Accidental Atheists?
Agent-Based Explanations for the Persistence of Religious Regionalism*

Laurence R. Iannaccone
Department of Economics
George Mason University

Michael D. Makowsky
Department of Economics
George Mason University

October 2004

Abstract

This paper presents a multi-agent simulation of religious regionalism. We contrast the multi-agent approach to traditional regression analysis, showing that the former better captures the actual process by which people make religious choices and better explains the persistence of distinctive regions in a mobile and pluralistic society. Starting with an inverted form of Schelling's classic model of mobility and segregation, we show that the (MARS) multi-agent religion simulation readily captures a wide range of insights from standard theories of rational choice, social influence, and preference formation.

JEL Codes: Z12, C6

Keywords: economics of religion, religious regionalism, micro-macro, emergent phenomena, social simulation, multi-agent, agent-based.

* Laurence R. Iannaccone, Department of Economics, George Mason University, Fairfax, VA 22030, Larry@EconZone.com. Michael Makowsky, Department of Economics, George Mason University, Fairfax, VA 22030, mmakowsky@gmu.edu. Downloads: www.EconomicsofReligion.com. We thank Andrew Iannaccone and Matthew Zilli for contributions to early versions of the MARS program.

Accidental Atheists?
Agent-Based Explanations for the Persistence of Religious Regionalism

Imagine the following magic trick: A statistician grabs a pair of urns, one with red balls and one with green. He shakes the first urn and pours some of its red balls into the second urn, then shakes the second urn and pours some of its now-mixed red and green balls back into the first, and continues this process, back and forth, ten or twenty times. Finally, he tips over both urns, showing that the first still contains only red balls and the second only green. He then repeats the shaking and pouring procedure starting with a 80-20 mix of red-to-green balls in the first urn and a 20-80 mix in the second, but no matter when we check the urns the color ratios remain unchanged.

Short of special props or slight-of-hand, a statistician really would need magic to pull this off. Yet a similar religious result rarely raises eyebrows. Visualize the American population as three-hundred million "balls" – red if they regularly attend church and green otherwise. The balls are distributed unevenly across regional "urns" – the proportion of religious "reds" is relatively high in the Southeast, whereas the proportion of non-religious "greens" relatively high in the Pacific West. Each year, millions of balls move from urn to urn.[1] And yet decade after decade the regions maintain their distinctive ratios. Despite massive nonstop mixing, the South remains relatively religious, the West relatively irreligious, and so forth.

Researchers have studied mobility and regionalism using numerous data sets and increasingly sophisticated statistics. Unfortunately, their work does more to confirm the magic than explain it. Smith et al. (1998) conclude that "[s]omething about merely 'being' in the South … leads one to have a stronger religious commitment [and] detectable religious homogenization has still not occurred." Welch (1983) likewise warns that his "results must be interpreted modestly." His regressions "explain" a substantial portion of Western irreligiosity in terms of its relatively transient population and other demographic characteristics, but "[e]ven after statistical adjustment, the Western church membership trough is still quite marked." Moreover, the relatively "unchurched" character of the American West dates back to the mid-1800's at least (Finke 1989).

Multi-agent methods cast new light on the puzzle of religious regionalism. Relative to regression and most other standard multivariate statistical methods, M-A does a far better job showing how macro-level outcomes emerge from micro-level forces and actions. In particular, M-A traces the "magic" of persistent regionalism back to the combined impact of personal experience and social conformity. If newcomers routinely adapt to their new social environment, then a region can remain "green" despite a large influx of "reds." As newcomers change their colors, the West acquires "accidental" atheists and the South gains "accidental" enthusiasts. As we shall demonstrate, however, getting from this simple story to a working theory is much harder than it seems – in part because social

forces strong enough to capture newcomers tend also to tip entire regions toward a single shared "color."

Multi-agent methods have many strengths in addition to their capacity to derive complex social patterns from simple individual actions. They are, however, relatively new to the social sciences and almost entirely new to the study of religion[2]. The puzzle of religious regionalism provides an excellent opportunity for M-A models to have an impact. We have designed this paper so readers can verify its claims using an online version of our Multi-Agent Religion Simulation. Our "MARS" framework also applies to other generalizations about social dynamics and religious outcomes – including commitment, conversion, conformity, social networks, social capital, geographical mobility, religious regionalism, and denominational growth.[3]

The Paradox of Persistent Regionalism

The paradox of persistent mobility and persistent regionalism defies easy resolution. America is a nation of high mobility – both geographic and social – and mobility does induce homogeneity *ceteris paribus*. The apparent implication is so obvious that scholars of the 1950s and 60s viewed religious convergence as a done deal. Will Herberg (1960) aptly summarized the scholarly consensus in his widely read book, *Protestant, Catholic, Jew* – which described America as a "triple melting pot" where regional, ethnic, and sectarian differences had become "distinctly secondary," and the three major religious traditions had themselves converged onto a shared "American way of life." Recall as well that the late-50s and early-60s were the heyday the ecumenical movement, mainline denominational mergers, and secularization theory.

By the late sixties, however, scholars could no longer ignore the evidence for persistent regional cultures, especially in the South (Glenn and Simmons 1967; Halvorson, Newman, and Nielsen 1978; Reed 1972); and since that numerous studies have confirmed the remarkable stability of America's religious landscape. (See, for example, Finke 1989; Newman and Halvorson; Smith, Sikkink, and Bailey 1998; Stump 1984a; Wuthnow and Christiano 1979). Despite these and many other studies of both migration and regionalism, to say nothing of vastly larger literatures on regionalism alone or migration alone, the puzzle persists.[4] The standard analytical methods hit a wall long ago. Little has changed since the early eighties, when Welch (1983) complained that "explanations of regional distinctiveness tend to be post-hoc or excessively historicist" and "well-specified models" are "relatively rare." In fact, "rare" is too generous; "non-existent" comes closer to the truth.

To make progress, we must overcome two problems. First, we must find a way to model the joint and interacting impact of multiple factors, all of which influence a person's religious beliefs and behavior. Prior research demonstrates that any realistic list must include: social ties, denominational affiliation; prior religious experience, and personal demographic attributes. Second, we must somehow model a complete social *landscape* comprising individual actors (each with their own history, location, and personal attributes) whose interactions over space and time shape the landscape and are in turn

shaped by it. It is here that we must address the sociologist's critique, oft cited as a fundamental flaw of economics, and "socialize" the model, linking micro and macro levels, by endogenizing network effects.

In sociology of religion, the standard response to such problems is multiple regression and other multivariate statistical methods. Armed with a large national survey, or better yet a national longitudinal survey, the researcher statistically models the religiosity of individuals in terms of their age, gender, race, and so forth, religious background, migration history, current location, and social network. After sifting through as many effects as possible, the researcher then focuses on those that prove "statistically significant" or relatively good at "predicting" variation in the data. For well-executed examples of this approach see Stump (1984b; 1986) or Smith (1998). A typical result has the form,

$$R_i^t = a + b \times (\text{demographi cs}) + c \times (\text{denominati on}) + d \times (\text{migration status}) + e \times (\text{region})$$

where R_i^t denotes the (surveyed) religiosity of individual i at time t. The demographic variables usually include age, gender, race, education, and marital status; denominational affiliation should at least distinguish among conservative Protestants, liberal Protestants, and Catholics. Migration status captures the respondent's migration history, and region variables indicate where the respondent currently resides.

Estimation procedures along these lines produce fairly consistent results. The region coefficients associated with current location invariably find that the South remains relatively religious, the West relatively irreligious, the Midwest and Northeast somewhere in between, and Utah exceptionally religious. Residential mobility tends to reduce religious activity, *ceteris paribus*, but in practice this effect is complicated by "regional adaptation" – whereby religiosity tends to rise among people who migrate from less religious areas (such as the West) to more religious areas (such as the South) and tends to fall when moving in the opposite direction. Finally, the regressions almost always yield standard denominational and demographic effects – including positive, statistically significant effects associated with being older, female, African-American, married, and affiliated with an evangelical/conservative Protestant denomination.

Methodological Problems

The statistical pitfalls associated with regression and other multivariate statistical methods are by now well known. But in this case the abuse of "statistical significance," "goodness of fit," and other problematic assumptions are of secondary concern.[5] For the sake of argument, let us assume that our data and statistical methods are not tainted by *any* standard statistical flaws. Even then, the standard published results tell us almost nothing about the social dynamics that give rise to persistent religious regionalism.

Consider first the regression coefficients for "region" effects. These simply confirm what we already knew – that large and stable regional differences cannot be explained away by

individual-level factors alone. Consider next the "mobility" effects. Residential mobility tends by itself to reduce religious involvement, but the net effect of a move depends also on whether the individual enters a region of higher or lower religiosity. At best, these results suggest that individual religiosity depends on social ties – again something already known.[6] As for the demographic and denominational effects, these are absolutely standard statistical results, consistent with hundreds of other survey-based studies.

Might we do better with more direct measures of social characteristics? Consider the effect of augmenting our regressions with measures of regional religiosity such as the average *R* in the respondent's current region. The estimated equation now becomes $R_i^t = a + b(\text{demographi cs}) + ... + e(\text{region}) + f(\overline{R}_i^t)$, and for the sake of argument let us assume that all other "region" effects now drop to zero. This result might well find its way into the a major journal inasmuch as it fully "explains" regional differences in terms of social effects, plus mobility, demographics, and so forth.

But consider the matter more closely. Our improved model has introduced a potentially serious specification problem, because the right-hand-side averages, $\overline{R}_i^t$, implicitly contain within them the left-hand-side levels, R_i^t. The average-R's will tend therefore to be influenced by the same unobserved "error" terms that influence individual-R's, thereby violating the OLS independence assumption so as to bias estimated coefficients, significance levels, goodness of fit, and so forth. If this seems like a minor technicality, note that the same statistical problem seriously biased numerous regression results published throughout the 1990s on religious pluralism versus religious adherence (Voas, Olson, and Crockett 2002). See Moutlon (1990) for an analysis of the underlying statistical issues.

Ironically, larger samples merely exacerbate this problem. As the sample approaches the entire population, the average-R values exactly equal average of the corresponding individual-R. Better measures of social influence can likewise make matters worse. When average-R for an entire region is replaced by average-R in the respondent's neighborhood, let alone the religiosity of the respondent's closest friends, then statistical bias becomes larger still. (The smaller the social circle, the more likely it is that any unobserved "error" that influences average religiosity will tend also to directly shift the individual's religiosity. The R_i^t error term will be correlated with $\overline{R}_i^t$, thus overstating the size and significance of the average-R social effect and possibly invalidating all other equation statistics as well.) Something is seriously wrong when methods of inquiry become less valid as the data become more complete and detailed.[7, 8]

The fundamental problem is really not statistical at all, though it most certainly leads to statistical hazards. The problem is that we lack a coherent model linking individual behavior to aggregate characteristics and vice-versa. Armed with such a model, even one that is stylized or simplistic, we could generate explanations, predictions, and meaningful statistics. Above all, we could explore the micro-macro links to see if persistent

regionalism does indeed emerge from "adaptation" and other relatively simple micro-level tendencies.

Regression and its more sophisticated cousins will never get us where we need to go because *they bypass the very links we seek to address*. The macro-micro link is reduced to regional dummy variables (or regional averages in our more sophisticated variant). This contradicts our own assumptions about how social-influence operates. Countless careful studies have already shown that religious social-influence operates primarily through close social ties which, in turn, depend on individual social networks.[9] The fundamental feature of social influence, documented in scores of case studies since the 1960s, is how much it varies over time and across people. People are heavily influenced by their closest friends, but scarcely influenced at all by neighbors they hardly know, much less the mass of strangers who make up 99.9 percent of the population in their city, state, or region.

Our statistical models reduce social influence to a homogeneous *regional* attribute (captured by single regional average). Averages and dummy variables ignore social structure; they collapse relationships and local networks down to a shared set of "social" parameters. In essence, each individual becomes an atomized particle, floating free in a homogenous soup of other such particles. Nothing really happens at the social level. Regional effects enter the regressions as standard exogenous variables with no local content and no reciprocal effects. In this last respect they behave like a static equilibrium price of a fully equilibrated competitive market.[10]

Ironically, failure to appreciate social structure is precisely the sin that sociologists most frequently (and correctly) attribute to economists. In an influential article on social "embeddedness", Mark Granovetter (1985) observes that "[c]lassical and neoclassical economics operates …

> with an atomized, undersocialized conception of human action [that disallows] … any impact of social structure and social relations [483] Even when economists do take social relationships seriously … they invariably abstract away from the history of relations and their position with respect to other relations … The interpersonal ties described in their arguments are extremely stylized, average, "typical" – devoid of specific content, history, or structural location. [486] … Social influence here is an external force … ongoing social relations and structures are irrelevant. [486]11

These are precisely the problems that plague the statistical models of migration and regionalism. Within those models, social influence and interpersonal ties are indeed reduced to "stylized, average" parameters "devoid of content, history of structural location." More appropriate models – whether in economics or sociology of religion – must recognize that

> culture is not a once-for-all influence but an ongoing process, continuously constructed and reconstructed during interaction. It not only shapes its members but also is shaped by them, in part for their own strategic reasons. [486]

The Multi-Agent Alternative

When used by social scientists, regression and related methods typically operate as a form of simulation. Researchers start with a collection of numbers and a pre-selected family of equations that link a "dependent" number Y to various "independent" numbers X. They then apply OLS or other mechanical methods to calculate specific parameter values for their simulation. If a great many assumptions about the data and the actual relationship Y and X happen to be correct, the simulating equations generate meaningful outcomes, which can be summarized in terms of "predicted values," "confidence intervals," and other statistics. As we have seen, however, these simulations fail in the present context, and they almost certainly fail in many other contexts important to the sociology of religion.

Multi-agent models[12] provide an alternative method of simulation – one that is well-adapted to studies of social environments populated by numerous actors, each of whom operates within a small "neighborhood." This makes it easier to study religion, migration, and social influence while not ignoring the lessons of sociology[13].

The Schelling Model: Thomas Schelling (1971) developed the first multi-agent model as a way of understanding residential segregation. To this day, it remains the best known and most informative example of the MA approach to social theory.[14] In Schelling's simulation, different types of people are represented by disks of different colors, potential housing locations are represented by the squares of a large grid, and each disk "lives" on a separate square. Each person likes living in a location where at least one-third of his "neighbors" (the disks adjacent to him) share his color. The simulation starts with people randomly distributed over the entire grid, filling most but not all of the squares. (Visualize the setup as an oversized chessboard in which three-quarters of the squares are occupied by disks, half red and half green, with no apparent pattern of spacing or color.) To "run" the simulation, we simply select a disk at random, examine the squares adjacent to it, determine whether this "person" is content with his immediate neighborhood, and if not randomly move him to any open location that has sufficiently many same-color neighbors. We then repeat the process, randomly selecting another disk, and so forth.

Schelling's simulation would inevitably lead to total segregation if everyone insisted on having a *majority* of neighbors share their color. But no such restriction applies when everyone displays only a mild, one-in-three preference for his own color. Nevertheless, it turns out even mild preferences almost always yield very high levels of segregation. This is not a "trick." In the simulation – and very probably in real life – residential segregation truly is an "emergent" characteristic of social systems. Though *caused* by micro-level behavior, it is not necessarily *desired* or reflective of micro-level preferences. Figure 1 shows a typical Schelling simulation, starting with the initial (and dispersed)

random placement of a 50-50 population in panel 1a and stabilizing around the heavily clustered (and effectively segregated) placement in panel 1b.[15]

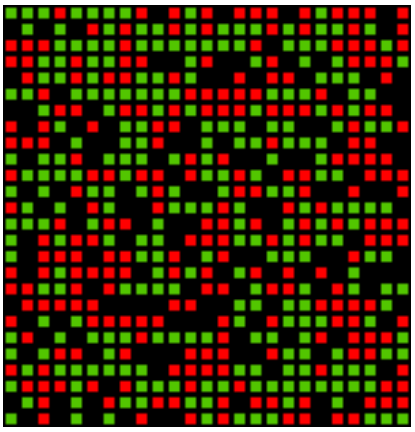

Figure 1a: Schelling Start

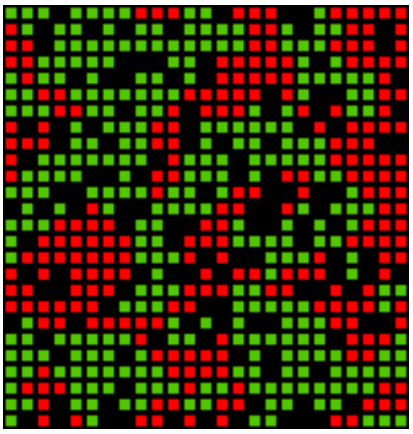

Figure 1b: Schelling End

The MARS model: The simplest variant of our proposed Multi-Agent Religion Simulation takes the Schelling's model and turns it on its head. In Schelling's simulated world individuals change locations based on their unchangeable personal characteristics. In the world of MARS, movement is *not* a matter of choice, but personal characteristics are. Upon being randomly moved to a new location, an individual has the option of changing his religious characteristics to better fit his new neighborhood. (In the context of religion, random movement appears to be the appropriate assumption. Unlike wages and housing costs, the religious character of a region seems to play scarcely any role in determining whether and where most Americans move. Moreover, very few Americans appear to know much about the religious character of different cities.[16] [c.f. Tooney].)

For a flavor of the simulation in action, consider figures 2a through 2d. In each panel, the "board" is now divided into two separate regions with reds initially dominating in the top region and greens dominating in the bottom. To begin running the simulation, we randomly select an individual from anywhere on the board, move him to an open location, also selected at random, and then let this individual "choose" whether to maintain or change his religious attributes (that is, his color) based on his own characteristics and the religious attributes of his new neighbors. We then repeat this procedure for the next randomly selected person, and so forth. As in the original Schelling simulation, the results will depend on the choice rules employed by the individuals.

At one extreme, we can imagine a world in which the influence of childhood socialization or simple inertia is so great that people always maintain their initial religious attributes no matter what their new neighbors are like. In this case, mobility quickly leads to homogeneity across regions – the "melting pot" result anticipated by most scholars of the 1950s and early-1960s. Figure 2b depicts the result just 100 moves after the initial position in panel 2a. In equilibrium, the red-green ratio in each region must equal the red-green ratio in the overall population.[17]

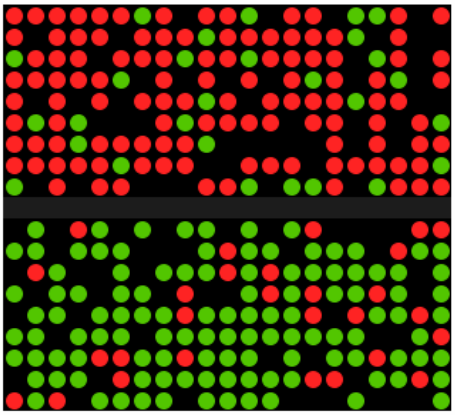

Figure 2a: Simple MARS (initial state)

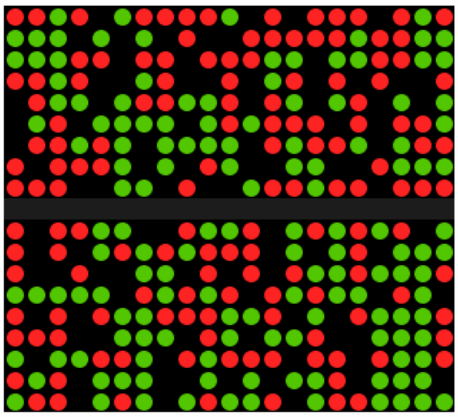

Figure 2b: Simple MARS (100 moves, without adaptation)

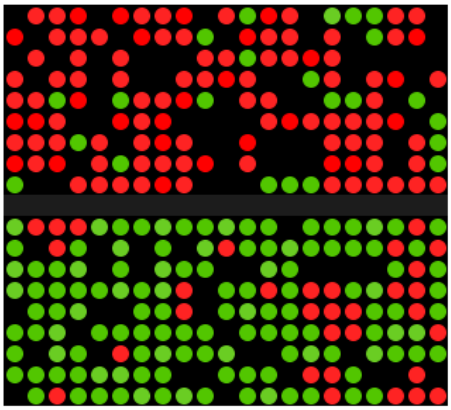

Figure 2c: Simple MARS (100 moves, with adaptation)

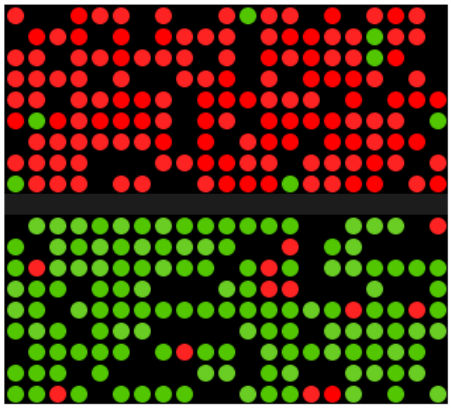

Figure 2d: Simple MARS (200 moves, with adaptation)

Consider next the opposite extreme: a simulation in which movers always adopt the majority religion of their new neighborhoods. Panel 2c displays the result of this alternative regime, starting at the same initial position (2a) as before and continuing again for 100 moves. Thanks to localized social pressures, the two major regions have become *less* similar, not more, and color-clustering has increased within most sub-regions. In this regime of pure social pressure the equilibrium result is just one color in each region. Although we cannot in general predict which color will dominate in any given region, we can predict that once any particular color takes over a region it will remain dominant thereafter (see figure 2d, which follows 2c after 100 more moves).

In short, social adaptation can cause regional differences to persist indefinitely in the face of migration that would otherwise induce homogeneity. In this sense, the simulation validates the explanation for persistent religious regionalism offered by Stump {, 1984 #3055}. But the simulations also highlight problem with prior explanations. Whereas pure inertia yields fails to account for persistent regional differences, pure adaptation

make the differences too stark, because each region "tips" into just one religion. Figures 2c and 2d illustrated this phenomenon, and the degree of tipping would be far more noticeable had we begun with less lopsided ratios of red versus green agents. Something more is needed to explain the observed macro-level combination of *both* persistent differences *and* persistent diversity. Likewise, something more is needed to bring the simulation's micro-level behavior into line with that of real-world movers (who do *not* always alter their religious behavior to fit their new neighborhood).

A successful simulation should yield macro-level results consistent with observed outcomes; and a credible simulation must incorporate micro-level assumptions consistent with observed behavior. Fruitful simulations also facilitate "what if" analysis and sensitivity checks by making it easy to incorporate and manipulate a wide range of effects. With these goals in mind, we turn to a more detailed description of MARS.

> [For an indication of what *might* be possible in a richer simulation framework, consider the two-stage account (inspired both by figures 1 and 2 and by actual American history). In stage one, large and relatively unpopulated areas grow rapidly due to the immigration of many people from just a few places. As these *groups* of immigrants arrive, they naturally settle with others of their type because they differ so radically from the members of other groups (in language, custom, and culture) and internal mobility in this stage tends to maintain or even increase the Schelling-like tendency to cluster. Soon, the area looks like figure 1b, with distinct regions dominated by different types of people. But when migration slows a different dynamic tends to take over. The children (or grandchildren) of immigrants adopt the national language and most other aspects of national culture, and so that restricting one's moves in Schelling-like fashion becomes less beneficial. Indeed, the optimal strategy is to move wherever (economic) opportunity leads and adapt to the new environment. In practice, however, this means adapting to the *religious* characteristics of regions established in stage one. So the initial religious patterns persist even after the dominant *social* dynamic reverses.]

Details: Life on MARS

As McCloskey (1998) and others have emphasized, even the most austere mathematical model is in essence, a story. M-A models are no exception, and within that genre MARS reads like a typical tale. MARS inhabitants move from time to time and do so without regard to religion. But upon arrival in their new neighborhoods, people reassesses their religious choices (because even the choice *not* to change means finding a new congregation) and in doing so they choose with regard to both their own attributes and their new environment. In the spirit of Einstein's dictum that "things should be made as simple as possible, but not any simpler," MARS captures only the most important features of region, religion, and migration. (These include: *diversity, disequilibrium, mobility, randomness, individuality, adaptation, localism, quasi-symmetry, and denominationalism*.[18]) The simulation occurs on a square "board" with divided into square "patches" that can be inhabited by individual "agents." As in other M-A models, the simulation is governed by attributes and rules. Each agent has various attributes including his current location and current and initial type of religion, which the simulation encodes as a list of parameter values associated with each agent. The most

important attributes – *location*, *type* of religion, and *rate* of religious involvement – are visually represented by screen location, color, and shade. Rules govern how agents act and interact, and these are coded as sub-programs and procedures that (randomly) move agents, determine how agents respond to their location and neighborhood.

Agent Movement and Choices: In the model's more simple variants movement is entirely random – agents exercise no control over when or where they move they move – and only movers consider changing their religious attributes (such as type of religion and rate of attendance). Agents choose based on their *original* attributes (determined at "birth"), their *current* attributes (just before moving), and *new* environment (as determined by the location and agents of their new neighborhood). Although this process is mathematically equivalent to maximizing an "objective function", agent "choice" is by no means restricted to rational choice or utility maximization in the standard economic sense.

Choice of Type: Agents here take on a traditional rational choice nature, executing a maximization function based on their own attributes and the attributes of their neighboring agents. However, in contrast to general equilibrium theorizing, an agent's rationality is bounded by his having only local knowledge (who his neighbors are) and no predictive power regarding the future. As such an agent may be perfectly rational in choosing a minority religion bordering on social extinction, so long as his neighborhood over-represented the religious type in question.

An agent evaluates his choice of religion each "move". This move often involves a physical change in location, depending on the movement rate parameter[19] Upon completing a move, an agent always evaluates his utility in each possible type-state, and then chooses the type with the highest utility.

Formal description:

Let ULjk(i) denote the utility that an agent will derive from choosing to become (or remain) type i, given that he was originally of type j, was previously of type k (in the last period, T-1), and is now located at patch L surrounded by NLi neighbors of type i. MARS operates with Moore neighborhoods, which entails a maximum number of neighbors, NL, of eight or less depending on how many border patches are adjacent to L. DL(i) is defined as the ratio [NL(i)/NL]. It can range from zero to one. The symbol "δij" is used in mathematical expression to switch terms "on" or "off" depending on whether i and j are equal or different. By definition, δ equals one if i = j and zero if $i \neq j$.

The utility of the function which agents attempt to maximize is as follows:

ULjk(i) = U(i, j, k, DL(i))

= (δij * [j_origin]) + (δik * [k_inertia]) + DL(i) * [i_community],

where [j_ origin], [k_ inertia], [i_ community] are parameter values[20] set by sliders which act as weightings in the agent's utility function.

Notes on parameter values:

[*j_ origin*] equals the utility associated with choosing to maintain (or return to) one's type-of-origin. If zero, then the agent displays no attachment to his religious upbringing.

[*k_ inertia*] equals the utility associated with *not* changing one's current type. If zero, then the agent displays no internalized "inertia" across periods.

[i_ *community*] equals the "social" utility that the agent receives from neighbors who share his type. If zero, then neighbors of type-i exert no "pull" on their co-religionists. This effect is weighted by the number of type-i neighbors *relative* to the number of potential (not actual) neighbors[21].

Choice of Rate: An Agent chooses his rate of attendance, a standard measure of religiosity, in a manner strongly analogous to his choice in religion "type." The parameters used in weighting the formulation are directly analogous to those used in the type formulation.

Formal description:

Let ULpq(r) denote the utility that an agent will derive from choosing to attend at rate r, given that he was originally attending at rate p, was previously attending at rate q (in the last period, T-1), and is now located at patch L surrounded by NLi neighbors attending at rate r.

$U^{L}_{pq}(r)$ = U(r, p, q, $D^{L}(r)$)

= (δrp * [p_ *upbringing*]) + (δrq * [q_ *habit*]) + DL(ri) * [r_ *flock*]

Upbringing, Habit, and Flock are analogous in effect to the Origin, Inertia, and Community used in the previous utility function. (Note: Although the preceding description, separates choice regarding type and rate of religion more sophisticated versions allow the two to interact.)

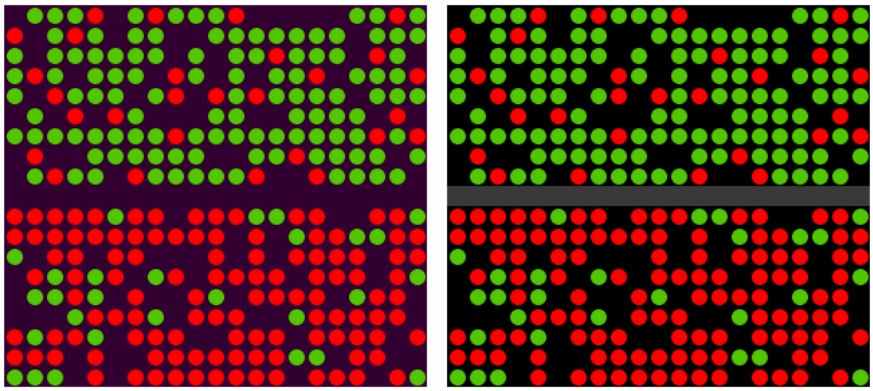

Figure 3: Simulation Board at the Start of a Run

The board: The simulation employs two panels. The right panel tracks the "actual" state of the world as agents move and make religious choices as described above. The left panel displays the world as it would be if agents moved but never changed their religious attributes. Hence the left panel provides a visible baseline or "control" against which to assess the impact of the choice rules employed in any particular simulation. As in figure 3, which displays both panels at the start of a run, both panels are always identical at the start. In this particular run, both boards are subdivided into two regions, upper and lower,

divided by a barrier strip. In setting up any given run, the user has the option of having the board divided into one, two, three, or four regions.[22] For each region, the user also decides the population *density* (the percentage of patches occupied by agents) and the shares on agent types (the relative number of agents of each color). The user also decides whether to run the simulation with one, two, three, or four different types of agents (colored red, green, blue, and yellow). Based on these and other user selections, the program then randomly allocates agents to each region. Figure 3 displays a two-region setup in which green agents predominate (80-20) in the upper region, red agents predominate (20-80) in the lower region, and the population densities are 75% in each region. At this point, of course, both the "actual" and "control" panels are identical.

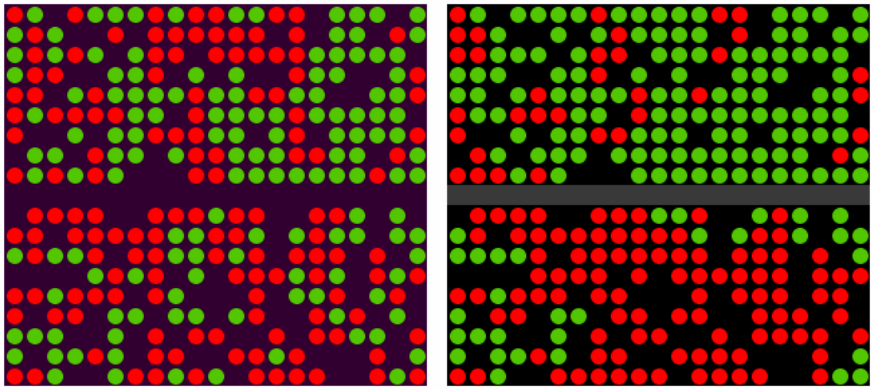

Figure 4: Simulation after 600 moves

Figure 4 shows the state of the world after 600 moves – enough "time" for the typical agent to have moved twice, and hence a period analogous to a full generation in contemporary America. We selected parameter values governing conversion that caused each mover to convert only if a substantial majority of its new neighbors were of another type – so regions do not readily "tip" to just one color. Moreover, the values were symmetric, so that neither color enjoyed an intrinsic advantage over the other in attracting new converts or holding onto old members. The net effect in this run is that regional characteristics persist quite strongly – compare the right panels in figures 3 and 4. In contrast, the left-side "control" panels show that in the absence of neighborhood effects, there would have been substantial convergence across the two regions, leading each region to become fairly homogeneous with a roughly equal mix of both types. Of course, this is just one particular run for one particular set of initial parameters. But in other runs starting from the same initial position and parameters, as well as other runs starting from different initial positions, or somewhat different initial parameters the same basic effect remains. Moderate levels of social adaptation lead to persistent religious regionalism – at least over simulated time spans that approximate two or three decades. Not surprisingly, this result is less robust over longer periods of time and broader ranges of parameter values. But that fact is, in itself, important and it is just one of the many issues that we can address by repeatedly running simulations over a range of different initial conditions.

"Tipping," Corner-Solutions, and other Parameter-Sensitive Results: *[this sub-section requires some additional work]*

A phenomenon of population "tipping" is quickly made apparent when analyzing regionalism, something we will discuss more fully in our conclusions. As an additional thought experiment, this tipping phenomenon was analyzed in a world without regions, made up equal numbers of agents of two different types, each with identical parameters. With 100% mobility and 2000 movement turns, the experiment was run 25 times. The results are summarized in table 1.

Runs	Starting share	Mean Finishing Share	Median Finishing Share	Leader with >50%	Leader with >40%	Leader with >33%
25	25% (4 types)	25%	19%	36%	64%	96%

In testing a variety of parameters and conditions, there arose a consistent theme of greater initial diversity making long-term corner solutions more likely in the market for religion, even with differing agent types endowed with identical characteristics. It would appear that initial diversity leaves each individual group relatively weak. Should the population drift in any direction, one group may quickly grow from a slim, anomalous advantage to near-complete market dominance.

Similarly we find that corner solutions are more likely to emerge from religions with greater strength of origin, all other parameters being equal. Simply being able to survive the tumult population migration gives them more time in what is essentially a waiting game. This would explain why religions such as Christianity, who teach that leaving the religion is equivalent to damnation will tend to grow, and brings to mind the cliché of the "lapsed Catholic" who nonetheless still identifies himself as Catholic, perhaps even having his children baptized and confirmed.

While not as overwhelming as corner solutions, a tipping phenomenon is relatively common. When dealing with only two agent types (religions) whose exogenously determined characteristics are or are not equal, the population will nonetheless arbitrarily drift, and tip, towards one religion or the other. What is interesting is, with only two agent types, small differences are often overwhelmed by numbers, and the relatively less attractive religion can persist.

And more to come …

Although this paper introduces MARS in the context of a particular research question, we will conclude this section with some notes on the richness and scope of the MARS environment. Specifically,

The next draft will include additional examples, illustrating different outcomes and the impact of different behavioral assumptions.

The next draft will also describe the full simulation screen, not just the two boards. The screen includes a large collection of buttons and sliders that control numerous setup options and parameters, start and stop the execution, alter the visual display of information, graph key variables (such as the average participation rate for each type of agent, the population share of each type, and the overall degree of clustering within all types), and much more.

The next draft will provide the reader with a better feel for the range of effects (social, habit, etc.) that can readily be controlled via sliders and buttons on the simulation screen. For example:

> *- different types of religions can have totally different characteristics, including: their average rate of attendance, their propensity to hold onto agents who were born into that religion, their propensity to hold onto agents from one period to the next, their social (neighborhood) influence on their own members, and their social influence on non-members.*
>
> *- analogous differences can be selected for parameters that govern rates of religious involvement*
>
> *- in even is possible to control a full n x n matrix of interactions (where n is the number of different religions) so that agents of a particular type (say red which we could view as "evangelicals") tend to strongly attract another type of agent (say the greens whom we might view as "liberal Protestants") but tend to repel another type (such as the yellows whom we might view as "atheists" or non-Christians). We can thus analyze very complex (but very realistic!) patterns of interaction across a spectrum of faiths.*

The next draft will describe some of the many additional features built into the current version of MARS. These include:

> *- having agents sometimes reevaluate their religious choices <u>without</u> moving (so that the actual number of choices linked to moves ranges from 0 to 100%, rather than being set at 100% as described above).*
>
> *- having movers sample up to eight new locations before deciding where to move, and then moving to the location with the most densely populated neighborhood or alternatively the greatest share of neighbors who share the agent's religion. In the former case, agents tend to cluster in "cities" and "towns". In the latter case, they cluster around others of their own kind – which makes it possible for a small (sectarian?) type to persist even it never dominates a region and never converts many outsiders.*
>
> *- having agents take account of <u>both</u> the religious type <u>and</u> participation rate of their neighbors, so that conversion becomes more likely to a*

(sectarian?) religion that has extremely active members. In practice, this interaction makes it much more likely that small but highly dedicated sects will flourish while large but "lazy" denominations decline.

- randomizing factors that cause agent to be born with unique type and rate values distributed around the mean values for their particular type. The examples described in the paper leave this feature "off", so that all agents of a particular type are effectively identical except for their personal location (and immediate past).

- the ability to stop and start runs at will, save and revert to selected times in a run, alter the speed of execution, add more agents in the middle of a run, and more.

Conclusions

Almost a century ago, Emile Durkheim sought to understand religion by analyzing its most "primitive" instances; and Max Weber sought a similar understanding through methods of "ideal types." Although neither approach has proved very successful, the quest for simple starting points remains a compelling strategy. In recent year, theories of rational choice and religious markets have provided one such set of starting points, leading to new theories, new hypotheses, and strategies for empirical research. At them same time, one cannot deny that standard economic perspectives and methods give short shrift to localized interactions, social networks, and "emergent" patterns of behavior. Multi-Agent models provide a very different starting point for social theory and empirical research.

Multi-agent models excel at building micro-to-macro links. They offer new and precise ways to explore the implications of almost any theory that embeds individual choice within networks of relationships and other social structures. Moreover, the process of recasting a theory into M-A form automatically demonstrates whether it is sufficiently clear and complete to yield observable implications. If so, we can "run" the theory and literally see whether its results match its claims. We can also see how the consequences change when we alter the theory's assumptions (or more precisely, when we alter the parameter values that capture the theory's underlying behavioral assumptions). A theory may, of course, work inside a simulation yet fail in the real world, but good theory must at least succeed within the world implied by its own assumptions.

The M-A approach helps identify gaps in sociological arguments and methods, particularly those that link collective outcomes to individual propensities. As we have seen, many standard "micro-to-macro" claims must be rethought, and perhaps totally reconstructed. Regionalism is a case in point. Although MARS ultimately sustains longstanding claims about the impact of social networks and adaptation, it does so in ways that overcome serious flaws in past research. The old arguments lacked precision,

making it difficult to know their implications, let alone test them. The statistical methods used to assess the arguments were weak, and many methods contradicted the very assumptions from which the arguments derived. The standard stories about religious regionalism (and much more) needed major repairs, especially where micro met macro.

[Next draft will include several statements about the specific insights obtained above. Additional remarks will note MARS' capacity to analyze all sorts of behavior and situations, as well as MARS' ease of use and availability for both instruction and research.]

Multi-agent methods overcome many limitations associated with standard theory and methods while in no sense requiring the rejection surveys, statistics, or formal models of choice. Indeed, it provides a new framework in which to compare the behavioral results of standard theories – both economic and sociological – against the real-world outcomes identified by standard data sources. It is this largely unexploited capacity to leverage existing research that may ultimately prove to be the most compelling feature of MARS and M-A in general.

Appendix/Links

The full MARS model is much more sophisticated than that which has been presented here. Subsequent drafts of the paper will provide access to a "MARS User's Guide" and the underlying NetLogo code. Subsequent versions the code will probably be programmed using more sophisticated Java-based languages.

References

Coleman, James S. 1990. *Foundations of Social Theory.* Cambridge: The Belknap Press of Harvard University Press.

Epstein, Joshua M., Robert Axtell, and 2050 Project. 1996. *Growing artificial societies : social science from the bottom up*. Washington, D.C.: Brookings Institution Press.

Finke, Roger. 1989. "Demographics of Religious Participation: An Ecological Approach, 1850 - 1980." *Journal for the Scientific of Religion* 28:45-58.

Glenn, Norval D. and J. L. Simmons. 1967. "Are Regional Cultural Differences Diminishing?" *Public Opinion Quarterly* 31:176-193.

Granovetter, Mark. 1985. "Economic Action and Social Structure: The Problem of Embeddedness." *American Journal of Sociology* 91:481-510.

Halvorson, Peter L., William M. Newman, and Mark C. Nielsen. 1978. *Atlas of religious change in America, 1952-1971*. Washington: Glenmary Research Center.

Herberg, Will. 1960. *Protestant, Catholic, Jew; an essay in American religious sociology*. Garden City, N.Y.,: Anchor Books.

Kahneman, Daniel, Paul Slovic, and Amos Tversky. 1982. *Judgment Under Uncertainty: Heuristics and Biases.* Cambridge: Cambridge University Press.

Kennedy, Peter. 1998. *A Guide to Econometrics*. Cambridge, MA: The MIT Press.

Leamer, Edward E. 1978. *Specification Searches: Ad Hoc Inference with Non-experimental Data.* New York: Wiley.

Manski, Charles. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60:531-542.

McCloskey, Donald N. 1998. *The Rhetoric of Economics*. Madison, WI: University of Wisconsin Press.

Moulton, Brent. 1990. "An Illustration of a Pitfall in Estimating the Effects of Aggregate Variables on Micro Unit." *The Review of Economics and Statistics* 72:334-338.

Newman, William M. and Peter L. Halvorson. 1984. "Religion and regional culture: Patterns of concentration and change among American religious denominations, 1952-1980." *Journal for the Scientific Study of Religion* 23:304-315.

Reed, John Shelton. 1972. *The enduring South; subcultural persistence in mass society.* Lexington, Mass.,: Lexington Books.

Schelling, Thomas C. 1971. "Dynamic Models of Segregation." *Journal of Mathematical Sociology* 1:143-186.

Smith, Christian, David Sikkink, and Jason Bailey. 1998. "Devotion in Dixie and beyond: A test of the 'Shibley Thesis' on the Effects of Regional Origin and Migration on Individual Religiosity." *Journal for the Scientific Study of Religion* 37:494-506.

Stark, Rodney and Roger Finke. 2000. *Acts of Faith: Explaining the Human Side of Religion*. Berkeley: University of California Press.

Stump, Roger W. 1984a. "Regional divergence in religious affiliation in the United States." *Sociological Analysis* 45:283-299.

—. 1984b. "Regional migration and religious commitment in the United States." *Journal for the Scientific Study of Religion* 23:292-303.

—. 1986. "Regional variations in the determinants of religious participation." *Review of Religious Research* 27:208-225.

Voas, David, Daniel V. A. Olson, and Alasdair Crockett. 2002. "Religious Pluralism and Participation: Why previous research is Wrong." *American Sociological Review* 67:212-230.

Welch, Kevin W. 1983. "Community Development and Metropolitan Religious Commitment: A Test of Two Competing Models." *Journal for the Scientific of Religion* 22:167-181.

Wrong, Dennis H. 1961. "The Oversocialized Conception of Man in Modern Sociology." *American Sociological Review* 26:183-193.

Wuthnow, Robert and Kevin Christiano. 1979. "The effects of residential migration on chruch attendance in the United States." Pp. 257-276 in *Robert Wuthnow (ed.), The Religious Dimension: New Directions in Quantitative Research. New York: Academic Press*.

Notes

[1] About 17% of Americans move each year, and the typical American makes about 12 moves during his or her entire life [Source: Kristin A. Hansen, "Geographical Mobility", U.S. Census Bureau, landview.census.gov/population/wwww/pop-profile/geomob.html.] About 45% of Americans move at least once during any given 5-year period. Although most moves are local, about 106 million people (20% of the total population) moves from a different county, state, or country between 1990 and 1995. [Source: Current Population Reports, by Jason P. Schachter. "Geographical Mobility: 1990 to 1995." September 2000. U.S. Census Bureau.] Rates of long-distance migration have declined throughout the past 150 years, so these numbers understate the actual mobility behind the observed regional distributions. [Source: Claude Fisher and Michael Hout, "USA: A Century of Difference." Berkeley, 2000.]

[2] The first that we know of is an early simulation created in the 1980s by William Bainbridge (198x).

[3] We welcome would-be "Martians" to adapt or extend the underlying program and publish their results.

[4] Immigration explains only a small portion of the overall pattern. Inflows of Hispanics keep the Catholic population growing in several border states, but elsewhere immigration's impact is relatively small. Over time, immigration has radically altered the religious-ethnic mix of many neighborhoods, but it has had less impact on cities (at least since the 1920s), and impact at the state and regional level has been smaller still. Moreover, more than 90% of all moves occur within the United States, and these internal flows really do mimic random religious mixing.

[5] For reasons why "significance" almost never deserves serious attention, see McCloskey McCloskey, Donald N. 1998. *The Rhetoric of Economics*. Madison, WI: University of Wisconsin Press. and Leamer Leamer, Edward E. 1978. *Specification Searches: Ad Hoc Inference with Non-experimental Data*. New York: Wiley.. For a more general discussion of estimation problems, including heteroscadasticity, autocorrelation, multicolinearity, misspecification, and much morem see Kennedy Kennedy, Peter. 1998. *A Guide to Econometrics*. Cambridge, MA: The MIT Press.. Although a massive body of theory and methods have been developed to address these problems, it remains unclear whether statistical innovations of the past generation have actually added much to our understanding of human behavior.

[6] One might argue that the relative *magnitude* of these effects was not previously known, but one must also concede that the estimated magnitudes vary considerably from one study to the next and that no research has ever converted these numbers into *behavioral* parameters that might show – at least in theory – how changing patterns of mobility affect regional religiosity

[7] One might try other tricks, but the old problems tend to remain and new one's get added. For example, it does not help to restrict the respondent s ample to recent immigrants while estimating the social effects based only on the non-migrants in region or neighborhood. Nor can we patch things up with intertemporal data that permit us to regress the current religiosity of individuals onto past religiosity of their regions or neighborhoods.

[8] Statistically-minded readers will note that there exist estimation methods to circumvent these types of estimation problems – see for example, Manski Manski, Charles. 1993. "Identification of Endogenous Social Effects: The Reflection Problem." *Review of Economic Studies* 60:531-542.. All we need is sufficient information about the structure of the data, such as the precise form of the equations that determine individual-R as a function individual- and group-level effects. But these "solutions" merely highlight the fundamental problem. In the case at hand and many others like it, social researchers employ standard statistical techniques precisely because they generate seemingly informative results in the absence of information about the underlying social dynamics. The entire point of the statistical exercise was, after all, to obtain such information on the cheap.

[9] For an overview of this literature in the context of conversion, see Stark and Finke Stark, Rodney and Roger Finke. 2000. *Acts of Faith: Explaining the Human Side of Religion*. Berkeley: University of California Press.

[10] This by no means exhausts the lists of serious omissions in the standard estimation models. (For example: It ignores the fact that newcomers' social networks look nothing like those of longtime residents. Newcomers initially exert relatively little social influence on those around them, whereas random others may heavily influence them because moving has disrupted their old ties. Over time, however, relationships and social influence become more symmetric. A good model should almost certainly address these and other "asymmetries" that exist across actors and over time.

[11] For more on the micro-macro problem in social-scientific research, see Coleman Coleman, James S. 1990. *Foundations of Social Theory*. Cambridge: The Belknap Press of Harvard University Press. and Wrong Wrong, Dennis H. 1961. "The Oversocialized Conception of Man in Modern Sociology." *American Sociological Review* 26:183-193..

[12] There are a variety of synonyms at work in the field, of which Multi-Agent model is only one. Other popular terms include Agent Based Model and Computational Model.

13 M-A methods are neither the only alternatives to statistical estimation, nor always the best. Appendix B reviews some other approaches, including Markov Chain models.

14 For the seminal work on the application of multi-agent computer simulations of social models, see J. Epstein and R. Axtell's *Growing Artifical Societies* Epstein, Joshua M., Robert Axtell, and 2050 Project. 1996. *Growing artificial societies : social science from the bottom up*. Washington, D.C.: Brookings Institution Press..

[15] The social system "tips" because almost any initial configuration includes a few individuals surrounded almost entirely by people not like them. When these people move, their departure tips their old neighborhoods even farther (hence making them even less attractive to people like them) and their arrival tends to tip their new neighborhoods in the opposite direction.

[16] Sources: Current Population Reports, by Jason P. Schachter. "Geographical Mobility: 1990 to 1995." September 2000. U.S. Census Bureau.] Schachter provides a breakdown of the reasons movers typically give in Current Population Survey interviews. The primary determinants of mobility include are family-related, work-related, housing-related, and together with college, health, climate, these reasons account for 98% of all stated reasons for moving. On the other hand, school quality and population characteristics (such as age, race, ethnicity, and socio-economic status) do influence the *neighborhoods* that movers choose, but movers can be reasonably confident that almost any major metropolitan areas will contain several neighborhoods that satisfy their demographic preferences. Moreover, the vast majority of Americans seem content to live within easy driving distance of one or two congregations of co-religionists. In contrast to other population characteristics, they voice little concern about the religious composition of their neighborhood or overall metropolitan area. The notable exceptions are members of very small and distinctive religious minorities, such as Orthodox Jews and Old Order Mennonites, who usually live in enclaves of co-religionists. Such people, however, make up a tiny share of the American population and an even smaller share of movers.

[17] On the other hand, as seen in figure 2b, random mixing almost always leads to many small clusters in which one group is far more or less prevalent than expected. This inevitable consequence of true randomness violates our *biased* intuition about randomness and often leads us to impute special conditions where none exists. The character and consequences and this and many other cognitive biases have been studied extensively by behavioral psychologists and economists Kahneman, Daniel, Paul Slovic, and Amos

Tversky. 1982. *Judgment Under Uncertainty: Heuristics and Biases*. Cambridge: Cambridge University Press.

[18] Briefly Defined: *Mobility* here indicates migration occurs continuously, as households move one-by-one, rather than intermittently or in groups (hence we do not model the system as in a Markov process that has various shares of the populations making "state transitions" each period). *Randomness* refers to the fact people move without regard for the religious characteristics of locations. *Individuality* means that personal attributes and personal history influence religious behavior. *Adaptation* means that people respond to social influence, and localism limits that influence to their networks of close contacts. *Quasi-symmetry* means that when people move, social influence initially runs only in one direction, from the existing neighbors to the newcomer. *Denominationalism* means that behavior varies across different types of religion. *Diversity* means that religious variation exists at every level – individual, neighborhood, and region. *Disequilibrium* means that the system keeps changing at every level, rather than converging to a static equilibrium.

[19] A "move" may also involve remaining in place. In such cases, the agent re-evaluates his type and rate choices based on his own current state and those of his neighbors. The *move-rate* slider determines what percentage of "moves" involve actual changes in location. The default *move-rate* = 100; hence by default all "moves" are to new locations. By setting *move-rate* = 0, no agent ever changes location; hence the simulation operates in "cellular automata" mode.

[20] For clarity it can be assumed that when we refer to a "parameter" or "parameter value" we are describing an exogenously controlled variable that is determined by the observer. Thus we can run a variety of "experiments" simply by adjusting the various parameters.

[21] Note that a different set of *_origin*, *_inertia*, and *_community* parameters are associated with each agent type. Hence, the "red" religion may be relatively effective in instilling loyalty (high *r_ origin*), but not very effective in holding on to converts (low *r_ inertia*), and only moderately effective in maintaining or attracting people via social effects (average *r_ community*).

By setting all *_ origin* and *_ inertia* parameters to zero and all *_ community* parameters equal to the same (positive) value, each agent will simply mimic the majority of his neighbors. This is equivalent to a standard agent-based model of social conformity. The result, over time, is complete uniformity within any given cluster of agents.

[22] As currently programmed, the MARS regions are fully bounded so that agents "feel" no social influence beyond the edges of each sub-region. This results in noticeable edge effects, because agents located along the sides and corners of a region have fewer neighbors than other agents. With a bit of additional programming, one may introduce alternative "topologies" which wrap each sub-region around its vertical borders, horizontal borders, or both. Our initial but limited experiments have not persuaded us that these alternatives produce very different results or provide much additional insight.