

*****Early Draft*****

For final published draft, see *Advances in the Economics of Religion*, edited By Jean-Paul Varvalo, Sriya Iyer, and Jared Rubin. Palgrave macmillan, 2019

When Average is Irrelevant: Computational Modeling of Religious Groups

Michael D. Makowsky*

March 12, 2018

*John E. Walker Department of Economics, Clemson University. Email:mdmakow@clemson.edu. I thank Charles North for useful comments and suggestions.

Consider a thought experiment. Two groups are competing for members in a small town religious market. Each of the groups, the Alphas and Omegas, count 100 members on its respective registry, but while the Alphas can boast 90 members in its pews every Sunday, the Omegas must demure to only having 60. Which group is more successful? Which receives more in yearly tithing? Which is more likely to still exist in 20 years?

While the Alphas would seem the immediately more likely pick, let me now present a slightly richer depiction of their registries:

Alphas	Omegas
10 (daily) Zealots	30 Zealots
80 (weekly) Regular Attendees	30 Regular Attendees
10 (yearly) Free Riders	40 Free Riders

As is often the case when the dimensionality of the problem is increased, the solution is far less clear. What is the value of zealotry? What costs (or benefits) do free riders bring to the group? Depending on the type of club good being produced, it may be only the output of zealots that matter (Iannaccone and Berman 2006). Free riders, ostensibly the bane of all club and public goods, may in fact be a net positive in the ledger (McBride 2015). Perhaps regular tithing is only realized from the stable, job holding Regular Attendees, 20% of whom account for 80% of the groups operating budget (Iannaccone 1997). The heterogeneity introduced by this relatively coarse typology has radically increased the predictive power of the model, but is also already beginning to threaten the tractability of the problem. Just as importantly, we haven't actually modeled the emergence of this typology, we've simply assumed it—we've defined agents as regular members and zealots, without assigning a meaningful, quantifiable, metric. Modeling a population of deeply heterogeneous agents offers the possibility of building a theory of behavior beyond the norm, theory predicting both the magnitude of deviation and its representation in broader society. It is the promise of accommodating some additional degree of agent and population

richness, while retaining sufficient theoretic parsimony, which motivates the introduction of computational modeling techniques to the study of religious groups.

In this chapter, I discuss prior use of computational modeling in the economic study of religion, as well as opportunities to extend and elaborate some of the core theoretical models underlying current research. Along the way, we will revisit influential models and discuss the theoretical assumptions, necessary for analytic tractability, that computational methods allow to be relaxed and, in some cases, specifically targeted as explanatory mechanisms for observable phenomenon.

Agent-Based Computational Models

At risk of disagreement over methodological taxonomy, the majority of the contributions discussed here-in can fall under the rubric of *agent-based computational models*. As in virtually any microeconomic model, a set of agents is individually constructed with attributes and objective functions. These agents will typically exist within a defined setting, from as simple as a linear projection or circle, to as complex as a fully-realized replication of a geographic or historical location. The set of agents and their environment will, with rare exception, exist within a dynamic simulation defined by a progression of time, either continuous or discretized into steps. The empirically testable predictions of the model are typically generated via Monte Carlo simulation, exogenously varying parameters while simulating stochastic “runs” of the model, and tracking the individual agent and population outcomes of interest. The nomenclature of “simulation” is appropriate here. The model remains a relatively sparse abstraction, but a run of the model has a distinct beginning and end – the arrival of an “equilibrium” is a possible, but not necessary, outcome. Rather, steady-state conditions, broken ergodicity, random walks, or cyclical dynamics all exist as potential model outcomes.

My colleagues in sociology, anthropology, or evolutionary biology would likely take a different view, but within the context of the economics of religion, it is worth emphasizing that these computational models are not especially different from standard microeconomic models.¹ What will distinguish the computational models discussed (and suggested) here is their abandonment of the representative agent. The unifying, theoretically salient character of these models is their *deep agent heterogeneity*, e.g. fully realized distributions of agent attributes, endowments, strategies, social networks, geographic location, and temporal context. The ability to not just cope with deep agent heterogeneity, but leverage it as a socially generative mechanism, has offered, and will continue to offer, useful and exciting insights into the scientific study of religious groups. Modeling n agents, of $m \leq n$ types, brings with it a new set of issues regarding analytic tractability, particularly when those richly distributed attributes cannot be collapsed to their statistical moments without loss of content. It is within these contexts in which agent-based computational simulation offers a compelling alternative. Given the character of religious groups— their association with charismatic leaders, dedicated zealots, entrepreneurial heretics, and

¹ While computational models in the *economic* study of religion has been relatively sparse, there have been a handful of models constructed within other social scientific disciplines, including sociology, anthropology, and evolutionary biology (Bainbridge 1995; Doran 1998; Upal 2005; Bainbridge 2006; Chattoe 2006; Dow 2008). In fact, *cognitive* models of religious groups and theology, well still early stage theoretical contributions, are perhaps the first to computationally model beliefs, symbolism, and the structure of ritual (Whitehouse 2002; Whitehouse, Kahn et al. 2012). In general, the consideration of beliefs in economic models of religion has been notably sparse (Montgomery 1996).

strictures of often incredible nuance— there can be scarcely few popular phenomena where the average agent is *less* sufficient to generate testable predictions of social outcomes.

As point of emphasis, the ambition here is not to just “weaken” the assumptions of prior theories for the sake of elegance or generalizability. Rather, it is to identify questions that can be newly asked or better answered when a dimension previously assumed away (or collapsed to homogeneity) becomes the subject of interest or explanation. In much of the discussion that follows, we will find that a variable previously homogenous across agents – adherence to parental beliefs, geographic location, network connectivity – becomes the fulcrum by which the model generates new and interesting testable predictions.

The Club Theory of Religion

Iannaccone’s (1992) theory of sacrifice and stigma was a watershed moment for the economic study of religion, serving as the seminal research introducing the “club theory” of religious groups (Iyer 2016). In classic Becker-ian fashion, the club theory of religion applies not just broad economic intuition to the study of religious groups, but also the specific mechanical tools of the economic theorist, e.g. utility maximization, the representative agent, and Nash equilibria.

The Iannaccone (1992) analytic equilibria hinges on a club of homogenous agents making simultaneous decisions based on an assumption of identical Nash Equilibrium behavior. The core insight serves as an invaluable existence proof of unproductive costs as a utility enhancing mechanism for club members, but it comes with a nontrivial blind spot: the corner solutions of the strictness space – requirements of zero or absolute sacrifice of private goods - are the only viable clubs. Any club of intermediate strictness is inferior to those permitting either all secular production or those permitting none.

This “corner solution only” outcome is a product of the (analytically necessary) assumption of homogenous agents.² The means by which moderate strictness clubs might offer a utility maximizing club good for a dominant subset of the population is demonstrated in Makowsky’s (2011) computational cellular automata model. The viability of moderate groups is a product of extending the club theory to accommodate agent heterogeneity, in terms of both income and social network, in both the broader population and within groups. Moderate sacrifice rates sufficiently mitigate member free riding by acting as a weak screening device that filters in a range of agent income types into the group. The resulting within-group heterogeneity in the opportunity costs of tithing and activity participation (i.e. time) allows agents to benefit from the differing comparative advantages of their fellow members. This results in both an interior optimum sacrifice requirement for any one group, and range of simultaneously viable group-types recruiting different member types from the broader population.

² The “corner solution” issue is accommodated in different manners in the club theory of religion literature. In his formal model of church and sect, Iannaccone (1988) assumes that the returns to personal religious conduct always have an interior maximum. Berman (2000), in point of contrast, simplifies the sacrifice to purely screening mechanism (omitting secular good-club, good substitution effects). Subsequent laboratory research found the sacrifice mechanism effective in a purely non-religious context, but also found the screening mechanism dominated (non-trivial) price effects (Aimone, Iannaccone et al. 2013).

Fundamentally Extreme

The dominance of corner solutions in the original club theoretic model is a more subtle concern, but no less costly, for insight into extremist groups. If, within the assumptions of the model, all successful groups impose extreme high sacrifice requirements, then there is no room within the model to gain insight under which groups become relatively stricter, or in which the relatively strictest gain market share in the religious marketplace. Computational models of religious groups, by simulating the behavior of large scale populations of heterogeneous agents can provide insight into the mechanisms generating rich and varied religious markets. New steady-state outcomes and insights are not limited to macroscopic phenomena – new behavioral patterns can be observed within the groups themselves. Within each club there is an ecology of members: a full distribution of zealots, free-riders, and everyone in between. We can observe the market share for extremism within the population, but also within different types of groups— from which types of groups extremists are likely to emerge and the conditions under which groups, subgroups, and individuals become more or less committed.

Identifying religious extremism demands context. Absent reference points, identifying someone or something as extreme is an exercise in normative, frequently pejorative, analysis. A useful model of extremists requires modeling everyone else who is *not extreme*. Makowsky (2012) and Friedman et al (2018) both build club-theoretic agent-based computational models of religious groups, within which populations of richly-attributed religious agents are realized and deeply heterogeneous distributions of religiosity are generated. While their respective models leverage different model attributes to generate their predictions, both models are only able to gain tractability via Monte Carlo simulation of populations of heterogeneous agents.

Both models are illustrative of the importance of deep agent heterogeneity. A standard dynamic stochastic general equilibrium (DSGE) model of heterogeneous agents would collapse the dimensions of agent heterogeneity to some finite number of statistical moments (Krusell and Smith 1998), typically just the mean. This “aggregate approximation” works well if the mean sufficiently captures the impact of the broader distribution on population dynamics (Krusell and Smith 1998; Winberry 2016). The population means (and standard deviations) of income, social networks, and familial religious history, are highly unlikely to be sufficient to capture the population dynamics that underlie behavior and prevalence of religious extremists.³

Collapsing the Makowsky (2012) model to the moments of the income distribution and local religiosity would preclude the conditions under which the most extreme religious groups can demand the highest levels of member sacrifice. Such simplifying assumptions would obscure the prediction generated by the model that extremist groups can be undermined through targeted public good provision. Friedman et al (2018) simulate the emergence of fundamentalist groups in tension with the broader population. They find that outcomes are sensitive to the influence of peer groups relative to individual interpersonal interactions. Approximately aggregated social network distributions would likely preclude the

³ One of the themes in Krueger’s (2007) empirical investigation of terrorism is the insufficiency of average rates of education and income in trying to empirically predict terrorism rates and other national level outcomes.

delineation of the impact of group interaction from individual interactions in a model of dynamically evolving peer groups.

The Church Sect Cycle

In James Montgomery's influential (1996) paper on the dynamics of religious markets, he extends the club theory to incorporate both dynamic entry and exit of groups within a religious market, concurrent with dynamic movement of members across different types of groups. Returning to the thought experiment, his model introduced agents choosing amongst heterogeneous groups whose own institutions of strictness were dynamically serving the median member agent within the group.

The model is notable for several features. First, it employs three income classes of agents (low, medium, and high) to produce simultaneous demand for low, medium, and high strictness religious groups. Religious capital and denominational affiliation is passed on from parent to child. Combined with random assignment of income class, agent denominational strictness may be suboptimal relative to their income. This tension leads to outcomes where populations can secularize, de-secularize, or end up in a cycle of group formation and dissolution analogous to the "Church-Sect Cycle" theory of the American religious market (Finke and Stark 1992)

To see how the Montgomery model and its insights might be extended further in a computational context, we can turn to a different model. Iannaccone and Makowsky (2007) presents an agent-based model of religious geographic regionalism. Their model is not club-theoretic, in the sense of Friedman et al (2018) or Makowsky (2012). Rather, it instantiates m groups to which n agents are assigned initial membership, and parameterizes a simple utility maximizing exercise by agents, each endowed with heterogeneous preferences for: 1) the size of their religious community, 2) continuing with their current religious identity, and 3) continuing in the religious identity of their parents. Their model demonstrates the possibility for steady-state religious regionalism, without tipping towards a uniform population, uniform regions, or pure mixing. It also demonstrates the counterintuitive population *stabilizing* effects of deep agent heterogeneity.

Integrated within the Iannaccone-Makowsky framework, the Montgomery model of the Church-Sect Cycle takes on entirely new context. The addition of a spatial setting, in this case a two-dimensional lattice, would allow for the possibility for groups to "move" across the terrain, following their target populations, or for the types of groups to take root (e.g. the evangelical south and the irreligious west), and wait for their would-be converts to come to them. Perhaps both. A computational setting would also allow for the three income class agent typology to be expanded to a full distribution of income more representative of the United States or another population of interest. The shape parameters, particularly those related to a growing divide between the highest and lowest earners, could generate predictions about future sectarian movements and their evolving bundle of proscriptions. Keeping in mind that tithing tends to follow a classic 80-20 Pareto Rule (Iannaccone 1997), the incentives for religious groups to retain their highest earners are likely to only grow stronger with income inequality, with important implications for churches trying to optimally balance their attractiveness to those most willing to dedicate their time and energy to the group, and those able to donate financial resources.

The Value of Free-Riders

Since economists first began thinking about public goods, it has begged the question of how to solve, or at least cope with, the problem of free riders. While the club theory of religious groups posits unproductive costs as a way of mitigating free-riding, McBride's (2007; 2015) model extends the theory to contexts in which free-riders, within bounds, can prove a net lifetime gain for the club good when group specific capital is acquired through exposure over time. Free-riders are the manner in which a group continues to grow its flock and expand its influence – rather than just a cost to be minimized, free-riders instead become an opportunity to be optimized.

McBride's model, for sake of analytic tractability, simplifies religious capital to two states: low and high. This capital is also entirely group specific – there is no prospect for moving to a lower or higher strictness group and bringing some non-zero religious capital. Ported to a computational setting, the model could be enriched in two interesting dimensions. First, religious capital could be made continuous and heterogeneous across individuals based length of exposure. This would allow the agents to steadily grow from low capital free-riders to valued contributors and community leaders. Second, groups could adjust their practices to increase or decrease the specificity of the capital formed to maximize the balance of retention (increasing the cost of leaving for capital-committed members) and recruitment (reducing the costs of switching for prospective recruits from rival groups). The optimal specificity of religious capital and tolerance of free-riders is likely to vary greatly for new versus established religious groups, and could speak to the theories of optimal “tension” with established churches for cults and other new religious movements (Stark, Bainbridge et al. 1985).

Deep agent heterogeneity, in the context of geography and social networks, would also allow for new dimensions in which the admittance of free-riders can be club-enhancing. McBride (2007) examines the role of welcoming free-riders with low religious capital in Mormon churches affiliated with the Church of Latter-Day Saints. Mormons, along with Jehovah's Witnesses and many evangelical churches, are known for the missionary and door-to-door recruitment. Within most simple cost-benefit models, the success rates of these operations would seem insufficient to justify them in any terms other than additional unproductive costs to reduce member free-riding. In a model of heavily clustered social networks, however, with cross-cluster connections limited by geography, cost, and technology, these largely futile missionary endeavors offer the potential for beachheads in previously untapped religious markets. In a population of agents heterogeneous in income, group-specific religious capital, rates of geographic displacement, and social connectivity, optimal patterns of missionary activity are likely to diverge greatly from simpler models. What might otherwise look like the radically counterproductive efforts of religious zealots might prove instead to be the uncannily sophisticated efforts of an organization whose growth seems unconstrained by political, cultural, or economic boundaries

Religion as a Network Good

While the club theory of religion is a rich source of useful insight and testable predictions, there can be little doubt the content of the “club” good has room for useful enrichment as well. Retaining our focus on purely social dimensions, religion is very clearly a *network good*. Membership provides access to a social network, but any one agent’s place in a network will vary greatly, and will be heterogeneous with regards to the density and importance of their connections. Founders, leaders, and charismatic figures will hold vastly different positions of influence and connectivity within any religious network. Further, they will vary not only in their endowment of connections, but also their ability to form new connections and retain existing ones. The value of these attributes in building and maintaining religious movements will depend on the attributes of the networks they are embedded in, the distribution of members within those networks, and the network-relevant technologies at their disposal.

Extending Rubin (2014) and Kuran (1987), Makowsky and Rubin (2013) embed agents within a social network and spatial landscape, and assess the impact of technologically extended and enhanced social networks on the influence of centralized (i.e. government) and decentralized (i.e. religious) authorities on the public expression of private preferences. The television, radio, and internet are all technologies which are likely to have had important influences on the nature and reach of religious group identity as a network good. As increasingly innovative religious movements seek to expand their membership and their broader social influence via social media, new insights from models of religious markets are likely to be had from redesigning the religious club good as an explicit network good, defined by nodes of heterogeneous influence, connectivity, and endowment.

The Cults are Alright

So much of the magic of religion is lost when the possibility of outsiders, heretics, and outsized characters exists beyond the scope of the theoretic construct. A new religion is, by definition, *a cult*, a term that seems to have carried a pejorative context since its inception. Cults are made feasible by the unmet social and spiritual needs of individuals on the periphery of existing social networks. When cults are successful, it often requires credit to founding individuals with extraordinary personal attributes finding themselves in extraordinary social settings. Regardless of what constitutes the representative agent in a standard economic model, it seems unlikely to represent the prospective members of a newly constituted religious group. Recruitment to groups varies with background religious affiliation, education, and specific social network conditions. Low density networks, broader social isolation, and weak ties to the group in question all correlate to receptiveness to cult recruitment (Stark, Bainbridge et al. 1985). Research on cults, as much as anything, is an opportunity to extend the rigor and testable predictions of economic models of religion to a subject where the barriers to formal modeling are the necessity of incorporating network properties and deep agent heterogeneity. Agent-based computational modeling methods offer the theorist a construct in which the weird, the ostracized, and the maybe-a-little-bit-scary outliers are no longer relegated to the “apologetically assumed away” footnote. Sometimes they’re even the star of the show.

Final Thoughts

Economic studies of non-market behavior often assume that behavior is uniformly rational and self-interested, while differences in behavior are a product of variations in exogenously determined endowments and preferences. Attention is seldom paid to how differences in preferences came to be. Computational models can attempt to grow, or evolutionarily select-for, those preferences. Models of *endogenous* religious preferences, as products of historic (even pre-historic) agent networks, geographic limitations, and resource distributions, are a natural next step for the social scientific study of religion. Religious preferences constitute too much of the story to remain relegated to the realm of exogenously assigned parameters whose only job is to facilitate mathematical elegance and expository simplicity. Given standard disciplinary modeling norms, it is not surprising that some of the earliest computational models of emergent belief and ritual structures have come from outside economics and mathematical sociology (Whitehouse 2002; Whitehouse, Kahn et al. 2012).

As with any technique, the time and need for methodological advocacy has quickly passed for agent-based computational modeling. The economic study of religion and religious groups is a perfect example of a field of study where computational models can explicitly build from the existing body of theory, opening up models along dimensions previously intractable, and contribute new avenues of theoretical insight and testable predictions.

References

- Aimone, J. A., L. R. Iannaccone, M. D. Makowsky and J. Rubin (2013). "Endogenous Group Formation and Unproductive Costs." Review of Economic Studies **80**(4): 1215-1236.
- Bainbridge, W. S. (1995). "Neural network models of religious belief." Sociological perspectives **38**(4): 483-495.
- Bainbridge, W. S. (2006). God from the machine: Artificial intelligence models of religious cognition, Rowman Altamira.
- Berman, E. (2000). "Sect, Subsidy, and Sacrifice: An Economist's View of Ultra-Orthodox Jews." Quarterly Journal of Economics **65**(3): 905-953.
- Chattoe, E. (2006). "Using simulation to develop testable functionalist explanations: a case study of church survival." The British journal of sociology **57**(3): 379-397.
- Doran, J. (1998). "Simulating collective misbelief." Journal of Artificial Societies and Social Simulation **1**(1).
- Dow, J. (2008). "Is religion an evolutionary adaptation?" Journal of Artificial Societies and Social Simulation **11**(2): 2.
- Finke, R. and R. Stark (1992). The Churching of America, 1776-1990: Winners and Losers in Our Religious Economy. New Brunswick, NJ, Rutgers University Press.
- Friedman, D., J. Fan, J. Gair, S. Iyer, B. Redlicki and C. Velu (2018). "How Fundamentalism Takes Root: A Simulation Study."
- Iannaccone, L. and E. Berman (2006). "Religious Extremism: The Good, the Bad, and the Deadly." Public Choice **128**(1-2): 109-129.
- Iannaccone, L. and M. Makowsky (2007). "Accidental Atheists? Agent-Based Explanations for the Persistence of Religious Regionalism." Journal for the Scientific of Study of Religion **46**(1): 1-19.

- Iannaccone, L. R. (1988). "A Formal Model of Church and Sect." American Journal of Sociology **94**(Supplement: Organizations and Institutions: Sociological and Economic Approaches to the Analysis of Social Structure): S241-S268.
- Iannaccone, L. R. (1992). "Sacrifice and Stigma: Reducing Free-riding in Cults, Communes, and Other Collectives." Journal of Political Economy **100**(2): 271-291.
- Iannaccone, L. R. (1997). "Skewness Explained: A Rational Choice Model of Religious Giving." Journal for the Scientific Study of Religion **36**(2): 141-157.
- Iyer, S. (2016). "The new economics of religion."
- Krueger, A. B. (2007). What Makes a Terrorist: Economics and the Roots of Terrorism, Princeton University Press.
- Krusell, P. and J. Smith, Anthony A (1998). "Income and wealth heterogeneity in the macroeconomy." Journal of political Economy **106**(5): 867-896.
- Kuran, T. (1987). "Preference Falsification, Policy Continuity and Collective Conservatism." Economic Journal **97**: 642-665.
- Makowsky, M. D. (2011). "A Theory of Liberal Churches." Mathematical Social Science **61**(1): 74-87.
- Makowsky, M. D. (2012). "Emergent Extremism In A Multi-Agent Model Of Religious Clubs." Economic Inquiry **50**(2): 327-347.
- Makowsky, M. D. and J. Rubin (2013). "An Agent-Based Model of Centralized Institutions, Social Network Technology, and Revolution." PLoS One **8**(11).
- McBride, M. (2007). "Club Mormon: Free-riders, Monitoring, and Exclusion in the LDS Church." Rationality and Society **19** (4): 395-424.
- McBride, M. (2015). "Why churches need free-riders: Religious capital formation and religious group survival." Journal of Behavioral and Experimental Economics **58**: 77-87.
- Montgomery, J. D. (1996). "Contemplations on the economic approach to religious behavior." The American Economic Review **86**(2): 443-447.
- Montgomery, J. D. (1996). "The Dynamics of the Religious Economy: Exit, Voice, and Denominational Secularization." Rationality and Society **8**(1): 81-110.
- Rubin, J. (2014). "Centralized institutions and cascades." Journal of Comparative Economics **42**(2): 340-357.
- Stark, R., W. S. Bainbridge and The Seymour Martin Lipset collection. (1985). The future of religion : secularization, revival, and cult formation. Berkeley, University of California Press.
- Upal, M. A. (2005). "Simulating the emergence of new religious movements." Journal of Artificial Societies and Social Simulation **8**(1).
- Whitehouse, H. (2002). "Modes of religiosity: Towards a cognitive explanation of the sociopolitical dynamics of religion." Method & Theory in the Study of Religion **14**(3): 293-315.
- Whitehouse, H., K. Kahn, M. E. Hochberg and J. J. Bryson (2012). "The role for simulations in theory construction for the social sciences: case studies concerning divergent modes of religiosity." Religion, brain & behavior **2**(3): 182-201.
- Winberry, T. (2016). "A Toolbox for Solving and Estimating Heterogeneous Agent Macro Models." Quantitative Economics. Work in progress.