DYNAMICS OF RELIGIOUS GROUP GROWTH AND SURVIVAL

A Dissertation
Presented to
The Academic Faculty

By

Tongzhou Chen

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Mathematics

Georgia Institute of Technology

August 2018

DYNAMICS OF RELIGIOUS GROUP GROWTH AND SURVIVAL

Approved by:

Dr. Martin Short, Advisor
School of Mathematics
Georgia Institute of Technology

Dr. Michael McBride
School of Social Sciences
University of California, Irvine

Dr. Sung Ha Kang
School of Mathematics
Georgia Institute of Technology

Dr. Wenjing Liao
School of Mathematics
Georgia Institute of Technology

Dr. Yingjie Liu
School of Mathematatcs
Georgia Institute of Technology

Date Approved: June 29, 2018

The way was long, and wrapped in gloom did seem,

As I urged on to seek my vanished dream.

Qu Yuan, translated by Yang Hsien-yi and Gladys Yang

PARENTIBVS MEIS

ACKNOWLEDGEMENTS

First and foremost, I would like to express my sincere gratitude to my Ph.D. advisor Dr. Martin Short. I am very lucky to have him as my supervisor. In the past five years, he leaded me into the world of beautiful applications of mathematical models in human society. His inspiring guidance helped me overcome the obstacles and difficulties in my Ph.D. research. Nothing can be achieved without his continuous support and encouragement during my Ph.D. study.

Many thanks go to our collaborator Dr. Michael McBride as he could always provide creative ideas and show big pictures of the field. I am also grateful to Dr. Yingjie Liu, Dr. Ronghua Pan, Dr. Molei Tao and Dr. Haomin Zhou, for all wise and beneficial advice in our discussions.

I would like to appreciate Dr. Martin Short, Dr. Michael McBride, Dr. Sung Ha Kang, Dr. Wenjing Liao and Dr. Yingjie Liu for accepting to serve on my dissertation committee. As a board member of the SIAM student chapter for three years, I thank Dr. Luca Dieci for the help and suggestions on holding seminars and conferences. I also want to thank Dr. Molei Tao for his idea of organizing the graduate student brainstorm seminars which give graduate students opportunities to communicate and discuss our researches. The English classes offered by Ms. Cathy Jacobson and Dr. Morag Burke, as well as the training and advice on teaching given by Ms. Klara Grodzinsky, Dr. Xu-Yan Chen and Dr. Gregory Mayor are very useful and impressive, which I will never forget.

I acknowledge all the professors that I have had at Georgia Tech, Paris-Sud and USTC. Without them, I am not able to discover the beauty and magic of mathematics and make up my mind to pursue a Ph.D. in applied mathematics.

I am fortunate to make so many friends at School of Mathematics who made my Ph.D. life much more colorful and enjoyable, including Renyi Chen, Fabio Difonzo, Rundong Du, Qianli Hu, Chenchen Mou, Ruilin Li, Wuchen Li, Qingqing Liu, Shu Liu, Longmei

Shu, Hagop Tossounian, Ruidong Wang, Xin Wang, Shijie Xie, Xin Xing, Yian Yao, Lei Zhang, Weiwei Zhang and many others. I am not able to list all of them, but I believe they know who they are. I am also very thankful to many other friends for sending their care to me from all over the world. I treasure all the fun moments spent with my friends.

Last but not least, my deepest gratitude goes to my parents, for their unconditional support and understanding since I was born.

The work in this dissertation is partially funded by ARO MURI Grant W911NF-11-1-0332 and AFOSR MURI Grant FA9550-10-1-0569.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

We model and analyze the dynamics of religious group membership and size. A groups 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 clear parameter combination that determines whether the non-zero strictness group can survive over time, which is more difficult at higher strictness levels. At the same time, we show that a higher than average birthrate can allow even the highest strictness groups to survive. We also study the dynamics of several groups, gaining insight into strategic choices of strictness values and displaying the rich behavior of the model. We then move to the simultaneous-move two-group game where groups can set up their strictnesses strategically to optimize the goals of the group. Affiliations are assumed to have three types and each type of group has its own group utility function. Analysis on the utility functions and Nash equilibria presents different behaviors of various types of groups. Finally, we numerically simulated the process of new groups entering the religious marketplace which can be viewed as a sequence of Stackelberg games. Simulation results show how the different types of religious groups distinguish themselves with regard to strictness.

CHAPTER 1
INTRODUCTION AND BACKGROUND

Religion can be highly dynamic. Unless stifled by government regulations that hinder competition, religious groups will come and go as they experience inflows and outflows of membership [1, 2]. As shown in a comprehensive study of American religion, religious market dynamics are seen in the high rates of religious switching and reaffiliation among churchgoers [3]. While some groups win in this competition for adherents, others lose, and the outcome is a vibrant 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 explain this vibrancy, 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 [4, 5]. This insight helps to explain why strict churches have grown faster during the last several decades [6]. 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 [7, 8, 9]. 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. Entrepreneurs supply this diversity, and high religious pluralism results from religious consumers with different tastes making their optimal affiliation decisions.

This dissertation 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 non-religious 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 [10, 11]. The second literature uses dynamic models of cultural transmission to understand the spread of religious practices within and across generations [12, 13]. Our model thus combines key elements of several different strains of analysis, namely, 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.

The mathematical and simulation analysis herein reveals that the dynamics of a religious market with these many features are rich but that several patterns are still emergent. 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 [14], 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 place them near the mean of the underlying distribution for non-religious goods. That moderate groups can survive and thrive has been noted before in explicit dynamic studies [15, 16], but we demonstrate how several other factors not mentioned in those studies can also contribute to the persistent success of moderate groups.

Three prior attempts to explicitly model religious market dynamics are most similar to ours. Montgomery [15] 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 [16] 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. Finally, Scheitle et al [11] 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.

We also study the competitions of groups in the religious market in this dissertation. In prior models, religious groups are often assumed to enter the marketplace at different times so that groups currently in the market must care about potential entrants. Barros and Garoupa [17] treated churches as Stackelberg leaders followed by sects. McBride [8, 9] used a two-stage representation of group entry allowing groups to postpone the decision at stage two instead of entering at stage one. In our work, we first analyze the simultaneous-move two-group games to gain some understanding of the behavior of different types of groups at the Nash equilibria. Then we adopt the idea that current groups have to be concerned about the future potential followers and simulate how groups sequentially enter the religious market. This process can be viewed as a sequence of Stackelberg games. Simulation results reveal how religious groups with different goals arrange themselves with regard to strictness.

CHAPTER 2
SINGLE GROUP MODEL

2.1 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. 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 amongst the group members evenly, regardless of their individual contribution.

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

$$U_{in} = c \left(\frac{\sum_j t_j}{N_g} \right)^{1/2}, \tag{2.1}$$

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 (2.1) 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

$$U_{out} = r_i(1 - t_i). \tag{2.2}$$

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 prefer. We will assume throughout that the values of r_i for the various individuals are 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

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

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 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 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 P. 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

$$\begin{aligned} U_i &= U_{in} + U_{out} - P \\ &= \left(\frac{\sum_j t_j}{N_g} \right)^{1/2} + r_i(1 - t_i) - \beta_g(\lambda_g - t_i)_+. \end{aligned} \tag{2.3}$$

2.2 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 his 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 which 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

$$t_a = \frac{1}{4Nr_i^2} - T_i \equiv a_i + (1-a_i)\lambda,\ 0 < a_i < 1, \tag{2.4}$$

$$t_b = \frac{1}{4N(r_i-\beta)^2} - T_i \equiv b_i\lambda,\ 0 < b_i < 1, \tag{2.5}$$

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

$$\frac{1}{2\sqrt{N(1+T_i)}} < r_i < \frac{1}{2\sqrt{N(\lambda+T_i)}}, \tag{2.6}$$

and t_b is only a valid critical point for individuals with

$$\frac{1}{2\sqrt{N(\lambda+T_i)}} + \beta < r_i < \frac{1}{2\sqrt{NT_i}} + \beta. \tag{2.7}$$

For fixed T_i , $U_i(1) = U_i(\lambda)$ at an r_i value within the range of values in (2.6), 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 (2.7), 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 non-decreasing in r_i. For t_a, which is a function of r_i, we have

$$U_i(t_a) = \frac{1}{4Nr_i} + r_i(1+T_i)\ , \tag{2.8}$$

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,

$$U_i(t_b) = \frac{1}{4N(r_i-\beta)} + r_i(1+T_i) - \beta(\lambda+T_i)\ , \tag{2.9}$$

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 non-decreasing function of r_i, and the optimal t_i smoothly 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 2.1. *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.*

Proof. 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 (2.6) 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 $t_i = \lambda$, $T_i = T - \lambda$ and the upper bound of (2.6) becomes R_1 while the lower bound of (2.7) 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 (2.4) will be satisfied regardless of t_a so long as $r_i = 1/2\sqrt{NT} = R_1$, which is the case for these individuals, so they are also playing their optimal choice. Finally, for those $t_i = t_b$, (2.5) will be satisfied regardless of t_b so long as

$r_i = 1/2\sqrt{NT} + \beta = R_1 + \beta$, which is the case for these individuals. Since all r_i values are accounted for and no individual can increase utility by unilaterally switching strategy, this is a Nash Equilibrium.

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

$$\begin{aligned} T = {} & [F(R_1) - P(R_1)] + [a + (1-a)\lambda]\, P(R_1) + \\ & \lambda\, [F(R_1+\beta) - F(R_1) - P(R_1+\beta)] + b\lambda P(R_1+\beta) \\ = {} & (1-\lambda)\,[F(R_1) - (1-a)P(R_1)] + \lambda\,[F(R_1+\beta) - (1-b)P(R_1+\beta)] \ . \end{aligned} \tag{2.10}$$

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

$$T = \frac{1}{4NR_1^2} \ . \tag{2.11}$$

Then, so long as an R_1 (and potentially corresponding values for a and/or b) exists that satisfies both (2.10) and (2.11), the Nash Equilibrium above exists. But this can always be made the case: (2.11) is a monotonically-decreasing, continuous function taking on all positive values as R_1 ranges from 0^+ to ∞, while (2.10) is a non-decreasing 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. $\square$

An illustration of the Nash Equilibrium is shown in Fig. 2.1. Here, ten individuals with r_i values shown as blue X marks on the horizontal axis are members of a group with $\lambda = 0.25$ and $\beta = 0.15$. The solid red curve is (2.11), while the discontinuous black curve is (2.10) with $a \approx 0.62$ and $b = 0.6$ (though the value of b is unimportant in this particular

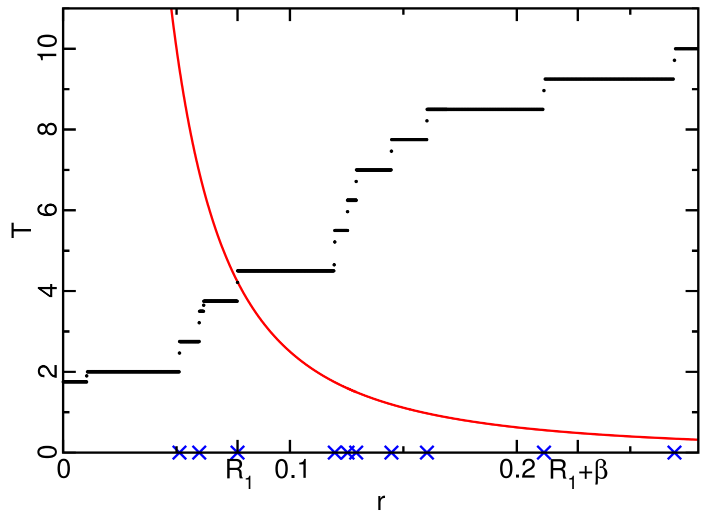

Figure 2.1: An illustration of the existence and determination of the Nash Equilibrium for the single group case. Here, the r_i of the group members are shown as blue X marks on the horizontal axis, the solid red curve is (2.11), and the discontinuous black curve is (2.10).

case). 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 $a \approx 0.62$, the six individuals with $R_1 < r_i < R_1 + \beta$ will choose $t_i = \lambda$, and the one remaining individual will choose $t_i = 0$.

2.3 Ideal Strictness and Punishment Levels

The previous section considers how a variety of individuals 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 the ideal strictness and punishment factor are 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

$$U = \sqrt{\lambda} + r(1-\lambda)\ . \tag{2.12}$$

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 paper we generally assume that all $r_i \geq 1/2$. If the group adopts the ideal strictness level, they 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 (2.11). 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

$$\sqrt{\lambda} + r_i(1 - \lambda) < \sqrt{\lambda N/(N+1)} + r_i - \beta\lambda \,, \tag{2.13}$$

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

$$\sqrt{\lambda N/(N+1)} + r_i - \beta\lambda > r_i \,. \tag{2.14}$$

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 $P = \beta\lambda \approx \sqrt{\lambda}$, 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 (2.10) and (2.11) can only occur at one of three places: at an $R_1 > r_i$ where (2.10) yields $\lambda N + 1$ and $t_i = 1$, at $R_1 = r_i$ where (2.10) yields $\lambda(N+1) + a(1-\lambda)$ and t_i is a linear combination of one and λ, or at an $R_1 < r_i$ where (2.10) 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 impossible 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 one 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.

CHAPTER 3
MULTIPLE GROUP MODEL

We now turn to a more dynamic situation, in which case 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 [11]. We cover each effect separately below, and summarize the model in Fig. 3.1.

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 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

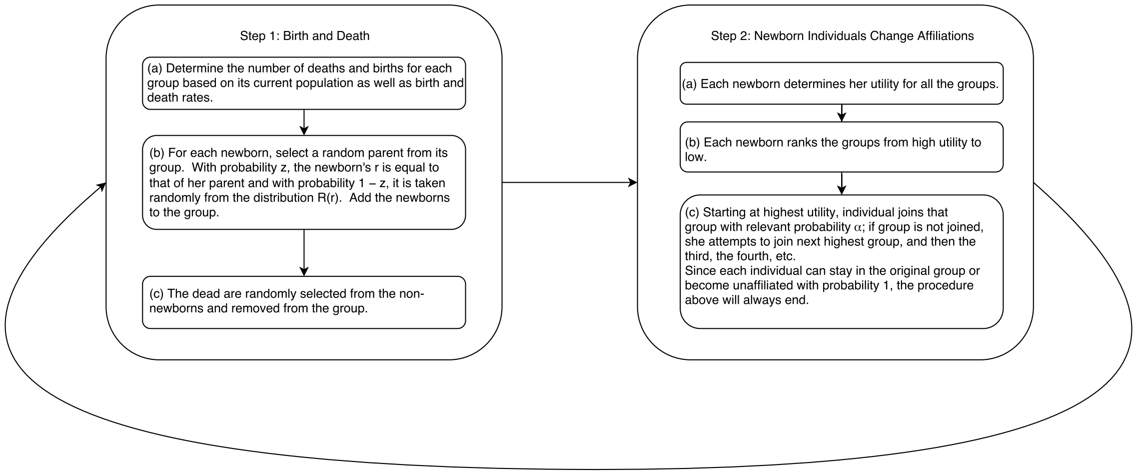

Figure 3.1: A flowchart detailing the various pieces of the multigroup model.

linear combination of one and λ, but this is a very borderline case that should not affect the remainder of the results.

3.1 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 his r_i is equal to that of his parent and with probability $1-z$ his 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 individuals exhibit the same birth rate, this mechanism will cause $R(r)$ to be stationary in time, in expectation.

3.2 Changing Affiliation

We assume that every individual has one chance in their life to change their 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 principal associate with any of the groups g' and thereby obtain utility

$$U_{ig'} = \sqrt{\lambda_{g'}} + r_i(1 - \lambda_{g'}) \,. \tag{3.1}$$

In a system with perfect and complete information, each individual would simply determine which g' provides the maximum utility and choose that group. For M groups with strictnesses $\lambda_0 = 0 < \lambda_1 < \lambda_2 < \cdots < \lambda_{M-1}$, we define cutoff values r_g in the out-group rate distribution such that individuals with r values in (r_{g+1}, r_g) would get their maximal utility within the group with strictness λ_g. By finding the specific r value that would have an equal payoff between the two groups λ_{g+1} and λ_g, we determine that the cutoff between these groups lies at $r_g = 1/(\sqrt{\lambda_{g-1}} + \sqrt{\lambda_g})$ where $g = 1 \ldots M-1$; we define $r_0 = \infty$ and $r_M = 1/2$ for notational convenience. If $R(r)$ is stationary in time and in expectation, then a model in which all players simply choose their optimal group would immediately place all individuals in their ideal group, and the system would remain in that same configuration for all time, in expectation, regardless of birth and death. We refer to this simple model as the "stationary model".

However, the simple "stationary model" neglects an important aspect of switching religious groups, which is the 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'

$$\alpha_{gg'} = 1 - \left[1 - \frac{(1-\lambda_{g'})N_{g'}}{\sum_k (1-\lambda_k)N_k}\right]^{s(1-\lambda_g)} , \tag{3.2}$$

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 (3.2) 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.

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 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 $g' \neq g,\ 0$ 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.

3.3 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

$$\frac{dn_{g\sigma}}{dt} = -n_{g\sigma} + \sum_{g'=0}^{M-1} b_{g'} \left[z n_{g'\sigma} + f_\sigma (1-z) N_{g'} \right] p_{g'g\sigma} \ . \tag{3.3}$$

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

$$p_{g'g\sigma} = \alpha_{g'g} \prod_{j=0}^{J-1} \left(1 - \alpha_{g'\sigma(j)}\right) \ . \tag{3.4}$$

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 (3.3) 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

$$\frac{d\tilde{n}_{g\sigma}}{dt} = \sum_{g'=0}^{M-1} b_{g'} \left(\left[z\tilde{n}_{g'\sigma} + f_\sigma (1-z) \tilde{N}_{g'} \right] p_{g'g\sigma} - \tilde{N}_{g'} \tilde{n}_{g\sigma} \right) \ , \tag{3.5}$$

where the p values are the same as above, and the α values still follow (3.2) but with N_j replaced with $\tilde{N}_j$. In general we will use (3.5) 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.

3.4 Two Group Case

In this section, we present some analytical results for the simplest non-trivial 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 $p_{110} = 0$.

3. People who are originally in group 0 and prefer group 1 can switch to group 1 with probability
$$p_{011} \equiv p = \alpha_{01} = 1 - \left[1 - \frac{(1-\lambda)N_1}{1-\lambda N_1}\right]^s. \tag{3.6}$$

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 without loss of generality. Then (3.5) becomes

$$\begin{cases} \frac{dn_{00}}{dt} &= -n_{00} + z(n_{00} + n_{10}) + (1-f)(1-z) \\ \frac{dn_{01}}{dt} &= -n_{01} + [zn_{01} + f(1-z)N_0)](1-p) \\ \frac{dn_{10}}{dt} &= -n_{10} \\ \frac{dn_{11}}{dt} &= -n_{11} + [zn_{01} + f(1-z)N_0]\,p + [zn_{11} + f(1-z)N_1] \end{cases} \tag{3.7}$$

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

$$\frac{dn}{dt} = (f - Kn)p + Kn - n \equiv g(n) = 0, \tag{3.8}$$

where

$$p(n) = 1 - \left[\frac{1-n}{1-\lambda n}\right]^s. \tag{3.9}$$

In the extreme case $\lambda = 1$, 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 result:

Theorem 3.1. *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 $(0, f)$.*

Proof. 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 $g(n)$:

$$g'(n) =(f - Kn)p'(n) - Kp(n) + K - 1 \tag{3.10}$$

$$g''(n) =(f - Kn)p''(n) - 2Kp'(n) \ , \tag{3.11}$$

where

$$p'(n) =s(1-\lambda)\frac{(1-n)^{s-1}}{(1-\lambda n)^{s+1}} \ , \tag{3.12}$$

$$p''(n) =s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}}\left[-(s-1)(1-\lambda) + 2\lambda(1-n)\right] \ . \tag{3.13}$$

Therefore,

$$\begin{aligned} g''(n) =& s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}}\left((f-Kn)\left[-(s-1)(1-\lambda) + 2\lambda(1-n)\right] - \right. \\ & \left. 2K(1-n)(1-\lambda n)\right) \\ =& s(1-\lambda)\frac{(1-n)^{s-2}}{(1-\lambda n)^{s+2}}[(2K + (s-1)(1-\lambda)K - 2\lambda f)n - \\ & (2K + (s-1)(1-\lambda)f - 2\lambda f)]. \end{aligned} \tag{3.14}$$

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

$$h(n) = (2K + (s-1)(1-\lambda)K - 2\lambda f)n - (2K + (s-1)(1-\lambda)f - 2\lambda f), \tag{3.15}$$

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

$$(1+\lambda)K + s(1-\lambda)K - 2\lambda f \geq (1+\lambda)K - 2\lambda f > 0, \tag{3.16}$$

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

$$2K + sf(1-\lambda) - (1+\lambda)f \geq 2K - (1+\lambda)f > 0. \tag{3.17}$$

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

$$n^* = \frac{2K + (s-1)(1-\lambda)f - 2\lambda f}{2K + (s-1)(1-\lambda)K - 2\lambda f}\,. \tag{3.18}$$

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 non-trivial 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

$$-h(f) = 2K + (s-1)(1-\lambda)f - 2\lambda f - f(2K + (s-1)(1-\lambda)K - 2\lambda f) \tag{3.19}$$

$$= (2K - 2\lambda f)(1-f) + (s-1)(1-\lambda)f(1-K) \tag{3.20}$$

$$> (2K - 2\lambda f)(1-f) \tag{3.21}$$

$$> 0 \tag{3.22}$$

That implies $f < n^*$. So similar to the first case, we have if $g'(0) > 0$ there exists one non-trivial zero point n_0 of $g(n)$ on $[0, f)$ and $g'(n_0) < 0$. Furthermore, since $g'(1) = -1 < 0$, g' remains negative on $(n_0, 1)$, so there are no other zero points of g in the interval $(n_0, 1)$. If $g'(0) < 0$ on the other hand, we only have a trivial zero point of $g(n)$ at $n = 0$.

□

Theorem 3.2. *If $g'(0) = fs(1-\lambda) + K - 1 > 0$, then the non-trivial equilibrium n_0 satisfies $\frac{\partial n_0}{\partial \lambda} < 0$, $\frac{\partial n_0}{\partial z} > 0$, $\frac{\partial n_0}{\partial s} > 0$.*

Proof. Applying implicit function theorem, we have

$$\frac{\partial g}{\partial n}(n_0)\frac{\partial n_0}{\partial \lambda} + (f - Kn_0)\frac{\partial p}{\partial \lambda}(n_0) + p(n_0)f'(\lambda) = 0 \tag{3.23}$$

$$\frac{\partial g}{\partial n}(n_0)\frac{\partial n_0}{\partial z} + (-p(n_0) + fp(n_0) + 1)n_0 = 0 \tag{3.24}$$

$$\frac{\partial g}{\partial n}(n_0)\frac{\partial n_0}{\partial s} + (f - Kn_0)\frac{\partial p}{\partial s}(n_0) = 0 \tag{3.25}$$

By Theorem 3.1, we have $\frac{\partial g}{\partial n}(n_0) < 0$. And $f - Kn_0 > f(1-K) > 0$ since $n_0 < f$. 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, $f'(\lambda) < 0$. Moreover, the probability

$$p = 1 - (\frac{1-n}{1-\lambda n})^s$$

also satisfies $\frac{\partial p}{\partial \lambda}(n_0) < 0$. Therefore, according to (3.23), $\frac{\partial n_0}{\partial \lambda} < 0$.

In (3.24), $(-p(n_0) + fp(n_0) + 1)n_0 > 0$ as $p(n_0) < 1$. This gives $\frac{\partial n_0}{\partial z} > 0$.

We also notice that since $\frac{1-n_0}{1-\lambda n_0} < 1$, p is increasing in s implying $\frac{\partial p}{\partial s}(n_0) > 0$. Therefore (3.25) provides that $\frac{\partial n_0}{\partial s} > 0$. □

Summarizing the results above, the group with strictness $\lambda > 0$ will only survive as a finite fraction of the population at equilibrium if $fs(1-\lambda) + f(1-z) > 1 - z$. Presumably the inheritance rate z and the parameter s are not under the control of any of the groups, but λ, and thereby f, are chosen by each individual group. Note that f is decreasing in λ, so that more strict groups are more apt to have a small population or even die out over time than less strict groups. And, for any given R, z, and s, there is some maximal strictness that the group can adopt and still continue to survive in the long run.

Raising the probability of inheritance z will tend to prevent group 1 from dying out; see Fig. 3.2. 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.

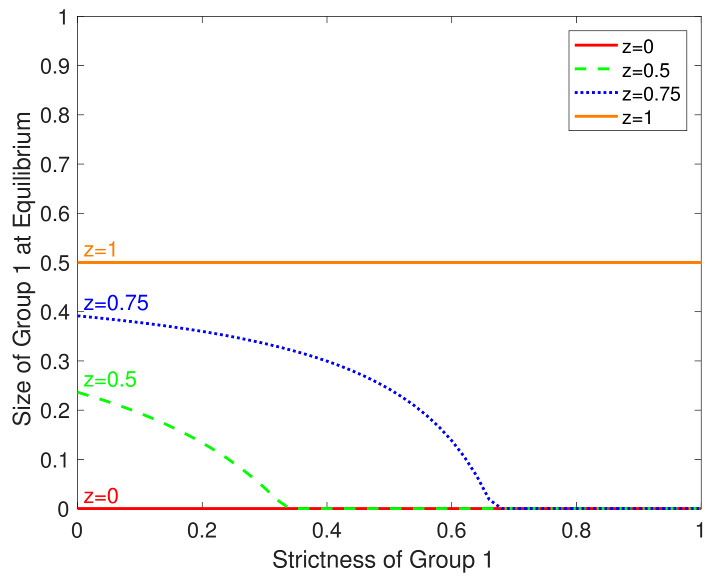

Figure 3.2: The size of group 1 at equilibrium is plotted as a function of its strictness λ for varying values of z, with $s = 0.75$ fixed and the distribution $R(r)$ chosen at each λ such that $f = 0.5$.

Increasing the model parameter s has a similar effect; see Fig. 3.3. 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 ∞, $\alpha_{0,1}$ 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, $\alpha_{0,1}$ 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

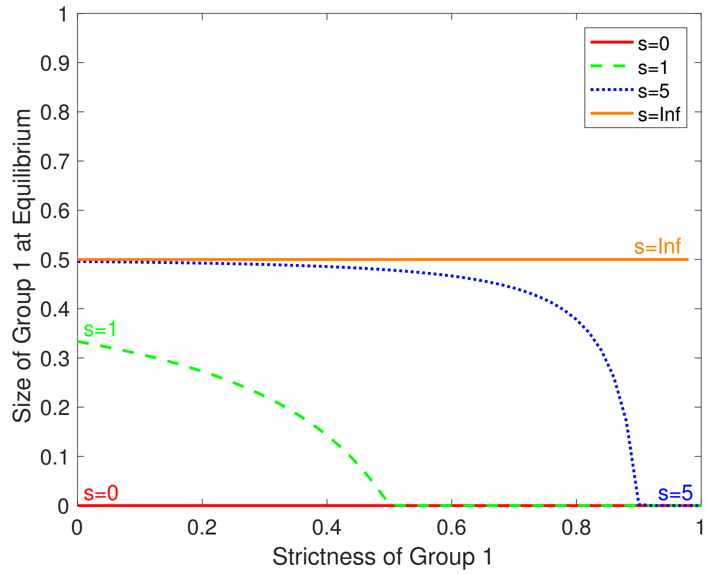

Figure 3.3: The size of group 1 at equilibrium is plotted as a function of its strictness λ for varying values of s, with $z = 0.5$ fixed and the distribution $R(r)$ chosen at each λ such that $f = 0.5$.

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 [11, 18]. A high birth rate in the strict group can counteract the fact that there is relatively low conversion of external individuals to the group, given low exposure 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

$$\begin{cases} \frac{dn_{00}}{dt} &= z(n_{00} + bn_{10}) + (1-f)(1-z)(N_0 + bN_1)(1 - n_{00}) \\ \frac{dn_{01}}{dt} &= [zn_{01} + f(1-z)N_0)]\,(1-p) - (N_0 + bN_1)n_{01} \\ \frac{dn_{10}}{dt} &= -(N_0 + bN_1)n_{10} \\ \frac{dn_{11}}{dt} &= [zn_{01} + f(1-z)N_0]\,p + b\,[zn_{11} + f(1-z)N_1] - (N_0 + bN_1)n_{11} \end{cases} \tag{3.26}$$

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

$$\begin{aligned}
&\left\{-K(b-1)n^2 + n\left[(b-1)(Kz + f(1-z)) - K(1-z)\right] + f(1-z)\right\} p(n) + \\
&\quad \left\{-(b-1)^2 n^2 + n(b-1)(bK + z - 2) + (1-z)(bK-1)\right\} n = g_b(n) = 0 \,. \qquad (3.27)
\end{aligned}$$

Then the following result holds:

Theorem 3.3. *For any $0 < \lambda \leq 1$, there exists a minimal birthrate b_{min} that allows for survival of the stricter group at equilibrium.*

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

The result above highlights the fact that for very strict groups, a higher than average birthrate may be necessary for long term survival, but also guarantees that this is always possible, at least in the case of two groups. Also, b_{min} is increasing in λ, so the stricter a group wishes to be, the greater the birthrate necessary for survival, assuming the group could not survive at $b = 1$. Finally, note that the largest possible value of b_{min}, occurring at $\lambda = 1$, could not be greater than $1/z$, so any group with a birthrate higher than this is guaranteed to survive regardless of their λ. Fig. 3.4 illustrates that if the strict group has a higher birth rate, while all the other parameters are fixed, it can still survive.

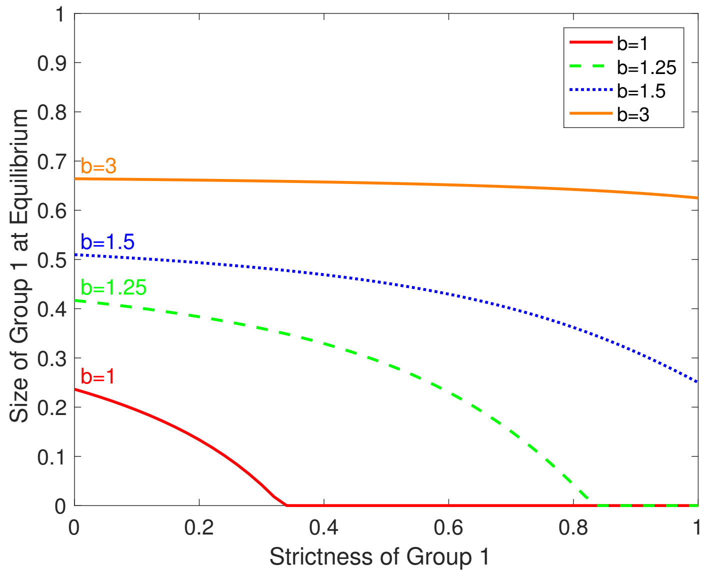

Figure 3.4: The size of group 1 at equilibrium is plotted as a function of its strictness λ for varying values of b, with $s = 0.75$ and $z = 0.5$ fixed and the distribution $R(r)$ chosen at each λ such that $f = 0.5$.

3.5 Three or More 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.

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 twelve-dimensional system, which can be reduced by one dimension down to 11 due to the constraint that 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 [15, 16, 19]. We solve the ode (3.3) numerically with the initial condition such that everybody starts in their favorite group and determine that the equilibrium is reached if the numerical derivative at this point is smaller than a very small threshold.

Consider first the results presented in Fig. 3.5, where we explore the equilibrium sizes of each of three groups as the strictnesses of the two affiliated groups vary, given the dynamics of (3.5) 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 $z = 0.5$, use uniform birthrates, 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 [20], 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 0.83 to survive alongside the unaffiliated group, without the need to increase their birthrate 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 Fig. 3.5. First, if the strictness of the new group is too high, they cannot sustain their population and eventually die out, which is to be expected. Further, the maximal strictness value that the new group can adopt and still survive is always less than the absolute maximal strictness of 0.83, again as expected.

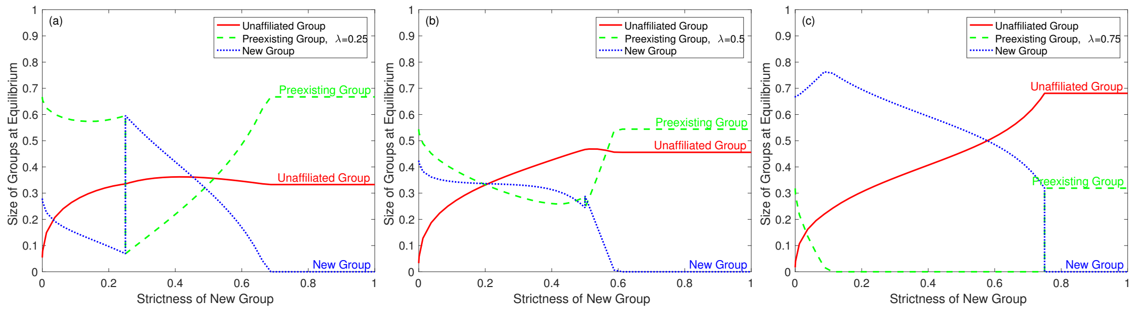

Figure 3.5: Equilibrium sizes of three groups as the strictnesses of the two affiliated groups vary. In all cases, $s = 3$, $z = 0.5$, birthrates are uniform, and the distribution R is lognormal with parameters given in the text. (Left) The preexisting group has strictness 0.25; (Center) the preexisting group has strictness 0.50; (Right) the preexisting group has strictness 0.75.

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 0.5 in this case, the maximal strictness of the new group is reduced. But, as the strictness of the preexisting group rises above 0.5, 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 – over strictnesses 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 0.5 in this case – their 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 groups 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 that those found in prior studies [8, 9] but with the added twist that the new group must avoid being too strict to prevent from eventually dying out due to loss of members by insufficient births into the group. We thus see strong market pressures toward moderate religion that can only be countered with sufficiently high birth rates in the strictest groups.

More insights into the behavior of the three group case can by seen by examining Fig. 3.6. 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 Fig. 3.5, have chosen $s = 5$ and $z = 0.5$ 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

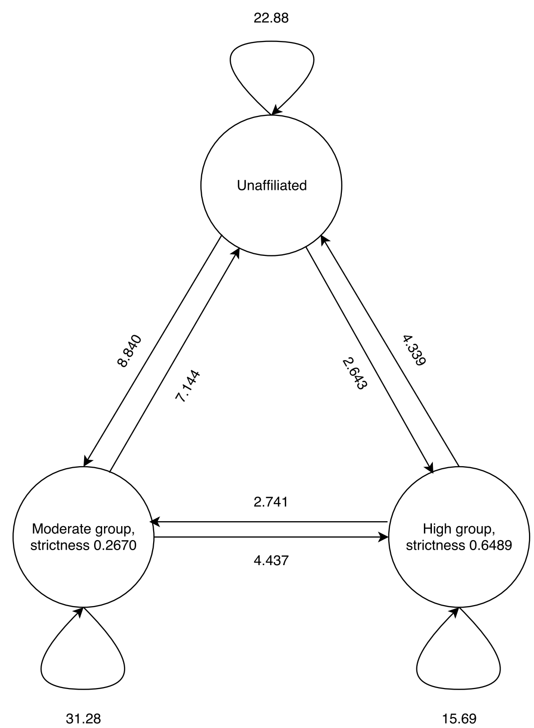

Figure 3.6: Scaled flowrates of newly born individuals between groups or remaining within a given group at equilibrium. See text for parameters used in this simulation.

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 3.6 reveals an interesting behavior not seen in prior models. Note that there is a greater flow from the high strictness group to unaffiliated than from high strictness to moderate strictness. Prior models based purely on ideal strictness levels [7, 8, 9] 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. In our model, however, the transition probabilities greatly affect the model outcome. 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.

Moving beyond three groups, the system becomes ever more complex and even simple numerical experiments become unwieldy as there are too many parameters to vary. But, as an example, we do provide simulated results in the case of eight groups here. For these simulations, a lognormal $R(r)$ was chosen, and the strictnesses of the eight groups are selected such that each of the eight groups has an equal fraction of the population that ranks that group most highly. We use initial conditions such that every subpopulation is equally sized, and plot equilibrium sizes of the eight groups as functions of s (with z fixed at 0.5) and z (with s fixed at 2) in Fig. 3.7. As is clear from Fig. 3.7, the relative sizes of the groups at equilibrium varies 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 0.001 has a smaller population than the group with strictness 0.003 under all parameters tested. Similarly, the 0.015 strictness group has a smaller size than the 0.035 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 0.107 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 0.252 group increases enough so that it now is able to recruit many of those who prefer it, as well as those that prefer the highest strictness group, all at the expense of the strictness 0.107 group, whose size drops accordingly.

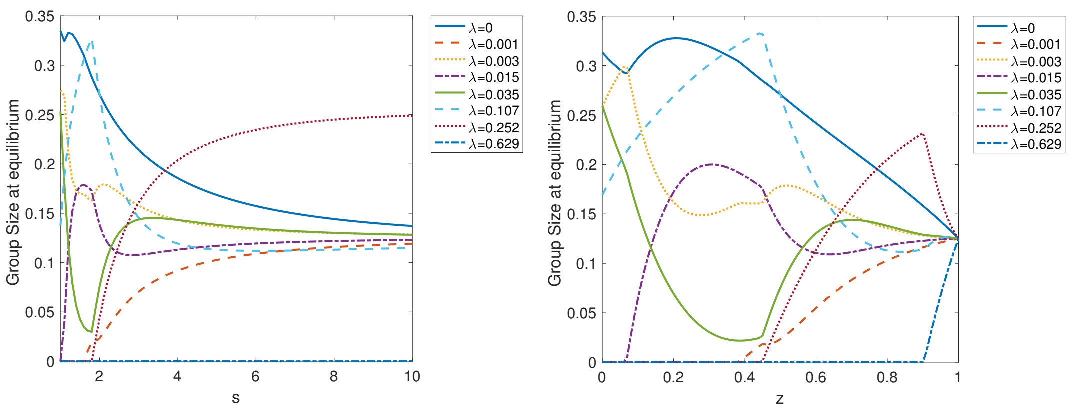

Figure 3.7: Equilibrium sizes for a system with eight groups, as parameters s or z varies. When not varying, $s = 2$ and $z = 0.5$. The distribution $R(r)$ is lognormal, and strictnesses are chosen such that every group is ranked most highly by an equal fraction of the population. Initial conditions set each subpopulation to an equal size.

CHAPTER 4

NASH EQUILIBRIUM IN SIMULTANEOUS-MOVE TWO-GROUP GAMES

We now move to a higher level case where groups have the flexibility to choose their strictness levels. There might be various types of groups with different utility functions and they compete with each other. The strictness level of each group will be selected strategically in order to maximize its own utility function in the group competition.

It is assumed that there are three types of groups:

Type A This type of group cares about the group production contributed by its members. Its utility function is the average time spent in group activities.

Type B This type of group is concerned about its membership or size, hence the utility function is the population of the group.

Type C This type of group would like to help everybody in the society to optimize their payoffs. So its utility function is the average individual payoff over the whole population.

The same simplification of the model as we discussed in the previous chapter is applied here. Groups will have a sufficiently large punishment level β to deter free-riding and we have the approximation that all the members of a group will play $t_i = \lambda$. Therefore, the goal of Type A groups is simplified to maximizing its strictness λ.

We are going to consider the simultaneous-move game between two groups of different types (A vs B, A vs C, B vs C). Through the choice of strictness λ, each group is trying to optimize the group payoff. However, due to the presence of the other group, the group goal might be affected, which makes the problem game-theoretical. The main objective of this chapter will be to study the Nash equilibrium of the two-group game.

4.1 Stationary Model

4.1.1 Two-Group Game without Unaffiliated Group

We start the discussion with a simple scenario, where there are only 2 groups and everybody can switch to her optimal group. We recall that since the individual payoff U_{ig} in group g is

$$U_{ig} = \sqrt{\lambda_g} + r_i(1 - \lambda_g),$$

if $\lambda_g < \lambda_{g'}$, then people with $r > r_{gg'} = 1/(\sqrt{\lambda_g} + \sqrt{\lambda_{g'}})$ will prefer group g to g', while people with $r < r_{gg'}$ will prefer group g' to g.

If group g has the same strictness level as g', we distinguish these two groups in the following way. The cutoff $r_{gg'}$ can be considered as the limit when the strictness λ_g approaches $\lambda_{g'}$ from the left. So people with $r > 1/(2\sqrt{\lambda_g})$ are in group g while people with $r < 1/(2\sqrt{\lambda_g})$ are in group g'. We refer to this scenario as undercutting. Alternatively, λ_g can approach $\lambda_{g'}$ from the right, leaving people with $r < 1/(2\sqrt{\lambda_g})$ in group g while people with $r > 1/(2\sqrt{\lambda_g})$ in group g'. This scenario is denoted as overcutting. The analysis below shows that undercutting or overcutting make no difference in the utility functions of a Type A or Type C group as the utilities are continuous. For a Type B group, its payoff when $\lambda_g = \lambda_{g'}$ is assumed to be the larger of the one by undercutting and by overcutting so that its utility is upper semi-continuous, as required by the proofs of Theorems 4.1 and 4.2.

To study the Nash equilibrium, we need to analyze the response of group strictness λ given the strictness level of the other group λ_g. The objective function of Type A group is independent of the choice of the other group. Hence the response of Type A group will always be $\lambda = 1$ regardless of the value λ_g.

Given λ_g, the utility function of Type B group with strictness λ becomes

$$U_B(\lambda) = \begin{cases} \int_{r_g}^{\infty} R(r)dr & \text{if } 0 \le \lambda < \lambda_g; \\ \max(\int_{1/(2\sqrt{\lambda_g})}^{\infty} R(r)dr, \int_{1/2}^{1/(2\sqrt{\lambda_g})} R(r)dr) & \text{if } \lambda = \lambda_g; \\ \int_{1/2}^{r_g} R(r)dr & \text{if } \lambda_g < \lambda \le 1, \end{cases} \tag{4.1}$$

where $r_g = 1/(\sqrt{\lambda_g} + \sqrt{\lambda})$ is the cutoff value of these 2 groups. If $\lambda \le \lambda_g$, to maximize U_B equates to minimizing the lower bound of the integral r_g, which is further equivalent to maximizing λ. So Type B group will undercut group g. On the other hand, if $\lambda > \lambda_g$, a similar analysis suggests that type B group will overcut group g. The optimal payoff of Type B group will be the larger of $U_B(\lambda_g^-)$ and $U_B(\lambda_g^+)$.

The objective function of Type C group with strictness λ is

$$U_C(\lambda) = \begin{cases} \displaystyle\int_{1/2}^{r_g} R(r)[\sqrt{\lambda_g} + r(1-\lambda_g)]dr + \\ \qquad \displaystyle\int_{r_g}^{\infty} R(r)[\sqrt{\lambda} + r(1-\lambda)]dr & \text{if } 0 \le \lambda < \lambda_g; \\ \displaystyle\int_{1/2}^{r_g} R(r)[\sqrt{\lambda} + r(1-\lambda)]dr + \\ \qquad \displaystyle\int_{r_g}^{\infty} R(r)[\sqrt{\lambda_g} + r(1-\lambda_g)]dr & \text{if } \lambda_g < \lambda \le 1. \end{cases} \tag{4.2}$$

Since everybody is able to switch to her favorite group, after introducing a Type C group, the members of this group are the only ones who have an increase in their payoffs. Therefore, the group goal can also be understood as to maximize the total increase in the payoffs of those who join it from outside. So $U_C(\lambda)$ can be rewritten as

$$U_C(\lambda) = \begin{cases} \displaystyle\int_{1/2}^{\infty} R(r)[\sqrt{\lambda_g} + r(1-\lambda_g)]dr + \\ \qquad \displaystyle\int_{r_g}^{\infty} R(r)[\sqrt{\lambda} - \sqrt{\lambda_g} + r(\lambda_g - \lambda)]dr & \text{if } 0 \le \lambda < \lambda_g; \\ \displaystyle\int_{1/2}^{\infty} R(r)[\sqrt{\lambda_g} + r(1-\lambda_g)]dr + \\ \qquad \displaystyle\int_{1/2}^{r_g} R(r)[\sqrt{\lambda} - \sqrt{\lambda_g} + r(\lambda_g - \lambda)]dr & \text{if } \lambda_g < \lambda \le 1. \end{cases}$$

$U_C(\lambda_g)$ is a continuous function. It is clear that undercutting or overcutting other groups by choosing the same strictness will give a Type C group no increase in the payoff of its members. Conversely, choosing a different strictness can at least attract and benefit some fraction of the population. Therefore U_C gets the global minimum at $\lambda = \lambda_g$.

Moreover, thanks to the fact $r_g = 1/(\sqrt{\lambda_g} + \sqrt{\lambda})$,

$$U_C'(\lambda) = \begin{cases} \int_{r_g}^{\infty} R(r)[\frac{1}{2\sqrt{\lambda}} - r]dr & \text{if } 0 \le \lambda < \lambda_g, \\ \int_{1/2}^{r_g} R(r)[\frac{1}{2\sqrt{\lambda}} - r]dr & \text{if } \lambda_g < \lambda \le 1. \end{cases}$$

We notice that $\lim_{\lambda \to 0} U_C'(\lambda) = \infty$, $U_C'(1) < 0$ as $r \ge 1/2$ and $\lambda = \lambda_g$ is the global minimum of U_C, so U_C attains the maximum on $(0, \lambda_g)$ at λ_1 and the maximum on $(\lambda_g, 1)$ at λ_2 where λ_1 and λ_2 are zero points of $U_C'(\lambda)$. The optimal payoff of Type C group should be the larger of $U_C(\lambda_1)$ and $U_C(\lambda_2)$. This group will choose between λ_1 and λ_2 whichever makes U_C larger. As λ_g varies, the global maximum might switch from one interval to the other. Hence the response of Type C group might not depend continuously on the given λ_g.

Based on the analysis above, we are able to study the Nash equilibrium of this two-group game. In the game between a Type A group and a Type B group, Type A group will respond $\lambda_A = 1$ at all time. In the meantime, Type B group will either undercut or overcut the Type A group. The cutoff point of r now becomes $1/(\sqrt{\lambda_A} + \sqrt{\lambda_B}) = 1/2$. As $r \ge 1/2$, Type B group will undercut Type A group by choosing $\lambda_B = 1^-$, taking all the population and eliminating its competitor completely.

The game between Type A group and Type C groups also has a pure Nash Equilibrium, $\lambda_A = 1$ and

$$\lambda_C = \arg\max_{\lambda \in [0,1)} \int_{1/2}^{r_g} R(r)dr + \int_{r_g}^{\infty} R(r)[\sqrt{\lambda} + r(1-\lambda)]dr$$

where the cutoff $r_g = 1/(1+\sqrt{\lambda})$. Type C group will choose a lower strictness than λ_A to attract a fraction of the total population and raise their payoffs.

There is no pure Nash Equilibrium in the simultaneous-move game between Type B and Type C groups. It is known that Type B group prefers the same strictness as the other group while Type C group would rather have a different one. As a result, at no point in $[0,1] \times [0,1]$ will both groups be at a utility maximum simultaneously, so there cannot be a pure Nash Equilibrium. Theorem 4.1 guarantees the existence of the mixed Nash equilibrium.

Theorem 4.1. *There exists mixed Nash equilibrium of the game between Type B and Type C groups*

$$G = ([0,1] \times [0,1], (U_B, U_C))$$

where U_B and U_C are defined in (4.1) *and* (4.2).

The proof of Theorem 4.1 can be found in the Appendix. To illustrate the mixed Nash equilibria, we discretize the strategy space $[0,1]$ and approximate the game G by finite games. We use Lemke-Howson algorithm [21, 22, 23] to compute the mixed Nash equilibria of the finite approximation. Fig. 4.1 presents the mixed Nash equilibrium when the r distribution is $r - 1/2 \sim \text{Lognormal}(\mu, \sigma^2)$ with $\mu = -1/2$ and $\sigma = 2$, the same as in the previous chapter. Type B group tends to choose a strictness near 0 while Type C group has a probability around 0.41 to choose λ slightly above 0, a probability around 0.53 to choose $\lambda \approx 0.455$ and some very small probabilities to choose λ in between these two values.

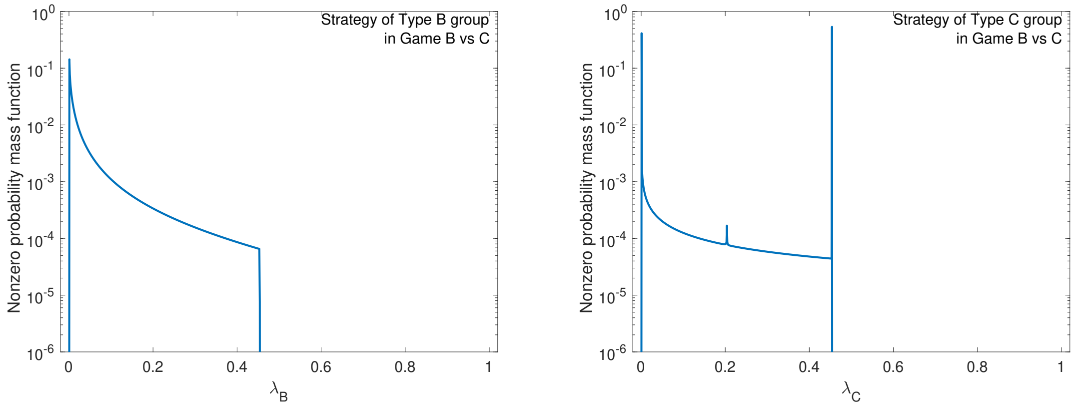

Figure 4.1: Mixed Nash equilibria of finite approximation of game B vs C. Stationary model without unaffiliated group is used. (Left) Strategy of Type B group; (Right) Strategy of Type C group.

4.1.2 Two-Group Game with Unaffiliated Group

In this subsection, the unaffiliated group is taken into consideration. It is assumed that a small minimum strictness level λ_{min} exists in order to prevent overcutting of the unaffiliated group.

Given the strictness levels of the three groups, $0 = \lambda_0 < \lambda_g < \lambda_{g'}$, there are 2 cutoff points $r_{0g} = 1/\sqrt{\lambda_g}$ and $r_{gg'} = 1/(\sqrt{\lambda_g} + \sqrt{\lambda_{g'}})$. The unaffiliated group will take the population with $r > r_{0g}$, group g will have the population with $r_{gg'} < r < r_{0g}$ and group g' will contain the rest.

The response of Type A group will still be $\lambda = 1$ in this case. The utility function of Type B becomes

$$U_B(\lambda) = \begin{cases} \int_{r_g}^{r_0} R(r)dr & \text{if } \lambda_{min} \leq \lambda < \lambda_g; \\ \max(\int_{1/(2\lambda_g)}^{1/\lambda_g} R(r)dr, \int_{1/2}^{1/(2\lambda_g)} R(r)dr) & \text{if } \lambda = \lambda_g; \\ \int_{1/2}^{r_g} R(r)dr & \text{if } \lambda_g < \lambda \leq 1, \end{cases} \tag{4.3}$$

where $r_0 = 1/\sqrt{\lambda}$ and $r_g = 1/(\sqrt{\lambda_g} + \sqrt{\lambda})$. On the closed interval $[\lambda_{min}, \lambda_g]$, the maximum of U_B can either be attained at the endpoints or at a point λ^* in the interior where

$U_B' = 0$. Hence the choice of Type B group will be the best of the following

1. $\lambda = \lambda_{min}$, the minimum of allowed strictness levels;
2. $\lambda = \lambda^* \in (0, \lambda_g)$ if U_B has local maxima on $(0, \lambda_g)$;
3. $\lambda = \lambda_g^-$ to undercut group g;
4. $\lambda = \lambda_g^+$ to overcut group g.

After adding the unaffiliated group, the utility function of Type C group becomes

$$U_C(\lambda) = \begin{cases} \int_{1/2}^{r_g} R(r)[\sqrt{\lambda_g} + r(1-\lambda_g)]dr + & \\ \int_{r_g}^{r_0} R(r)[\sqrt{\lambda} + r(1-\lambda)]dr + \int_{r_0}^{\infty} R(r)rdr & \text{if } \lambda_{min} \le \lambda < \lambda_g; \\ \int_{1/2}^{r_g} R(r)[\sqrt{\lambda} + r(1-\lambda)]dr + & \\ \int_{r_g}^{r_{0g}} R(r)[\sqrt{\lambda_g} + r(1-\lambda_g)]dr + \int_{r_{0g}}^{\infty} R(r)rdr & \text{if } \lambda_g < \lambda \le 1. \end{cases} \tag{4.4}$$

Since undercutting group g will never benefit Type C group, similar results apply here that the response of Type C group is the better of the local maximum point on $[\lambda_{min}, \lambda_g)$ and the local maximum point on $(\lambda_g, 1]$.

We investigated the Nash equilibrium of this two-group game using some lognormal distributions $r - 1/2 \sim \text{Lognormal}(\mu, \sigma^2)$. The minimum strictness λ_{min} is 0.001. If Type A group is in the game, it will always set its strictness λ_A to 1 regardless of the choice of the other groups. Fig 4.2 shows the optimal choices of Type B and Type C groups given that $\lambda_A = 1$ as μ and σ vary. When σ is small, the r values are close to the mean value. It is easy to attract most people by placing the strictness level close to the average preferred strictness. As μ grows, people tend to prefer less strict groups. So the response of Type B group λ_B drops continuously from 1^- to λ_{min}. On the other hand, when σ is big, r values are more widely spread. Putting strictness in $(\lambda_{min}, 1)$ might not be so appealing to Type B group. Thus it instead chooses $\lambda_B = 1^-$ for small μ values and $\lambda_B = \lambda_{min}$ for big μ values.

The response of Type C group λ_C decreases both in μ and in σ. That is because, when μ increases, Type C group can obtain a larger population by choosing a smaller strictness; when σ increases, more people prefer the Type A group, therefore Type C group will lower its strictness to benefit more people from the unaffiliated group.

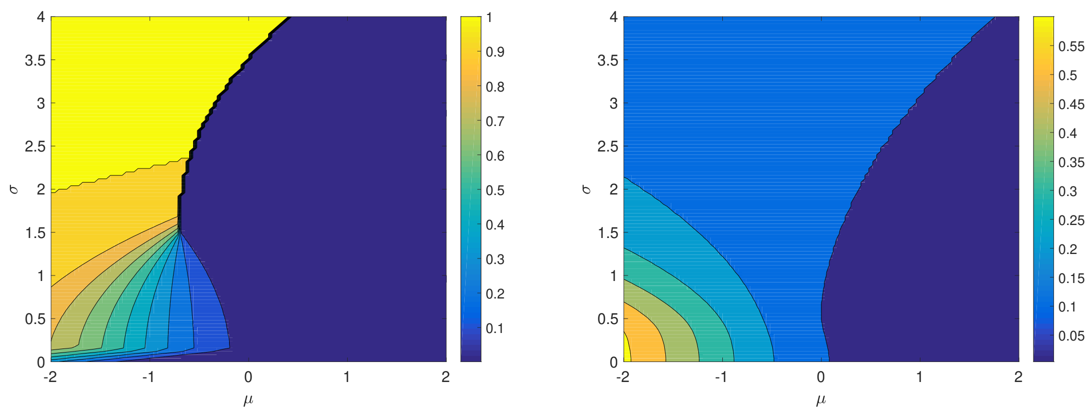

Figure 4.2: Contour plots of the optimal choices of Type B and Type C groups given that $\lambda_A = 1$ as μ and σ vary. (Left) Type B group; (Right) Type C group.

Theorem 4.2 guarantees the existence of the mixed Nash equilibrium in the game between Type B and Type C groups as long as λ_{min} is sufficiently small. The proof is in the appendix.

Theorem 4.2. *There exists $\lambda_* > 0$ such that if $\lambda_{min} < \lambda_*$ then there exists mixed Nash equilibrium of the game between Type B and Type C groups*

$$G = ([\lambda_{min}, 1] \times [\lambda_{min}, 1], (U_B, U_C))$$

where U_B and U_C are defined in (4.3) *and* (4.4).

However, there may also exist pure Nash equilibria in the game between Type B and Type C groups. For example, with the same lognormal distribution of R as in the previous discussions, namely $r - 1/2 \sim \text{Lognormal}(\mu, \sigma^2)$ with $\mu = -1/2$ and $\sigma = 2$, there exists a pure Nash equilibrium. Fig 4.3 shows the optimal responses of Type B group and Type C group as λ_g varies. The two curves intersects at $(\lambda_B, \lambda_C) = (0.001, 0.382)$, which is a pure Nash Equilibrium.

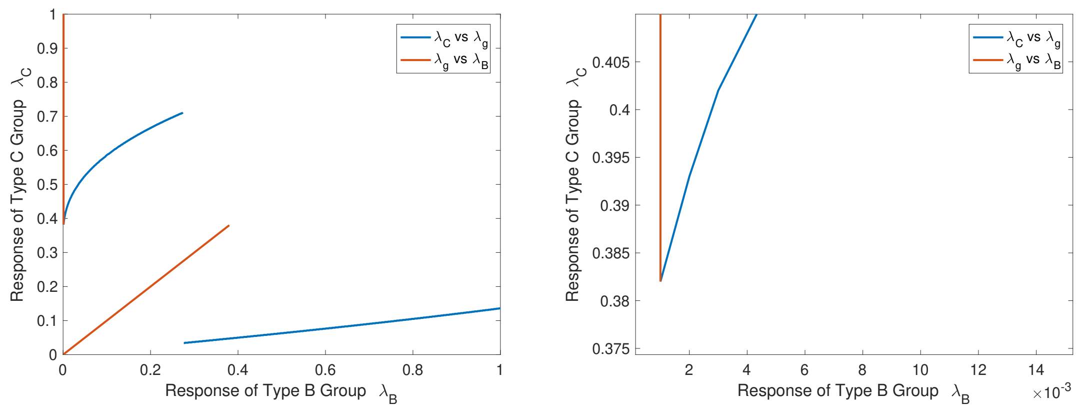

Figure 4.3: The responses of Type B group and Type C group as λ_g varies and a zoomed-in version near the intersection. The blue line presents the response λ_C in the vertical axis given λ_B. The red line shows the response λ_B in the horizontal axis given λ_C.

4.2 Dynamic Model

In this section, we investigate the best responses of all types of groups given the strictness of the other group λ_g on the prior dynamic model that assumes birth, death and inheritance among group members and allows each individual to have one chance in their lifetime to change their affiliation. The initial condition is chosen so that everybody is in their favorite group.

We choose parameter values $s = 3$ and $z = 0.5$, use uniform birthrates, and use a lognormal distribution for R so that $r - 1/2 \sim \text{Lognormal}(\mu, \sigma^2)$, where $\mu = -1/2$ and $\sigma = 2$. The minimum allowed strictness for any affiliation is $\lambda_{min} = 0.001$. In the dynamic model, the λ that maximize the group utility might possibly lead to zero population at the equilibrium, but the group goal will never be accomplished in an empty group. As a result, it is assumed as well that groups will attempt to maintain a minimum population $\epsilon = 0.01$ in addition to optimizing their goals. We recall here that in the case with the unaffiliated group and an affiliated one, a single group with strictness up to approximately 0.83 is able to survive alongside the unaffiliated group under these parameters and distribution. This strictness is referred as the absolute maximal strictness.

Fig. 4.4 presents the optimal responses of all types of groups. The immediate observation comes out that if λ_g is larger than the absolute maximal strictness, it will be eliminated in the game. In this case, each type of group is able to perform their optimal choice assuming that only the unaffiliated group is present. So type A group will choose approximately 0.83; type B group will undercut the unaffiliated group from the right; and type C group will select around 0.31.

The choice of Type A group λ_A depends continuously on λ_g, but not monotonically. When λ_g is close to 0, λ_A is close to the absolute maximal strictness. As λ_g grows, Type A group tends to lower its strictness at first. By doing so, this group will still maintain a considerable fraction of the population that prefers it in order not to die out in the group competition. When λ_g gets beyond approximately 0.58 but less than the absolute maximal strictness, the best choice of Type A group becomes to undercut the other group. Any strictness above that will result in a group population smaller than the threshold ϵ.

There is no continuous relationship between λ_B and λ_g. A small λ_g will cause type B group to overcut it, stealing most of its members. When λ_g grows to around 0.5, it is more appealing for type B group to select λ_{min} instead. For an even larger λ_g, a choice between 0 and λ_g will give the largest population. This can also be seen in Fig. 3.5. If we view the preexisting group as group g and the new group as the Type B group, then the figure shows the population at the equilibrium with λ_g in $0.25, 0.5, 0.75$ and λ_B varying in $[0, 1]$. It is clear from the figure that when $\lambda_g = 0.25$, the optimal response is $\lambda_B = 0.25$; when $\lambda_g = 0.5$, the optimal response becomes $\lambda_B = 0$; when $\lambda_g = 0.75$, the optimal response lies near 0.1.

Type C group still will not pick the same strictness as other groups, which is as expected. Type C group will choose approximately 0.31 if there is only the unaffiliated group existing. We refer to this baseline strictness as λ_C^*. When $0 < \lambda_g < 0.22$, the response

$\lambda_C > \lambda_C^*$; when $0.22 < \lambda_g < 0.83$, the response $\lambda_C < \lambda_C^*$. When $\lambda_g > 0.83$, $\lambda_C = \lambda_C^*$ as the group g dies out. Comparing this with the response in the stationary model in Fig. 4.3, we have several observations. (1) With λ_g being small, Type C group tends to be more strict in the stationary model. In the dynamic model, the probability α to switch to a high strictness group is comparatively small, so Type C group needs to lower its strictness to increase the chance of each individual switching to their optimal group in order to maximize the group objective; (2) with a relatively larger λ_g around 0.7, the response in the dynamic model is larger. That is because in this case, group g will have a much smaller size at the equilibrium in the dynamic model than in the stationary model. A slightly larger strictness will benefit those who prefer group g but end up on the Type C group.

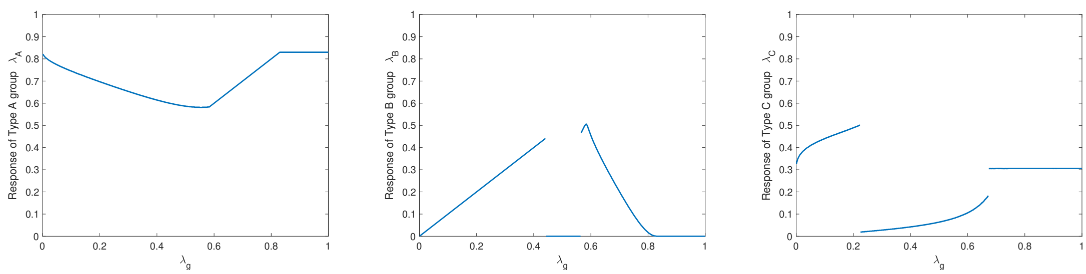

Figure 4.4: Optimal responses of the three types of groups as the strictness λ_g varies. In all cases, $s = 3$, $z = 0.5$, birthrates are uniform, and the distribution R is lognormal with parameters given in the text. (Left) the response of Type A group; (Center) the response of Type B group; (Right) the response of Type C group.

By combining the responses of two different types of groups in one figure, we are able to analyze the pure Nash equilibrium of the game by looking at the intersection of response curves in Fig. 4.5. There is an intersection of the curves of Type A group and Type B group around $(\lambda_A, \lambda_B) = (0.587, 0.5)$, which is a pure Nash equilibrium of the game between these two groups. Clearly no intersections exist in the figure A vs C or B vs C, which implies that there are no pure Nash equilibria in the game involving Type C group. We use Lemke-Howson algorithm to obtain mixed Nash equilibria of finite approximations of games A vs C and B vs C respectively, as shown in Fig. 4.6. In the game A vs C, Type A group chooses strictness levels in the interval $[0.65, 0.71]$, while Type C group has a very large probability to play 0.185. In the game B vs C, the average strictness of the strategy of Type B group is 0.27 while the average strictness of the strategy of Type C group is 0.37. It is notable that the Nash equilibria of the games on the dynamic model are obtained assuming that once the group strictnesses are selected, they will never be changed. However, in the real religion market, groups can change their strictness levels over time, leading to more complicated results.

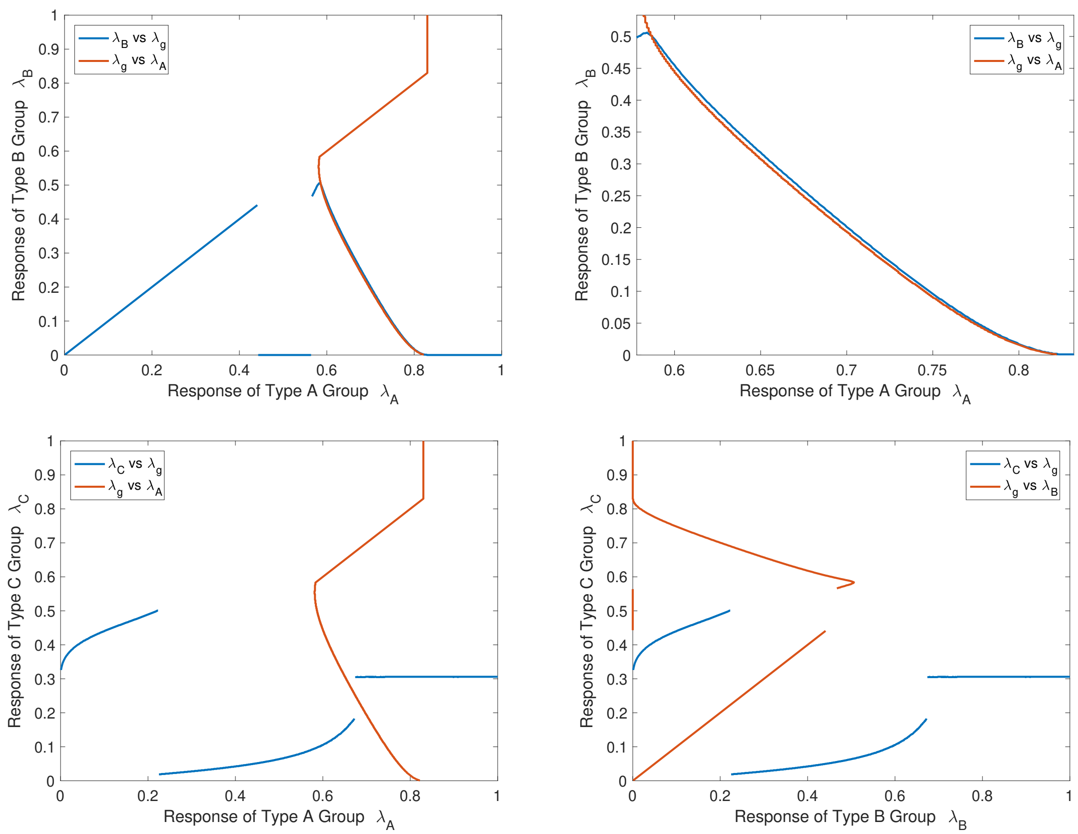

Figure 4.5: Responses of two of the three groups to various λ_g. In all cases, $s = 3$, $z = 0.5$, birthrates are uniform, and the distribution R is lognormal with parameters given in the text. (Top) the response of Type A and Type B group and the zoomed-in version near the region where the two curves are close; (Bottom Left) the response of Type A and Type C group; (Bottom Right) the response of Type B and Type C group.

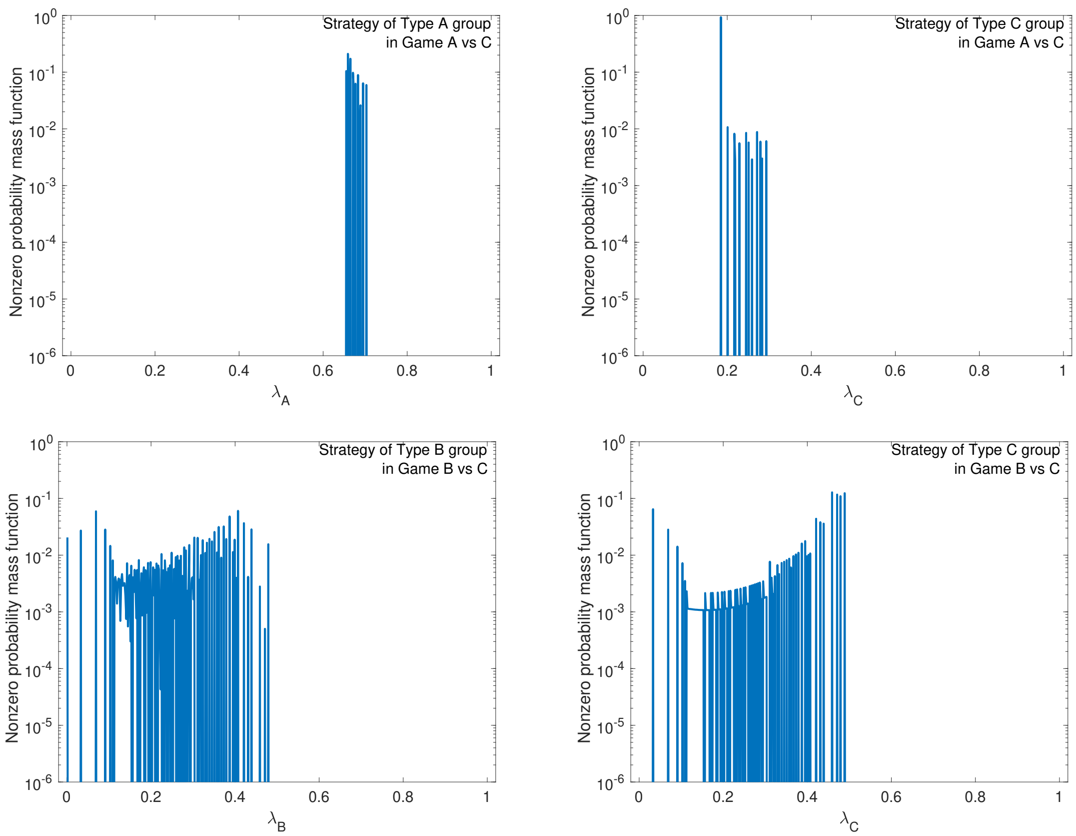

Figure 4.6: Mixed Nash equilibria of finite approximations of games A vs C and B vs C. Dynamic model is used. (Top Left) Strategy of Type A group in the Game A vs C; (Top Right) Strategy of Type C group in the Game A vs C; (Bottom Left) Strategy of Type B group in the Game B vs C; (Bottom Right) Strategy of Type C group in the Game B vs C.

CHAPTER 5

SIMULATIONS ON SEQUENTIAL STACKELBERG GAMES

In the previous chapter, we analyzed the strategies of choosing strictness of various types of groups in the game of two or three groups which come into existence simultaneously. However, in the real religious marketplace, new affiliations are founded one at a time and enter the market sequentially and successively. In this chapter, we are going to present several simulation results of such procedures.

We will consider the same three types of groups as in the previous chapter. Originally everybody is unaffiliated, in other words, there is only the unaffiliated group. The groups then come into the marketplace sequentially. Each group has to take a different strictness level from all the other existing groups in the market in order to distinguish itself. The prioritized goal of all the group is to maintain a size above some population threshold ϵ because if a group becomes too small, there is a high probability that it will disappear completely, given the group switching rules of the model. Once groups meet the requirement of minimum population, they will maximize their own utilities. Moreover, when a group is joining the marketplace, it thinks one more step ahead that there will be another group following it. We refer to these two groups as the current group and the following group, respectively, for short. The following group has an equal chance to be any one of the considered types. For example, if it is assumed that all the groups are of Type B, then the following group can only be a Type B group; if all three types are taken into account, then the following group has $1/3$ chance to be each of those three types. The following group also prioritizes getting a size at least ϵ. Given a λ value of the current group, the optimal strictness of the following group λ_{foll}^{opt} of any given type can be calculated numerically. The current group then chooses a strictness level λ_{curr}^{opt} that can maximize the expectation of its utility function given λ_{foll}^{opt} of each type while keeping the size above ϵ. The whole procedure can be viewed as a series of Stackelberg games [24].

It is also assumed that after a group enters, the groups that die out will be removed from the market. So a new group can retake the same strictness as a previous group that has been eliminated before. We remark here that the current group will not come to the marketplace if no λ values can satisfy the minimum population condition. In addition, a Type C group will also not enter if existence of this group will lower the average payoff of individuals within the society, that is, if its maximum utility is negative, as this is against the goal of a Type C group.

We discretize the domain of λ into meshes and only allow groups to choose the λ values on the grid points denoted by the points Λ. We refer to the set of strictnesses of all the previous groups that have not died out by Λ_{active}. The minimum possible strictness λ_{min} of any affiliation is automatically given by the mesh. The algorithm for adding one group is presented in Algorithm 5.

5.1 Stationary Model with Switching Penalty

We first simulate this procedure using the "stationary multigroup model". We do not take birth, death, or inheritance into account. Further, we allow everybody to switch groups infinite times, but individuals will only switch to a group with a utility higher than their current group. Specifically we also introduce a penalty p in group switching. People will switch to a new group only if they can gain p more there. This is done to prevent undercutting and overcutting to some extent. Otherwise, a new group of the same type as a pre-existing group could under or over cut them and gain all of their members, without really offering anything new to those individuals. This penalty term also causes this stationary model to exhibit some history dependence, as the order in which group arise will affect where individuals end up; this would not happen in a penalty-free model if infinite switching were allowed. The cutoff r value for members of a preexisting group to switch to a new group is

Algorithm 5.1 Algorithm of adding one group in the simulation of sequential Stackelberg games.

for λ_{curr} in $\Lambda - \Lambda_{active}$ **do**
  if The population of current group is below ϵ **then**
    continue
  end if
  for Group type of following group in set of group types **do**
    for λ_{foll} in $\Lambda - \Lambda_{active} - \{\lambda_{curr}\}$ **do**
      Evaluate the goal of following group U_{foll}. Goal is $-\infty$ if population is below ϵ.
    end for
    Obtain λ_{foll}^{opt} and U_{foll}^{opt} in the for loop.
    if λ_{foll}^{opt} exists **then**
      Evaluate the goal of the current group $U_{curr,type}$ given λ_{foll}^{opt}. Goal is $-\infty$ if population is below ϵ.
    else
      Evaluate the goal of the current group $U_{curr,type}$ without λ_{foll}^{opt}. Goal is $-\infty$ if population is below ϵ.
    end if
  end for
  $U_{curr} \leftarrow$ the average of all the $U_{curr,type}$'s
end for
Obtain λ_{curr}^{opt} and U_{curr}^{opt} in the for loop. λ_{curr}^{opt} does not exist if no groups can be added.
if λ_{curr}^{opt} exists **then**
  $\Lambda_{active} \leftarrow \Lambda_{active} \cup \{\lambda_{curr}^{opt}\}$
end if
Remove the strictnesses of all the groups with zero population from Λ_{active}

given by

$$r = \frac{\sqrt{\lambda_{preexisting}} - \sqrt{\lambda_{new}} + p}{\lambda_{preexisting} - \lambda_{new}}.$$

We note that if $\lambda_{preexisting}$ is fixed, r does not decrease monotonically as λ_{new} grows in some neighborhood of $\lambda_{preexisting}$, unlike the penalty-free switching. This means that in some extreme cases a group can choose its strictness between 2 other preexisting groups and eliminate one of them as shown in the example below.

Suppose 4 groups with strictnesses $\lambda_1 = 0.04$, $\lambda_2 = 0.033$, $\lambda_3 = 0.004$, $\lambda_4 = 0.021$ are sequentially added, with penalty $p = 0.01$. The r ranges of people in each group after a group is added are listed in Table 5.1. Note that, after the joining of the fourth the group, the second group dies out. The second group will survive in this scenario if there is no penalty in switching.

Table 5.1: An example of the stationary model with switching penalty where a group joins between two other groups and eliminates one of them. The intervals in the cells indicate the range of r values in each group.

With switching penalty	First Group	Second Group	Third Group	Fourth Group
After the Third Group Joins	$(0.5, 4.05)$	$(4.05, 4.43)$	$(4.43, 13.31)$	
After the Fourth Group Joins	$(0.5, 3.43)$	$\emptyset$	$(4.43, 13.31)$	$(3.43, 4.43)$
Without switching penalty	First Group	Second Group	Third Group	Fourth Group
After the Third Group Joins	$(0.5, 2.62)$	$(2.62, 4.08)$	$(4.08, 15.81)$	
After the Fourth Group Joins	$(0.5, 2.62)$	$(2.62, 3.06)$	$(4.80, 15.81)$	$(3.06, 4.80)$

The following r distribution is used in the simulation: $r - 1/2 \sim \text{Lognormal}(\mu, \sigma^2)$ with $\mu = -0.5$ and $\sigma = 2$. The population threshold is $\epsilon = 0.01$ and the penalty term is $p = 0.01$. The λ domain is discretized evenly into 1000 parts such that all the feasible λ's are a positive multiple of $1/1000$.

We first assume that all the groups that will join the marketplace are of the same type. In this simulation, when a groups thinks one more step ahead, it only needs to consider a group that has the same type as itself. The simulation terminates if no group can be added. Fig.5.1 shows the strictnesses of the group added at each step and their final sizes for all the three types.

Seventy-six Type A groups can coexist under this distribution. The first Type A group being added has strictness $\lambda = 1$. When the first type A group is added, it takes around 0.458 of the population leaving 0.542 to the unaffiliated one. All the following groups have to locate themselves between two existing groups and tend to be as close to the stricter one as possible. We observe that the strictnesses of following groups decreases monotonically until no more groups can be added. The sizes of some groups still drop below the threshold ϵ. That is because when a group joins the market, it does not consider that there is a second group coming after the following group. However when the following group is being added, it does consider the second group. So in order to survive, the λ value that the following group chooses might be a little bit smaller than what the current group thinks. As the r cutoff of the current group with λ_{curr} and a following group with λ_{foll} does not depend monotonically on λ_{foll} when a switching penalty exists, an even smaller λ_{foll} might harm the size of the current group instead of helping it. For example, the second group has strictness 0.978 and thinks the third group will play 0.957. But the third group instead chooses 0.956, making the population of the second group drop from 0.011 to 0.002.

Unlike Type A groups, only nine Type B groups can be added to the marketplace. The strictness of the first Type B group is $\lambda = 0.311$. The Type B groups that join later just select a strictness level in between other preexisting groups attempting to maintain a minimum population while maximizing their size. All the groups can end up with a size larger than ϵ. The average strictness of all the Type B groups is 0.2826.

Eight Type C groups coexist in the end. The strictness of the first Type C group is $\lambda = 0.552$, larger than that of first Type B group. That is because by choosing such a strictness, it might not gain a population as large as the Type B group, but it will benefit more on the payoffs of its members and increase the overall average payoff. We also observe that the first Type C group even has a larger population in the end than the first of the Type B groups. Because purely aiming at a high size does not necessarily give the best increase in the overall average payoff, Type C groups do not compete for population size as much as Type B groups, which gives opportunities for the first Type C group to preserve a relatively large size. All the Type C groups that have entered the marketplace have an average strictness at 0.2445.

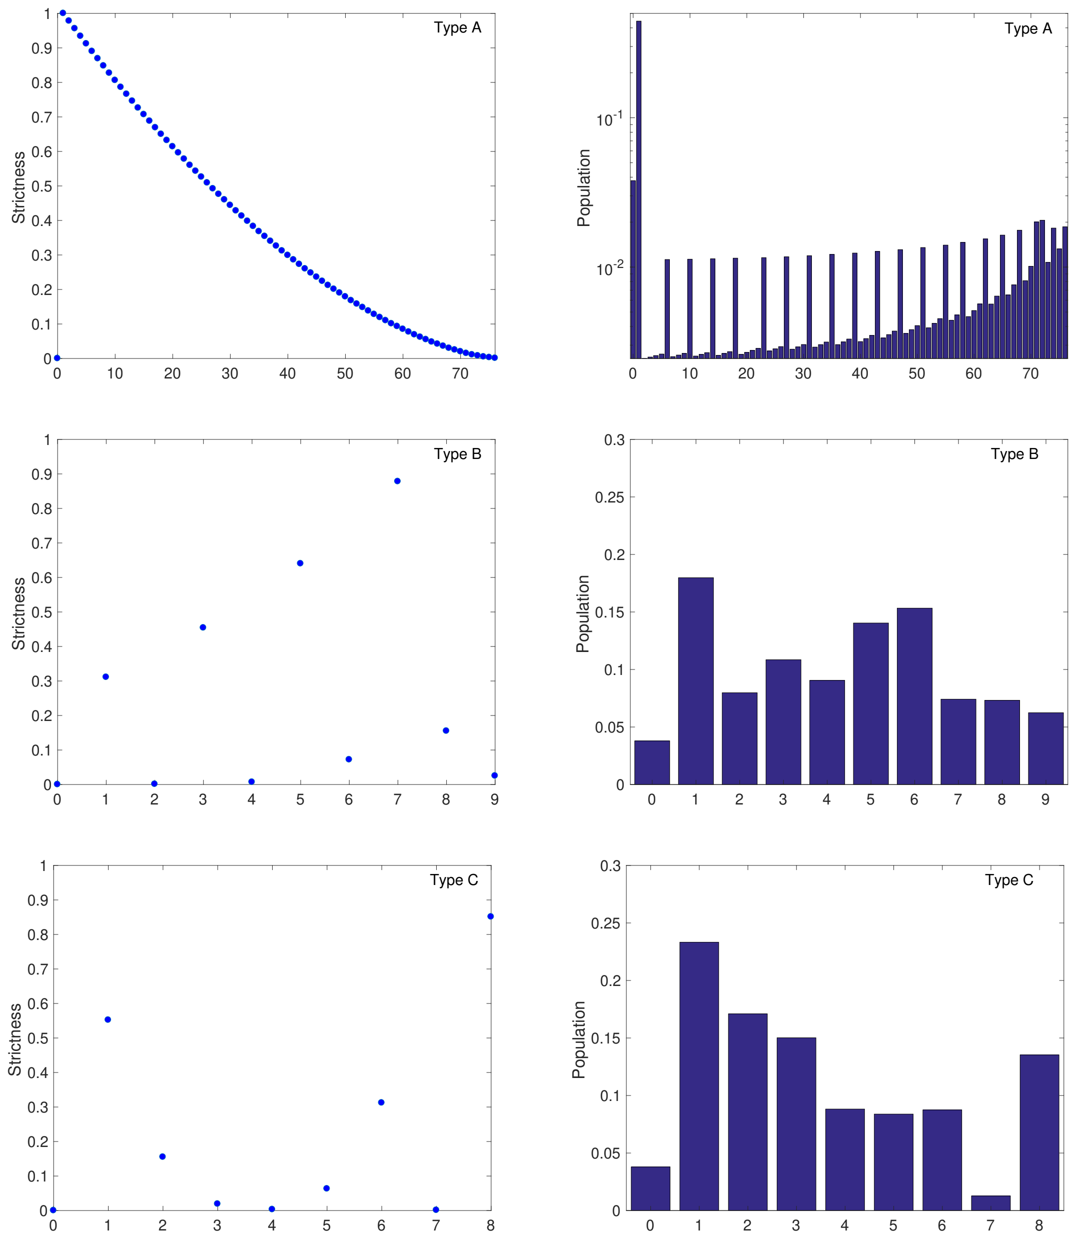

Figure 5.1: Simulation of the sequential Stackelberg game on the stationary model with switching penalty. The left column shows the strictnesses of groups added at each step and the right column presents the population of each group in the end of the simulation. (Top) Type A groups only; (Middle) Type B groups only; (Bottom) Type C groups only.

We also run the simulation with all three group types available. The procedure will end if 10 affiliations have been added. We iterate over all possible combinations of group types and present the statistics of group strictnesses in Table 5.2 and Fig. 5.2. We notice that as expected, the Type A groups are much stricter than the other two types. Type C groups are slightly stricter than Type B groups.

Table 5.2: Statistics of simulation of the sequential Stackelberg game on the stationary model with switching penalty. Groups have an equal chance to be any of the three types.

Strictnesses	Type A Groups	Type B Groups	Type C Groups
Maximum	1	0.754	0.778
75-percentile	1	0.311	0.297
50-percentile	0.76	0.07	0.069
25-percentile	0.34	0.007	0.017
Minimum	0.001	0.001	0.001
Average	0.6504	0.1721	0.1747

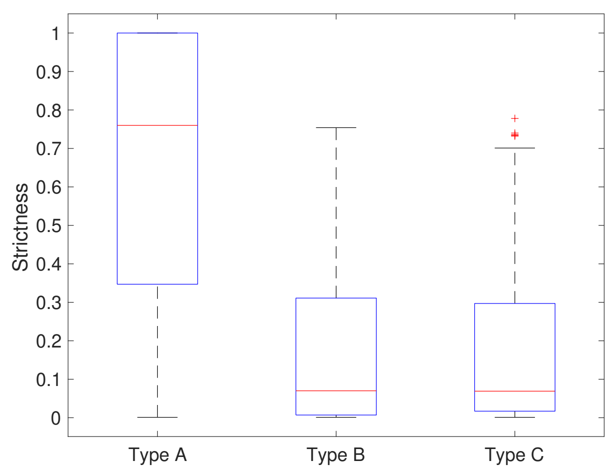

Figure 5.2: Boxplot of simulation of the sequential Stackelberg game on the stationary model with switching penalty. Groups have an equal chance to be any of the three types.

5.2 Dynamic Model

The procedure of adding groups is also simulated on the dynamic model described in Chapter 3. When a new group is founded, it initially has an entirely new set of individuals according to the population-wide r distribution R with size $\epsilon/(1-\epsilon)$. We then renormalize the whole population to 1. After renormalization, the new group has a population which is equal to the minimum population threshold ϵ and the sizes of all the preexisting groups are multiplied by $(1-\epsilon)$ while the r distribution in each group remains intact. The equilibrated population after the entering of a new group can be obtained by solving the ODE (3.3) with the initial condition mentioned above. In numerical simulations, we set $n_{g\sigma}$ in the ODE to zero if the value drops below some small threshold and a group g is considered to be eliminated if $\sum_\sigma n_{g\sigma} = 0$. Moreover, to avoid people switching to a group with similar strictnesses, the probability of switching from group g to group g' $\alpha_{g,g'}$ (3.2) is multiplied by a factor $\frac{1}{1+e^{10-1000\triangle\lambda}}$ where $\triangle\lambda = |\lambda_g - \lambda_{g'}|$ is the difference of the strictness levels of those two groups. When $\triangle\lambda$ is very small, the probability of switching is also close to zero; when $\triangle\lambda$ is larger than around 0.02, the factor is almost 1, so the probability of switching is nearly the same as the original one. By doing so, undercutting and overcutting are prevented to some extent, but not completely. That is because by choosing a similar strictness level to a preexisting group, even if the new group can steal negligible population from the preexisting group, it can still beat that preexisting group in the group rankings of people in some other groups and attract those people.

The same $R(r)$ and population threshold as the previous section are used, namely $r-1/2 \sim \text{Lognormal}(\mu, \sigma^2)$ with $\mu = -0.5$ and $\sigma = 2$ and $\epsilon = 0.01$. In the ordinary differential equation, the parameters are $s = 3$ and $z = 0.5$. The λ domain is discretized evenly into 500 parts such that all the feasible λ's are a positive multiple of 1/500.

The first simulation is done under the assumption that all the groups that will enter the market are of the same type. Due to the intensive nature of this computation, the program was run on a compute cluster in parallel. The results are presented in Fig. 5.3. In all simulations the strictest group has λ no more than 0.6. If a group chooses an even higher strictness, it will easily be eliminated by the following group. The first Type A group has strictness 0.564. At each step, there is always only one group other than the unaffiliated one that has population larger than 0.1 and the strictness level of that group is decreasing along the simulation. We also notice that the affiliation with the largest size is the least strict among all the affiliations. The first Type B group has strictness 0.468. After 100 steps, around nine groups including the unaffiliated one can coexist with size larger than 0.01. The strictness levels of the eight affiliations are spaced at a distance around 0.08. The simulation of Type C groups has a quite different behavior from the other two. It terminates at 10 steps. No more groups can enter the market because they will either fail to maintain a minimum population size or will decrease the average payoff of individuals within the society by eliminating some preexisting groups and causing those members to leave their ideal group .

Similar to the previous section, we also run the simulation with the assumption that the groups have an equal chance to be any of the three types and a random type of group will join the market at each round. We end the procedure if 7 affiliations have been added and search over all possible combinations of group types. The results are shown in Table 5.3 and Fig. 5.4. Only groups with population above the threshold $\epsilon = 0.01$ are taken into account. We observe that there are clear distinctions among the strictnesses of these three types of groups. Type A groups have an average strictness 0.5722 with over 75% of them being no more than 0.03 below or above the average. Type B groups are generally less strict than Type A groups. The average strictness of Type B groups is 0.3436. Type C groups have an average strictness 0.1322, which are the least strict. The strictnesses of 75% of the Type C groups are no more than 0.220. By choosing low strictnesses, Type C groups can raise the chance of people who prefer it to join it in the dynamic model. The overall payoff of the whole population will grow if more people can stay in their favorite groups. It is also notable that among the 3^7 possible combinations, on average 5.27 groups can have final population above ϵ.

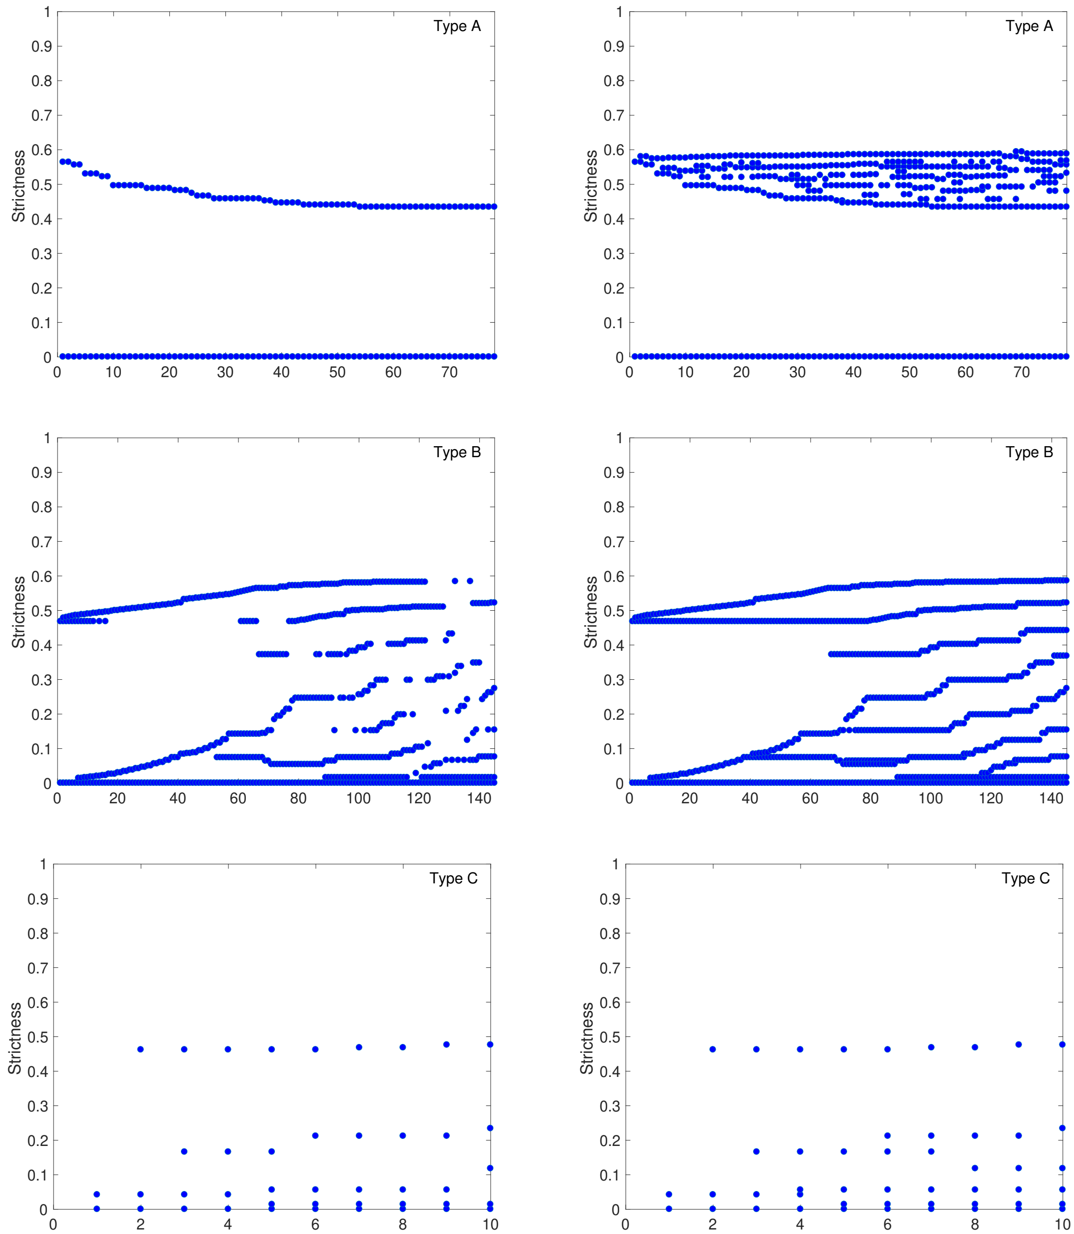

Figure 5.3: Simulation of the sequential Stackelberg game on the dynamic model with modified α. The left ones show the strictness levels of groups with size larger than 0.1 at each step while the right ones present the strictness levels of groups with size larger than the population threshold 0.01 at each step. (Top) Type A groups only; (Middle) Type B groups only; (Bottom) Type C groups only.

Table 5.3: Statistics of simulation of the sequential Stackelberg game on the dynamic model with modified α. Groups have an equal chance to be any of the three types.

Strictnesses	Type A Groups	Type B Groups	Type C Groups
Maximum	0.608	0.562	0.590
75-percentile	0.594	0.480	0.220
50-percentile	0.584	0.394	0.056
25-percentile	0.558	0.164	0.028
Minimum	0.288	0.026	0.014
Average	0.5722	0.3436	0.1322

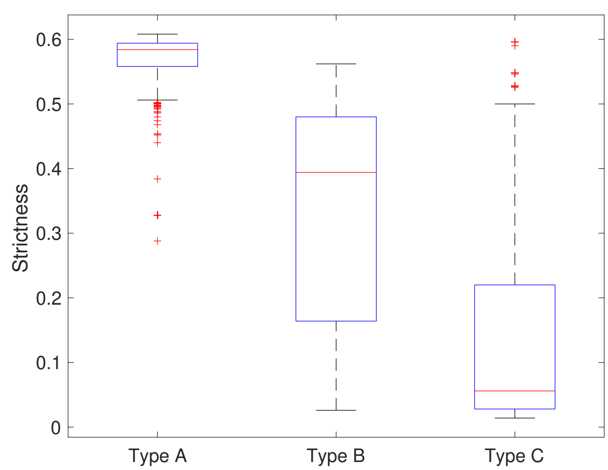

Figure 5.4: Boxplot of simulation of the sequential Stackelberg game on the dynamic model with modified α. Groups have an equal chance to be any of the three types.

CHAPTER 6
CONCLUSIONS

In this dissertation, we have constructed a dynamic model for the sizes of religious groups, based on a unidimensional categorization of groups by their strictness level λ, 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 as in the stationary model 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 its members by sheer number, and to join lower strictness groups, as those individuals spend more time out of group during which you might encounter them. 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 birthrates of the various groups, are summarized by a system of ordinary differential equations for the various population sizes in time.

Analysis of the dynamic model has confirmed several phenomenon seem 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 their 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 birthrate of its members is high enough in relation to the birthrate of non-members.

We also briefly examined from a numerical viewpoint a scenario with several (8) 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. 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.

Going beyond the multi-group model, we studied how groups might arrange themselves with regard to strictness to optimize the goals of the group, be they simply maximizing the contribution of their members (Type A), gaining the largest size (Type B), or maximizing the average utility of the whole population in the society (Type C). Pure Nash equilibria are not guaranteed to exist in the simultaneous-move games between two different types of groups. Analysis of the games has shown that Type A group has a high strictness; Type B group often tends to undercut or overcut the other group in the stationary model; Type C group wants to distinguish itself from other groups.

Finally, we analyzed the effects of new groups of these three types entering the religious market. This procedure can be viewed as a sequence of Stackelberg games. Simulations show how the three types of religious groups differentiate themselves with regard to strictness. In the stationary model with switching penalty, Type A groups have the highest strictness levels while Type B groups and Type C groups have roughly the same strictness level. In the dynamic model, Type A groups are the strictest; Type B groups are moderate; Type C groups tend to have low strictness levels.

Appendices

APPENDIX A

PROOFS OF THEOREMS 4.1 AND 4.2

Before proving the theorems, we present some basic concepts in game theory as well as some definitions and results introduced by Reny et. al. [25] and Allison et. al. [26].

We denote an N-player game by $G = (X_i, u_i)_{i=1}^N$ where X_i is the strategy space and $u_i : X \to \mathbb{R}$ where $X = \times_i X_i$ is the payoff function of of player i. The subscript $-i$ means all the other players except i. If each X_i is a pure strategy space, we can extend each u_i to $M = \times_i M_i$ where M_i is the set of probability measures on the Borel subsets of X_i by defining $u_i(\mu) = \int_X u_i(x)d\mu$. The mixed extension of G is denoted by $\bar{G} = (M_i, u_i)_{i=1}^N$.

Definition A.1. *If X_i is a nonempty compact subset of a topological vector space and u_i is bounded for every i, then the game G is called a compact game. If X_i is further a Hausdorff set, then the game G is a compact Hausdorff game.*

Definition A.2. *Player i can secure a payoff of $\alpha \in \mathbb{R}$ at $x \in X$ if there exists $\bar{x}_i \in X_i$, such that $u_i(\bar{x}_i, x'_{-i}) > \alpha$ for all x'_{-i} in some open neighborhood of x_{-i}. A game G is better-reply secure if whenever (x^*, u^*) is in the closure of the set $\{(x, u) \in X \times \mathbb{R}^N | u = u(x)\}$ and x^* is not an equilibrium, some player i can secure a payoff strictly above u^* at x^*. A game G is payoff secure if for every $x \in X$ and every $\epsilon > 0$, each player i can secure a payoff of $u(x) - \epsilon$ at x.*

Definition A.3. *A game G satisfies disjoint payoff matching if for all $x_i \in X_i$, there exists a sequence of deviations $\{x_i^k\} \subset X_i$ such that the following holds:*

1. $\liminf_k u_i(x_i^k, x_{-i}) \geq u_i(x_i, x_{-i})$ *for all* $x_{-i} \in X_{_i}$*;*

2. $\limsup_k D_i(x_i^k) = \emptyset$ *where* $D_i(x_i) = \{x_{-i} \in X_{_i} : u_i$ *is discontinuous in* x_{-i} *at* $(x_i, x_{-i})\}$ *is the set of discontinuities of player* i.

Definition A.4. *A game G is reciprocally upper semi-continuous if whenever (x, u) is in the closure of the set $\{(x, u) \in X \times \mathbb{R}^N | u = u(x)\}$ and $u_i(x) \leq u_i$ for every i, then $u_i(x) = u_i$ for every i.*

Proposition A.1. *Suppose that G is a compact game and G satisfies disjoint payoff matching, then G is payoff secure.*

Proposition A.2. *If G is reciprocally upper semi-continuous and payoff secure, then it is better-reply secure.*

Proposition A.3. *Suppose that G is a compact, Hausdorff game, then G possesses a mixed strategy Nash equilibrium if its mixed extension $\bar{G}$ is better-reply secure.*

With all the definitions and propositions above, we now present the proofs of Theorems 4.1 and 4.2.

Proof of Theorem 4.1. As the strategy space $[0, 1]\times[0, 1]$ is a compact Hausdorff set and the utilities U_B and U_C are bounded, the game G is a compact and Hausdorff game. Denoting the set of probability measures on the Borel subsets of $[0, 1] \times [0, 1]$ by M, then the mixed strategy μ can be defined as $U_g(\mu) = \int_{[0,1]\times[0,1]} U_i d\mu$ where $g \in \{B, C\}$. The mixed extension of the game is referred as $\bar{G} = (M, (U_B, U_C))$, which is also compact.

$U_C(\lambda_B, \lambda_C)$ is continuous on $[0, 1] \times [0, 1]$. On the other hand, $U_B(\lambda_B, \lambda_C)$ has discontinuity at $\lambda_B = \lambda_C$ but $U_B(\lambda_B, \lambda_B) = \max(U_B(\lambda_B^-, \lambda_B), U_B(\lambda_B^+, \lambda_B))$, thus U_B is upper semi-continuous in (λ_B, λ_C) on $[0, 1]\times[0, 1]$, so it is with sum U_B+U_C. Therefore U_B+U_C is also upper semi-continuous in μ on M. So the mixed extension $\bar{G}$ is reciprocally upper semi-continuous.

Since U_C is continuous, to check whether G satisfies disjoint payoff matching, we only need to study the utility of the Type B group. The set of discontinuities of Type B group $D_B(\lambda_B)$ is

$$D(\lambda_B) = \{\lambda_C \in [0, 1] : U_B(\lambda_B, \lambda_C) \text{ is discontinuous in } \lambda_C \text{ at } (\lambda_B, \lambda_C)\},$$

so $D(\lambda_B) = \{\lambda_B\}$. For a given λ_B, without loss of generality, we assume that $U_B(\lambda_B, \lambda_B) = U_B(\lambda_B^-, \lambda_B)$ (We will have a similar proof otherwise). Then clearly $\lambda_B \neq 0$ and we can find an strictly increasing sequence $\lambda_k \uparrow \lambda_B$. So

$$\liminf_k U_B(\lambda_k, \lambda_C) = U_B(\lambda_B, \lambda_C),\ \forall \lambda_C \in [0, 1].$$

Further, we have

$$\limsup_k D_B(\lambda_k) = \limsup_k \{\lambda_k\} = \emptyset,$$

since all the λ_k's are distinct. So the game G satisfies disjoint payoff matching. Hence its mixed extension $\bar{G}$ is payoff secure.

Reciprocal upper semi-continuity and payoff security of $\bar{G}$ ensure that $\bar{G}$ is better-reply secure by A.2. Therefore according to $A.3$, the original game G has a mixed Nash equilibrium. □

Proof of Theorem 4.2. It is sufficient to prove that

$$U_B(\lambda_{min}, \lambda_{min}) = U_B(\lambda_{min}^+, \lambda_{min}),$$

then U_B is upper semi-continuous on $[\lambda_{min}, 1] \times [\lambda_{min}, 1]$. We can use the same method as in Theorem 4.1 to show that the game has mixed Nash equilibrium.

We recall that

$$U_B(\lambda_B^-, \lambda_B) = \lim_{\lambda \uparrow \lambda_B^-} U_B(\lambda, \lambda_B) = \int_{1/(2\lambda_B)}^{\infty} R(r)dr,$$

and

$$U_B(\lambda_B^+, \lambda_B) = \lim_{\lambda \downarrow \lambda_B^+} U_B(\lambda, \lambda_B) = \int_{1/2}^{1/(2\lambda_B)} R(r)dr.$$

Thus the former function is increasing while the latter one is decreasing. Moreover, $U_B(0^-, 0) = 0 < U_B(0^+, 0)$ and $U_B(1^-, 1) > U_B(1^+, 1) = 0$. Thus there exists a $\lambda_* > 0$ such that

$U_B(\lambda_*^-, \lambda_*) = U_B(\lambda_*^+, \lambda_*)$.

Hence for $\lambda_{min} < \lambda_*$, we have $U_B(\lambda_{min}^-, \lambda_{min}) < U_B(\lambda_{min}^+, \lambda_{min})$, then

$$U_B(\lambda_{min}, \lambda_{min}) = \max(U_B(\lambda_{min}^-, \lambda_{min}), U_B(\lambda_{min}^+, \lambda_{min})) = U_B(\lambda_{min}^+, \lambda_{min})$$

which completes the proof. $\square$

REFERENCES

[1] J. Fox and E. Tabory, “Contemporary evidence regarding the impact of state regulation of religion on religious participation and belief”, *Sociology of Religion*, vol. 69, no. 3, pp. 245–271, 2008.

[2] L. R. Iannaccone, “The consequences of religious market structure: Adam smith and the economics of religion”, *Rationality and Society*, vol. 3, no. 2, pp. 156–177, 1991.

[3] R. D. Putnam and D. E. Campbell, *American Grace: How Religion Divides and Unites Us*. Simon and Schuster, 2012.

[4] L. R. Iannaccone, “Sacrifice and stigma: Reducing free-riding in cults, communes, and other collectives”, *Journal of Political Economy*, vol. 100, no. 2, pp. 271–291, 1992.

[5] M. McBride, “Club mormon: Free-riders, monitoring, and exclusion in the lds church”, *Rationality and Society*, vol. 19, no. 4, pp. 395–424, 2007.

[6] L. R. Iannaccone, “Why strict churches are strong”, *American Journal of Sociology*, vol. 99, no. 5, pp. 1180–1211, 1994.

[7] R. Stark and R. Finke, *Acts of Faith: Explaining the Human Side of Religion*. Univ of California Press, 2000.

[8] M. McBride, “Religious pluralism and religious participation: A game theoretic analysis”, *American Journal of Sociology*, vol. 114, no. 1, pp. 77–106, 2008.

[9] ——, “Religious market competition in a richer world”, *Economica*, vol. 77, no. 305, pp. 148–171, 2010.

[10] M. Hout, A. Greeley, and M. J. Wilde, “The demographic imperative in religious change in the united states”, *American Journal of Sociology*, vol. 107, no. 2, pp. 468–500, 2001.

[11] C. P. Scheitle, J. B. Kane, and J. V. Hook, “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*, vol. 50, no. 3, pp. 470–482, 2011.

[12] A. Bisin and T. Verdier, “beyond the melting pot: Cultural transmission, marriage, and the evolution of ethnic and religious traits”, *The Quarterly Journal of Economics*, vol. 115, no. 3, pp. 955–988, 2000.

[13] J.-P. Carvalho, “Veiling”, *The Quarterly Journal of Economics*, vol. 128, no. 1, pp. 337–370, 2012.

[14] M. McBride, “Why churches need free-riders: Religious capital formation and religious group survival”, *Journal of Behavioral and Experimental Economics*, vol. 58, pp. 77–87, 2015.

[15] J. D. Montgomery, “Adverse selection and employment cycles”, *Journal of Labor Economics*, vol. 17, no. 2, pp. 281–297, 1999.

[16] M. D. Makowsky, “A theory of liberal churches”, *Mathematical Social Sciences*, vol. 61, no. 1, pp. 41–51, 2011.

[17] P. P. Barros and N. Garoupa, “An economic theory of church strictness”, *The Economic Journal*, vol. 112, no. 481, pp. 559–576, 2002.

[18] T. Frejka and C. F. Westoff, “Religion, religiousness and fertility in the us and in europe”, *European Journal of Population/Revue européenne de Démographie*, vol. 24, no. 1, pp. 5–31, 2008.

[19] L. R. Iannaccone and M. D. Makowsky, “Accidental atheists? agent-based explanations for the persistence of religious regionalism”, *Journal for the Scientific Study of Religion*, vol. 46, no. 1, pp. 1–16, 2007.

[20] P. Liberati, “The world distribution of income and its inequality, 1970–2009”, *Review of Income and Wealth*, vol. 61, no. 2, pp. 248–273, 2015.

[21] C. E. Lemke and J. T. Howson Jr, “Equilibrium points of bimatrix games”, *Journal of the Society for Industrial and Applied Mathematics*, vol. 12, no. 2, pp. 413–423, 1964.

[22] L. S. Shapley, “A note on the lemke-howson algorithm”, in *Pivoting and Extension*, Springer, 1974, pp. 175–189.

[23] B. Codenotti, S. De Rossi, and M. Pagan, “An experimental analysis of lemke-howson algorithm”, *arXiv preprint arXiv:0811.3247*, 2008.

[24] S. Heinrich Von, *Market structure and equilibrium*, 2011.

[25] P. J. Reny, “On the existence of pure and mixed strategy nash equilibria in discontinuous games”, *Econometrica*, vol. 67, no. 5, pp. 1029–1056, 1999.

[26] B. A. Allison and J. J. Lepore, “Verifying payoff security in the mixed extension of discontinuous games”, *Journal of Economic Theory*, vol. 152, pp. 291–303, 2014.