actually exist for a given $\vec{r}$ and set of parameters λ and β, since our proposed behavioral threshold R_1 is defined in terms of T, but T is in turn determined by the actual optimal strategies t_i adopted by the individuals, which are based off of R_1.

We now show that such an equilibrium always exists. First, let $F(r)$ be the number of individuals with $r_i \le r$, and let $P(r)$ be the number of individuals with $r_i = r$. Then the total amount of time spent in group activities at the above Nash equilibrium is

$$\begin{aligned} T &= [F(R_1) - P(R_1)] + [a + (1-a)\lambda]\, P(R_1) \\ &\quad + \lambda\, [F(R_1+\beta) - F(R_1) - P(R_1+\beta)] + b\lambda P(R_1+\beta) \\ &= (1-\lambda)\,[F(R_1) - (1-a)P(R_1)] + \lambda\,[F(R_1+\beta) - (1-b)P(R_1+\beta)]\,. \end{aligned} \tag{A1}$$

At the same time, we know from the definition of R_1 above that T and R_1 must be related via

$$T = \frac{1}{4NR_1^2} \tag{A2}$$

for the system to be at a Nash equilibrium. Then, so long as an R_1 (and potentially corresponding values for a and/or b) exists that satisfies both Equations (A1) and (A2), the Nash equilibrium above exists. But this can always be made the case: Equation (A2) is a monotonically decreasing, continuous function taking on all positive values as R_1 ranges from 0^+ to ∞, while Equation (A1) is a nondecreasing function that can be made to take on any value between its minimum of $\lambda[F(\beta) - P(\beta)]$ to its maximum of N by adjusting a and/or b as needed as R_1 ranges from 0 to ∞. Hence, the two curves can be made to intersect, and this intersection point is unique with regard to R_1 and therefore T, so the Nash equilibrium exists and the various N_x values are all unique. If specific values of a and/or b must be chosen so that the two curves intersect, then upon doing so both Equations (4) and (5) will automatically be satisfied, as by moving T_i to the left-hand sides of these equations and substituting the r_i value of the individuals playing t_a or t_b, both equations simply become $T = 1/4NR_1^2$, which will be true. Note, though, that if the Nash equilibrium includes both those playing t_a and those playing t_b, the values of a and b that can be used to satisfy these equations may not be unique, though the R_1 and T values they give rise to will be.

Proof of Theorem 2

To prove our claim, we first note that $g(0) = 0$ and $g(f) = f(1-K)(p(f) - 1) < 0$ since $p(f) < 1$ when $f < 1$. We will then need to take the first- and second-order derivatives of $g(n)$:

$$g'(n) = (f - Kn)p'(n) - Kp(n) + K - 1, \tag{A3}$$

$$g''(n) = (f - Kn)p''(n) - 2Kp'(n), \tag{A4}$$

where

$$p'(n) = s(1-\lambda)\frac{(1-n)^{s-1}}{(1-\lambda n)^{s+1}}, \tag{A5}$$

$$p''(n) = s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}}\left[-(s-1)(1-\lambda)+2\lambda(1-n)\right]. \tag{A6}$$

Therefore,

$$\begin{aligned} g''(n) &= s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}}((f-Kn)[-(s-1)(1-\lambda)+2\lambda(1-n)] \\ &\quad -2K(1-n)(1-\lambda n)) \\ &= s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}}[(2K+(s-1)(1-\lambda)K-2\lambda f)n \\ &\quad -(2K+(s-1)(1-\lambda)f-2\lambda f)]. \end{aligned} \tag{A7}$$

Note that $s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}} > 0$ on [0, 1). Let us consider the function

$$h(n) = (2K+(s-1)(1-\lambda)K-2\lambda f)n-(2K+(s-1)(1-\lambda)f-2\lambda f), \tag{A8}$$

which is a linear function of n. The slope of h can be rewritten as

$$(1+\lambda)K + s(1-\lambda)K - 2\lambda f \geq (1+\lambda)K - 2\lambda f > 0, \tag{A9}$$

since $0 < \lambda < 1$, $s > 0$, and $0 < f \leq K \leq 1$. Similarly, the negative intercept of h can be rewritten as

$$2K + sf(1-\lambda) - (1+\lambda)f \geq 2K - (1+\lambda)f > 0. \tag{A10}$$

So $h(n)$ is an increasing function that attains 0 at

$$n^* = \frac{2K+(s-1)(1-\lambda)f-2\lambda f}{2K+(s-1)(1-\lambda)K-2\lambda f}. \tag{A11}$$

If $s \leq 1$, then since $K \geq f$, $n^* \geq 1$. So in this case $g''(n) < 0$ on [0, 1) so that $g'(n)$ is strictly decreasing on [0, 1). So, if $g'(0) > 0$ there exists one nontrivial zero point n_0 of $g(n)$ on [0, 1) with $n_0 < f$ and $g'(n_0) < 0$; otherwise we only have a trivial zero point of $g(n)$ at $n = 0$.

If $s > 1$, then $n^* < 1$, so $g'(n)$ is decreasing on $[0, n^*)$ and increasing on $(n^*, 1)$. We notice that

$$-h(f) = 2K+(s-1)(1-\lambda)f-2\lambda f - f(2K+(s-1)(1-\lambda)K-2\lambda f) \tag{A12}$$

$$= (2K-2\lambda f)(1-f)+(s-1)(1-\lambda)f(1-K) \tag{A13}$$

$$> (2K-2\lambda f)(1-f) \tag{A14}$$

$$> 0. \tag{A15}$$

Recall that $g''(n)$ is positively proportional to $h(n)$, so since $h(f) < 0$, $g''(f) < 0$, so $f < n^*$ since $g''(n) < 0$ only for values less than n^*. Then all the arguments for the case above with $s \leq 1$ still hold on the region $[0, f]$. Namely, if $g'(0) > 0$ there will be one nontrivial zero point n_0 of g on $[0, f)$ and $g'(n_0) < 0$, whereas if $g'(0) < 0$ there are no zero points of g on $[0, f)$ except for the trivial one at $n = 0$; in both cases $g'(f) < 0$ and therefore $g'(n^*) < 0$. Furthermore, since

$g'(1) = -1 < 0$, g' must be negative on all of $[f, 1]$, so there are no zero points of g on $[f, 1]$, and n_0 is the only possible nontrivial zero point of g.

Proof of Theorem 3

Note $g_b(0) = 0$, while $g_b(1) = b(K - 1) < 0$. Then if $g_b'(0) > 0$, g_b must have at least one root on the interval (0,1). The derivative

$$g_b'(0) = fs(1 - \lambda) + bK - 1.$$

So, if $b > b_{min} \equiv [1 - fs(1 - \lambda)]/K$, the stricter group can survive with a finite fraction of the population at equilibrium.