

Human Capital and the Supply of Religion*

Joseph Engelberg
Raymond Fisman
Jay C. Hartzell
Christopher A. Parsons‡

January 2013

Abstract: We study the role of labor inputs in religious attendance using data on Oklahoma Methodist congregations during 1961-2003. Pastors play a significant role in church growth: replacing a 25th percentile pastor with a 75th percentile one increases annual attendance growth by three percent. A pastor's performance in his first church (largely the result of random assignment) predicts future performance, suggesting a causal effect of pastors on growth. The deployment of pastors by the church indicates efficient use of labor: low-performing pastors are more likely to be rotated or exit the sample, and high-performing pastors are moved to larger congregations.

*We have benefited from discussions with Dan Hamermesh, Jonah Rockoff, Douglas Staiger, and seminar participants at Columbia University, Rice University, University of California at San Diego, University of Maryland, University of Southern California, and University of Texas at Austin. All errors are our own.

‡ Contact: Joseph Engelberg, 9500 Gilman Drive, University of California at San Diego, La Jolla, CA 92037. Raymond Fisman, Uris 605B, Columbia University, New York, NY 10025. Jay Hartzell, 1 University Station B6600, Austin, TX 78712. Christopher Parsons, 9500 Gilman Drive, University of California at San Diego, La Jolla, CA 92037.

Religion plays a central role in the lives of many, even in the modern world. A significant majority of the world's population remains affiliated with a major religion (*CIA Factbook*, 2012), and a 2009 Gallup survey involving 143 countries indicates that it “plays an important part [in the] daily lives” of over 80 percent of respondents.¹ Moreover, whether by promoting education (Becker and Woessmann (2009)), maintaining ethical systems (Weber (1905), McCleary and Barro (2006)), or fostering subjective well-being (Ellison (1991)) religion has been associated – if not credited – with a variety of welfare improvements. Over a third of U.S. charitable contributions go to religious organizations, greater than \$100 billion in 2008 (Giving USA, 2009). Given its broad and deep impact on the economy and society, economists have increasingly sought to understand what motivates individuals to devote resources to religious activities, versus their secular alternatives.

In this paper, we provide an analysis of the *production* side of religious attendance utilizing a unique human resource database from the United Methodist Church of Oklahoma.² Whereas other research focusing on the supply side of religion has considered the effects of market structure and regulation, we believe we are the first to look within the church organization to analyze the determinants of religious participation.

Using internal records from every Methodist congregation in Oklahoma during 1961-2003, we estimate the role of individual pastors in religious participation. Because Methodist pastors are typically assigned to serve at several churches (within Oklahoma) over their careers, we can observe the same church over the tenure of several pastors,

¹ <http://www.gallup.com/poll/114211/Alabamians-Iranians-Common.aspx>

² While the formal title of the denomination is “United Methodist” since 1969, we use “Methodist” for ease of exposition.

and the same pastor over several different churches, allowing for credible identification of pastor effects on church performance.

There is good reason to focus on labor inputs into religious attendance. In the economics of organizations literature, there is an emerging body of work that finds a critical role of managers and leaders for organizational productivity (see, in particular, Branch et al. (2012), Lazear et al. (2012) and Mollick (2011)). One might expect this to extend to religion: independently run congregations have grown into the tens of thousands, and their leaders – ‘superstar’ pastors like Joel Osteen and Rick Warren – are widely viewed as essential to their successes.³

In this paper, we document the impact of pastors on attendance growth, our primary measure of pastor success in furthering the church’s mission. Further, given that we find pastors to be important determinants of church growth, we then ask whether the Methodist Church organization utilizes pastors in a way that is consistent with attendance maximization.

Our findings can be summarized as follows: Using the empirical Bayes methodology to assess the performance contributions of individual pastors, we find that they collectively play an important role in church growth. Our estimates imply that replacing a 25th percentile pastor with a 75th percentile one increases annual attendance growth by three percent, similar to the effect of a 10 percent increase in the surrounding county’s population. The inter-quartile range of pastor effects is 0.15 of a standard deviation in attendance growth, a magnitude comparable to the impact of individual teachers on student performance (e.g., Rockoff (2004)) which, to date, has

³ See, for example, http://www.forbes.com/2003/09/17/cz_lk_0917megachurch.html.

been the primary application of the empirical Bayes method in economics. Moreover, we find that pastors are important in both small and large churches, and further that performance in small congregations is correlated with performance in larger ones. This result is surprising if personal contact with parishioners is crucial for a pastor's success. On the other hand, many of a pastor's activities (e.g., giving sermons) scale easily, suggesting that skills related to these latter types of activities drive the relations between specific pastors and performance.

An important caveat to these results is that pastors and churches are generally not matched randomly, allowing unobservable determinants of church growth to conflate our estimates of individual pastor productivity with church quality.⁴ Similar to a teacher being awarded credit for inheriting a class of promising students, a pastor's measured skill may be biased upward if he is assigned to a church with high growth potential. In this case, at least part of our estimated pastor value-added would be spurious, reflecting the effect of determinants of church performance omitted from our regressions.

The rich panel structure of the data provides two ways to address this concern. The first can be likened to a regression discontinuity. Whereas many determinants of church growth trend gradually – e.g., income growth, population changes, or competition from other sects – pastor rotations are discrete events in a church's life that, if causal, should be associated with abrupt spikes or drops in attendance. We conduct a simple variance decomposition of attendance growth which shows that volatility is concentrated in years in which pastors rotate: variance in attendance growth

⁴ For a discussion of this issue in the context of estimates CEO fixed effects on firm performance, see Bertrand and Schoar (2003).

is roughly 60 percent higher in rotation years than in years when there is no pastor change. While this does not completely rule out a possible role of unobservable church attributes, it does narrow the plausible set to those that are foreseeable approximately a year in advance (when pastor assignments are made), discrete, and church-specific.

Our second approach to addressing the issue of non-random matching aims to approximate the ideal scenario of exogenous pastor rotations, based on the assumption that “freshman” pastor assignments are “almost” random. As a pastor launches his career, there is very little information about his capabilities or fit with a particular church type. Under this assumption, a pastor’s first church assignment is largely a matter of chance.⁵ We proceed as follows: we first take the abnormal attendance growth from a pastor’s first church assignment as his measured ability, and then use this residual to predict attendance growth in his later church assignments. We find that a pastor’s first-church performance strongly predicts church growth in subsequent placements. Further, echoing our earlier findings on volatility, the link between first and subsequent church performance is concentrated in the initial year of placement in these later assignments (i.e., the first year of a pastor’s third, fourth, etc., churches). This pair of findings is difficult to reconcile with explanations based on unobserved heterogeneity, and further supports the claim that pastor ability is important in explaining religious participation and church growth.

Given the significant and persistent effects of pastors, we conclude by examining how the Oklahoma Conference utilizes past performance in making pastor appointments. We find that underperforming pastors are rotated earlier to new

⁵ We will argue as well that the data validate this assumption. We find that pastors that perform well in later church assignments are not given larger congregations as initial assignments, nor churches with high growth in the years leading up to the initial assignment.

churches, or exit the sample entirely, consistent with the church bureaucracy searching for better matches and greater attrition of low-performing pastors. For those remaining in the sample, we find little evidence of performance improvements across congregations or across years within a given pastor-church match, suggesting that pastor learning and better match quality over time play limited roles in explaining pastor productivity. Finally, given the apparent scalability of pastor ability in larger congregations, an attendance-maximizing church bureaucracy should appoint high-performing pastors to larger churches. Consistent with this view, we find that a pastor's lagged residual attendance growth predicts church size at future assignments. Our evidence thus broadly favors a view of the Oklahoma United Methodist Conference as utilizing pastor ability to maximize overall attendance growth.

The primary contribution of this paper is to provide what we believe to be the first microeconomic analysis of religiosity by utilizing individual-level data from the Methodist Church – a crucial step forward in understanding religiosity, given the importance often ascribed to the “production” side of religion. By contrast, most prior empirical work in the economics of religion has focused on the “consumer” side of religiosity. Azzi and Ehrenberg (1975) provided the first theoretical framework for analyzing religious participation, framed as a household production problem. Religious activity increases until the marginal cost of time or money offsets the marginal benefit, be it in this life (e.g., social value) or the next. A number of studies provide empirical support for this model. For example, church attendance follows a U-shaped profile with age (Neuman (1986)), and falls when statutes that forbid Sunday shopping are repealed (Gruber and Hungerman (2008)); both patterns are consistent with opportunity costs influencing religious decisions.

We are certainly not the first to consider the supply-side determinants of religiosity. Finke and Stark (1992), for example, credit the growth of some congregations over others – such as Baptist over Episcopal – to superior incentive schemes. In earlier work using the same data employed in this paper, Hartzell, Parsons, and Yermack (2010) show that Methodist pastors’ pay is sensitive to performance. Much of the supply-side literature has focused on market-wide considerations, in particular the role of market structure and government regulation in predicting religiosity (see, for example, McCleary and Barro (2006)). Hanson and Xiang (2011) examine cross-country patterns in attendance for protestant denominations and conclude that their governance structures are an important determinant of a denomination’s market share. Our work provides micro-level support for the view that it is important to consider intra-organizational dynamics in understanding the productivity of religious organizations. Further, our findings on the scalability of pastor ability may additionally provide some empirical grounding for the mega-church phenomenon, as it suggests the presence of possible superstar effects in the spirit of Rosen (1981).

Finally, we contribute to the emerging literature in personnel economics that tries to understand the role of individual managers in explaining residual organizational performance. Beyond examining a particularly important application in this paper, our data are well-suited to the task of analyzing individual contributions to organizational performance given the high frequency of rotation across churches, and the long panel of pastors that we observe.

The rest of this paper is organized as follows: Section I provides some background on the Methodist Church in Oklahoma. Section II provides a description of the data. Section III provides evidence that human capital in the Methodist Church is

important, describing both our empirical framework and providing estimates of pastor-specific impacts on church productivity. Section IV then considers how the Methodist Church allocates human capital, analyzing pastor rotation and exit across churches and over time. Section V concludes.

I. Institutional background

The Methodist Church is the second largest Protestant church in the U.S., with 7.77 million members domestically as of late 2011.⁶ Founded originally by John and George Wesley as a “methodological” offshoot of the Anglican Church, Methodists first appeared in the U.S. in the late 1700s. The Methodists take their cue from Matthew (28:19-20), where Jesus tells his followers: “Go therefore and make disciples of all nations, baptizing them in the name of the Father and of the Son and of the Holy Spirit, and teaching them to obey everything that I have commanded you to attract disciples.”

The Methodists have embraced this “Great Commission” and focus on finding new members, primarily through “disciple making” in local churches.⁷ Reflecting this stated objective, we take growth in attendance as our primary measure of pastor performance. By adding new attendees, a pastor generates a stream of future benefits to the church in the form of revenues and participation in church services and activities.

The basic unit of organization in the Methodist Church is a *conference*, of which 63 exist within the United States. Conferences are often organized by state – as is the case for Oklahoma – but there are exceptions; large states are sometimes partitioned into multiple conferences, and likewise, conferences can span several smaller states. A

⁶ 2011 Yearbook of American and Canadian Churches.

⁷ http://www.umc.org/site/c.lwL4KnN1LtH/b.2295473/k.7034/Mission_and_Ministry.htm, accessed November, 2012.

single bishop governs each conference.⁸ Conferences are further split into districts, each headed by a superintendent who is appointed by the bishop to a six-year term. Oklahoma has 12 districts, each comprised of 35 to 75 local churches.

All Methodist congregations have a single “senior” or “head” pastor. The pastor’s most visible contribution is the delivery of prepared remarks (a sermon) on Sundays, although numerous other tasks fall under his control. These include administrative obligations, such as oversight of the church’s finances, building campaigns, and clerical personnel. In addition, our communication with the Oklahoma Conference reveals that the bulk of a pastor’s time is spent interacting with church members – visiting hospitals, taking food to the elderly or infirmed, providing marriage or divorce counseling, presiding over funerals, speaking at graduations, and so forth.⁹

Unlike many competing sects such as Presbyterians or Baptists, where pastors are free to move across churches at will, pastor assignment in the Methodist Church is governed centrally by the conference, at the discretion of the bishop. Indeed, a primary function of the conference hierarchy, perhaps its most important, is to allocate pastors to churches. Usually, Methodist ministers work only for one conference over their careers,¹⁰ although the typical pastor will serve at several churches in that time.

Ordained pastors are typically hired from graduate school (seminary) by a particular conference – in our case, the Oklahoma Conference – and assigned to work with a local congregation. Both a pastor’s initial placement and subsequent assignments

⁸ Smaller conferences sometimes share bishops. There are 50 bishops overseeing the 63 conferences within the U.S.

⁹ See *The Book of Discipline of the United Methodist Church* (2008) for a formal description of a Methodist Pastor’s disparate responsibilities.

¹⁰ While we do not have the data to test for the frequency of inter-conference movements, our conversations with officials from the Methodist Church confirm that such movements are rare.

are determined by the district superintendent and conference bishop, not by individual congregations.¹¹ Although local churches have some responsibility for setting the pastor's pay and other aspects of his compensation, they do not directly influence where he serves.

Typically, a pastor is assigned to a single church, but in some cases, he may serve simultaneously at two or more small churches located near one another. Because these so-called circuit churches share a pastor and are often coordinated in other ways, we consider all congregations affiliated with a particular pastor in a given year as a single unit.¹² Our results are not sensitive to this aggregation.

II. Data

The Methodist Conference of Oklahoma has, since at least 1961, collected detailed annual records on each congregation including data on its membership, attendance, finances, personnel, and other activities. The causes of membership changes are also detailed, with records on changes occurring, for example, via baptisms and members' deaths. Our main performance measure is the change in the average weekly attendance at all Sunday worship services. This variable is one of several pieces of information that the conference requires each local congregation to report annually, along with data such

¹¹ From our conversations with church leadership, the annual appointment process begins with each pastor and each church's pastor-parish relations committee submitting to the conference respective forms, privately indicating whether each side would prefer that the current pastor stay in place, move, or indifference between the two. The conference takes these responses into consideration when reallocating pastors, but has full authority over both the timing of pastor moves and their destinations. We do not have access to the indicated preferences over moving, just the final outcomes of the conference's decisions.

¹² Circuit churches comprise about 25% of our observations. In most cases, circuit composition does not change year-to-year, but they are occasionally recombined – e.g., a given pastor is assigned to churches *A* and *B* in year *t*, but is assigned to churches *A*, *B*, and *C* in year *t+1*. When these occur, we treat recombined circuits as *de novo* assignments.

as average participation in Sunday school (e.g., weekly Bible study classes and children’s education programs), and attendance in youth programs. The main financial variables include a rough balance sheet, pastor pay (salary and expense allowances), donations to foreign missions, and building expenditures. Importantly, church personnel – specifically the church’s head minister – are identified by full name, allowing us to track specific ministers over their careers with the Oklahoma Conference.

These data are assembled and housed centrally at conference headquarters in Oklahoma City. In 2004, we were granted access to one district’s entire catalog, covering 1961-2003. Over the course of roughly two years, a third party constructed an electronic dataset using a combination of optical character recognition software and hand checking. Our audit suggests the data is of very high quality, on the order of 1:10,000 data entry errors, and the accuracy is further supported by verification that calculated totals of various items match the listed totals in the raw sources.

Table 1 presents summary statistics for our full sample, with the top panel showing data organized by pastor, and the bottom panel by church, combined across circuit churches when they occur.¹³ The first row indicates that in the average year, 459 unique Methodist pastors are employed in Oklahoma, with relatively mild year-to-year fluctuations; the inter-quartile range is 447 to 471.¹⁴

The second and third rows show summary statistics for pastor tenure, both over his entire lifetime (row 2) and at specific churches (row 3). With the caveat that the data are right-censored, we find that pastors serve, on average, for 2.64 years at each church

¹³ The data the Oklahoma Conference provided has two holes. First, no data on attendance are available for 1963-64; further, no data are available for the Stillwater District for 1982 and 1990.

¹⁴ Outright church closures are rare. In personal correspondence, the Conference estimates that there were probably fewer than a dozen closures in the period we examine – approximately one church every four years. Further, the Conference only shuts low-attendance churches so the impact on our analysis is very unlikely to be significant.

assignment, with fewer than five percent of pastors lasting greater than six years with a given congregation. This implies that the typical pastor is rotated at an annual rate of about $1/2.64 = 38$ percent, which slows down a bit as a pastors gain experience. We also define a variable, *Time_in_Sample*, which captures the number of years since a pastor first appeared in the conference's records. A pastor appears in our dataset for nearly nine years on average.

In the fourth row, we show the total number of lifetime church assignments for a pastor. These rotations will play a crucial role in our analysis, as they allow us to separately identify pastor effects and church effects in explaining church performance. Although the table indicates that an average pastor serves at slightly more than three churches, roughly one quarter of pastors are assigned to five or more churches over their careers.

The next set of rows summarizes the flow of pastors in and out of the Oklahoma Methodist Conference. In an average year, roughly 41 pastors exit the system, almost identical to the number of pastors that enter (40). To put the entry and exit rates in perspective, compared to the total number of pastors (459), about nine percent of churches or church circuits will inherit a rookie pastor, i.e., one with no previous experience as a head pastor. A similar number will lose their pastors to exit from the Oklahoma Conference; finally, as indicated by the large fraction of pastors with low overall tenure, many pastors are likely switching professions rather than retiring outright.

In the second panel of Table 1, we summarize several church-level performance metrics. One primary measure of a church's health is the average number of people that attend Sunday morning worship services at church c in year t , *Avg_Attendance_{ct}*. Across

all churches and years, the mean of *Avg_Attendance* is 145, although this is somewhat skewed by a few large churches, such as St. Luke's in Oklahoma City, with average attendance of 3,559 in 1962. However, with an inter-quartile range of 58 to 152, it is clear that the typical observation in our sample is a local, neighborhood church.

Most of our analysis focuses on the annual change in attendance at each church c , $Attendance_Growth_{c,t} = \log(Avg_Attendance_{c,t}) - \log(Avg_Attendance_{c,t-1})$. We focus on attendance rather than membership because it is apparent from our discussions with Church leaders that membership figures are seen as less representative of the current health of a church than attendance data. For example, because "members" can stop attending or participating years before they are formally removed from the rolls, using membership changes will sometimes misattribute one pastor's poor performance to his successor. We focus on changes rather than levels in attendance to normalize churches of very different sizes, and to allow for different long-term trend rates across churches (through the inclusion of church-level fixed effects).

As the second row of Table 1 shows, average attendance growth is slightly negative over our sample.¹⁵ Because we are focused on attendance growth attributable to individual pastors, it is important to control for exogenous determinants of church performance, such as fluctuations in local population. The third row standardizes each church's attendance growth by county-level population growth using data from the U.S. Census (<http://www.census.gov/popest/estimates.php>). This proxy for religious demand is imperfect, in part, because churches often draw parishioners from across county lines. Although population growth in the typical Oklahoma county appears to

¹⁵ More generally, church attendance in the Oklahoma region has been declining for decades. As one indication of this, using data from the General Social Survey, we find that the correlation between an ordinal measure of church attendance frequency and year of interview is -0.52 for respondents in the West South Central region, which contains Oklahoma.

have slightly outstripped church attendance growth over our sample, this adjustment is minor, reducing the normalized average attendance growth to -1.90 percent.

In the next few columns, we provide summary statistics on several alternative measures of pastor performance: the log change in church membership, $Membership_Growth = \Delta \log(Membership_{c,t})$; the log of additions to membership through baptisms, $\log(1+Baptisms_{c,t})$; and $Revenue_Growth = \Delta \log(Revenue_{c,t})$. As noted above, we believe *Attendance Growth* is the most appropriate measure of church performance; we present results based on these alternative measures – all highly correlated with one another – as robustness checks. We note that the last of these, *Revenue_Growth*, cannot be calculated directly because revenue figures are not reported directly to the conference by each church. However, we are able to calculate a rough estimate of donations by comparing each church’s expenses, which are reported, and changes in assets and liabilities, which are also reported, albeit imperfectly.¹⁶ Finally, we use $\log(1 + Deaths_{c,t})$ in a placebo test, because it is the component of attendance growth that is most plausibly beyond the pastor’s control.

III. Pastor-specific church performance

We now turn to our primary research question: how much does individual pastor ability affect Sunday attendance? We provide an answer in two steps. We first present a set of raw correlations relating pastor assignments to church attendance. Following pastors’ careers from church to church within Oklahoma, we show that some pastors are consistently associated with attendance growth, while others tend to coincide with lower

¹⁶ See Hartzell, Parsons, and Yermack (2010) for more discussion about inferring church revenue in this context.

performance. In Section V, we then provide a set of further results – building on the institutional details of the pastor assignment process – to argue that the pastor-attendance relation is likely causal.

a. Non-parametric evidence

We begin with a set of non-parametric analyses on pastor assignment and church performance. Denoting individual pastors with i , and each of pastor i 's successive church assignments with $a_i=1, 2, 3, \dots, A_i$, we look for serial correlation in performance across consecutive church postings. We ask, for example, if it is the case that a pastor who rapidly grew attendance at his first church is likely to grow attendance at his second church assignment as well.

Our methodology is as follows. Among the $\sum_{i=1}^I a_i$ unique pastor-church assignments in the data set, we calculate the average yearly attendance growth, rank pastor-church assignments from lowest to highest, and then form equally sized quartiles. Average annual attendance growth cutoffs for the 75th, 50th, and 25th percentiles are 2.6%, -1.6%, and -7.0%. With these breakpoints in hand, we then assign each unique pastor-church assignment a_i to the appropriate quartile, based on average yearly attendance growth during a_i . For instance, suppose that a pastor was assigned to five churches over his career which experienced average yearly attendance growth under his tenure of 8%, -2%, 2%, -10%, and 3%. Combining these with the breakpoints above, the appropriate performance quartiles would be 4, 2, 3, 1, and 4, respectively.

In Table 2, Panel A we present the results of a transition matrix, with columns showing quartiles of attendance growth for a pastor's current assignment, and rows indicating performance quartiles for his immediately prior assignment. For every

column and row, the diagonal element contains the highest value, indicating a high degree of serial correlation in performance across assignments.

While there is a tendency for pastors' performance to persist in the same quartile across assignments, this pattern is most striking among the worst (quartile 1) and best (quartile 4) performers. In quartile 1, over a third of the worst-performing quartile will rank in the worst quartile in their next assignment, with over 61% ranking below the median, i.e., in either quartile 1 or 2. Further, a bottom-quartile (1) pastor has only a 15% chance of ranking in the top quartile (4) during his next assignment (column 4, row 1), nearly identical (16%) to the chance that a top-performing pastor (quartile 4) will fall to the lowest group (column 1, row 4).

The transition probabilities shown in Panel A of Table 2 are suggestive of persistence in pastor performance, but also consistent with other explanations based on omitted variable bias – such as the consistent placement of some pastors in growing areas and others in declining ones. While we take up this set of concerns in greater detail in Section V, Panel B of Table 2 provides some preliminary evidence against the most straightforward alternative explanations based on omitted variables. In this panel, we subtract the percentage change in population in the surrounding county from church attendance growth, thus filtering out the effects of area's population growth, over which the pastor clearly has no influence. This adjustment has a negligible effect on the transition matrix: The diagonal elements remain the highest values in any row or column, and the chance that an extremely high or low performer experiences a dramatic change of performance remains very unlikely.

b. Fixed effects regressions

The non-parametric comparisons in Table 2 can be extended to a regression framework that identifies pastor-specific effects in church performance through the inclusion of a group of dummy variables, one for each pastor in the data set. This methodology has been used in the teacher value-added (e.g., Rockoff (2004)) and personnel economics (e.g., Lazear, Shaw, and Stanton (2011)) literatures, among others.

We start with a standard fixed effects model to examine the incremental impact of pastor fixed effects on adjusted R^2 . These results are reported in Table 3, Panel A. We begin by regressing attendance growth on county-level population growth in Column 1. As expected, there is a very strong correlation between the two: growing communities have an increasing population base to draw upon for church attendance. However, we note that the elasticity is much less than one. This may derive from the coarse level of aggregation, or from the differential rates of church attendance from migrants versus long-term residents.

Column 2 adds year fixed effects, resulting in a modest increment in adjusted R^2 to 0.015, suggesting that after controlling for county population, statewide common shocks are relatively unimportant.¹⁷ Column 3 adds church fixed effects, which generates a further increment in adjusted R^2 to 0.102.

Column 4 shows the main result. When we add pastor fixed effects to a regression of *Attendance_Growth*, the adjusted R^2 increases to 0.26. Thus, while individual churches explain some of the sample variation, as a group, pastors are collectively the most important determinant of church growth. While significant caution is in order in

¹⁷ In unreported results, we also find that the attendance growth of one church in a district is uncorrelated with the attendance growth of other churches in the district after controlling for zip-code level population growth. This suggests there is little evidence of district-wide common shocks to religiosity after controlling for population dynamics.

comparing our findings to those derived from vastly different institutional contexts with different data constraints, it is useful to benchmark the role of pastors against those from secular settings. We find the incremental explanatory power of pastors is much higher than that of CEOs (Bertrand and Schoar, 2003), and comparable to that of supervisors at a technology services firm (Lazear, Shaw, and Stanton, 2011).

In the next two columns, we examine whether the role of pastor performance differs by church size, a relationship that is ambiguous given the range of pastor tasks and their differing degrees of scalability – “non-rival” jobs like sermonizing are scale invariant, while counseling and other “high touch” activities will not scale. If the pastor is first and foremost a church manager, his talents may face diminishing returns in parish size, in the tradition of Simon (1971); other administrative tasks like management of finances are somewhat independent of scale. To the extent that pastor ability does scale, one might predict that an attendance-maximizing church will promote high-performing pastors to larger parishes, a point we return to below.

In the final two pairs of columns in Panel A of Table 3, we investigate this scalability, quantifying the marginal explanatory power conferred by the introduction of pastor fixed effects for small and large churches separately. For small churches (those below the median in attendance), the inclusion of pastor fixed effects increases the R^2 from 9 percent (column 5) to 35 percent (column 6). As a percentage improvement, this is twice what we observe for big churches (17 percent to 35 percent), although the main difference is simply that church fixed effects appear more important for large churches. In both cases, models that account for pastor quality allow us to ultimately explain over a third of the total variation in church growth.

As emphasized by the teacher value-added literature, the estimated fixed effects in OLS regressions such as those shown in Panel A of Table 3 will overstate the true dispersion in performance among pastors, because the fixed effects themselves are measured with noise. We thus provide estimates based on the empirical Bayes method (Morris (1983)) that has become standard in the teacher performance literature.¹⁸ Intuitively, the empirical Bayes approach shrinks estimated fixed effects to account for the noise in each pastor's measured performance. For example, a pastor who consistently had above average attendance growth at every church in his career would have little shrinkage. By contrast, a pastor who had a similar average effect on attendance growth, but with high year-to-year and church-to-church variability, would have his fixed effect reduced to reflect uncertainty over whether the positive effect could truly be attributed to the pastor. In practice, the model is implemented using a mixed multilevel model with church fixed effects and pastor random effects.

In Panel B of Table 3, we present the percentile distribution of our empirical Bayes estimates, with the estimated random effects derived using the best linear unbiased predictors (Goldberger (1962), Morris (1983)). To ensure that the random effects model is well behaved, we limit the sample to churches with at least 20 years of data, although we note that the patterns reported here are not sensitive to the precise choice of cutoff.

For the entire church sample (row 1), the interquartile range of the empirical Bayes estimates is approximately 2.7 percent (-1.4 percent to 1.3 percent), or roughly 15 percent of the sample standard deviation of *Attendance_Growth*. Another way to

¹⁸ For recent examples, see Raudenbush and Bryk (2002), Rockoff (2004), Gordon, Kane, and Staiger (2006), and Kane, Rockoff and Staiger (2008).

appreciate the magnitude is to compare it to the effect of local population growth on attendance changes. Taking either of the first two columns in Panel A of Table 3 as a guide, a 10 percent increase in *County Population Growth* translates to an increase of approximately $0.1 * 0.28 = 2.8$ percent in *Attendance_Growth*, similar to replacing a pastor in the bottom quartile with one from the top quartile. Again, with a caveat on the difficulties in making comparisons across institutional domains, this is similar (though slightly lower) in magnitude to the importance of teachers in explaining student test score outcomes based on findings in the teacher value-added literature, as reviewed in Chetty et al. (2012).

In the next two rows, we present the same quantities, allowing for two separate random effects for each pastor, one each for large and small churches. The data suggest that the pastor's role does indeed scale with congregation size – the interquartile range of pastor effect estimates for small and large churches are very similar, with both around three percent. We note, furthermore, that there is some evidence that pastor talent in small churches carries over to larger ones – the correlation between the estimates of pastors' individual effects for large versus small churches is 0.102 (p-value=0.017).

IV. Causality and non-random pastor assignment

While the univariate evidence in Table 2 and regression results in Table 3 are consistent with individual pastors collectively having a causal influence on church attendance, it is also possible that pastor assignments are systematically related to unobservable determinants of church growth. To the extent that we have omitted such factors from our analysis, our estimates of pastor value-added in the last section may be biased.

The direction of this bias, however, depends on the correlation structure between a pastor's true value added and the omitted determinants of attendance growth, which in turn depend on the conference's ultimate objective when assigning pastors to congregations. If the conference attempts to match the best pastors with churches whose unobservable (to us) factors suggest they are most primed for growth, we will overestimate the importance of pastors for church performance. On the other hand, if the conference matches in the opposite direction (i.e., sending better pastors to “rescue” struggling churches), then our estimates for pastor value-added are biased toward zero.

While we cannot model this matching process directly, we present a pair of analyses that are useful in identifying a causal effect of pastors on performance.

a. Discrete changes in attendance around pastor rotations

We first examine whether there are discontinuous performance changes around pastor transitions. If a disproportionate fraction of a pastor’s effect on attendance coincides abruptly with his arrival (or departure), it would suggest that the pastor’s arrival – rather than other omitted factors – are responsible for the change. This does not eliminate the possibility of omitted variables biasing our results, but it does narrow the set of plausible candidates. In particular, it rules out a number of potentially endogenous determinants expected to trend more smoothly, such as secular trends in religiosity or demographic shifts.

Consider a simple variance decomposition of attendance growth. We calculate

$$\sum_{c=1}^C \sum_{s_c=\underline{S}_c}^{\overline{S}_c-1} \left(Attendance_Growth_{c,s_c+1} - \overline{Attendance_Growth}_{c,s_c+1} \right)^2$$

where we index each of our C individual churches with c . Because not every church spans the full sample period 1961 - 2003, we use a church-specific time index in the second summation, s_c , which starts when we first observe church c , $\underline{S}_c \geq 1961$, and ends when we last observe it, $\overline{S}_c \leq 2003$. $\overline{Attendance_Growth}_{c, s_c+1}$ is therefore simply the sample mean over all 17,289 church-year observations in our dataset. The sum of squared deviations may then be decomposed straightforwardly into the contributions from subsamples of pastor-year observations.

Table 4 shows the results when we split the years into two groups – those where we observe a pastor change, and those where we do not. The first column indicates that about 28 percent of church-year observations involve pastor rotation. The second column lists the fraction of overall variance accounted for by each group: 38 percent for observations with pastor changes, and 62 percent for observations without. Together, these estimates imply that the volatility in attendance growth is about $(38/62) \cdot (72/28) - 1 = 58$ percent higher in pastor-change years, compared to those where the pastor has not changed. Tests for heteroskedasticity – such as the Bartlett test, the Brown and Forsythe test, and the Levene test – all reject the null of equality of variances between the change and no-change years with p-values less than 0.0001.¹⁹

The fraction of variance explained by pastor-change years is similar if we look at a decomposition of *within-church* sum of squared deviations (i.e., where each church is given its own mean for attendance growth), as illustrated by the figures listed in the second set of columns in Table 4. Examining county population-adjusted attendance growth yields a near-identical decomposition as well.

¹⁹ A placebo test that shifts the transition years back by one year does not generate a significantly higher level of variance, highlighting the uniqueness of transition years in explaining the sample variation.

Together, the evidence in Tables 3 and 4 indicates that pastor arrivals, which are discrete events, are simultaneously accompanied by abrupt shocks to attendance growth, and that the directions of such changes are persistent within pastors.

We reiterate that this set of findings helps to rule out the possibility that the patterns we are attributing to pastor ability are instead the result of slow-moving trends such as demographic shifts or secular trends in an area's religiosity. It also helps to rule out spurious correlations from occasional "slumps" or "booms" that overlap with a pastor's tenure at a particular parish. We now turn to examine the persistence of performance using what we will argue is quasi-random assignment of pastors in their first placement out of seminary.

b. Rookie assignments

The prior analysis does not fully address the concern that our measure of pastor ability conflates both individual pastor quality and the effects of church-pastor matching. The ideal setting for generating an unbiased measure of pastor ability would be lottery-based random assignment, where "fit" with a congregation is not considered.

We approximate this setting by focusing on a pastor's initial placement, which we argue can be seen as largely random, for several reasons. First, initial assignments are dictated by the small set of openings that occur in entry-level congregations in a given year. Further, given the limited information that the Oklahoma Conference has at the time of hiring, there is limited scope for matching. There are several pieces of evidence that reinforce the notion that initial placements are quasi-random.

First, there is little evidence in general that managers are able to make fine-grained predictions of future workers' performance at the time of hiring. Jacobs and

Lefgren (2007), for example, find that principals correctly predict that a new teacher will have above-median performance 51 percent of the time, scarcely better than random guessing.²⁰

Second, we find that, based on observables, initial assignment is uncorrelated with performance at subsequent placements. In Figure 1, we divide pastor attendance growth residuals into terciles, based on all placements following *initial assignment*. We then assess whether these pastor-specific performance residuals in later placements are predictive of first-assignment church attributes. We find that future performance is uncorrelated with recent church growth at initial placement, and also with district-level population growth – both prior to and concurrent with initial placement. Finally, we also find that subsequent performance is uncorrelated with the level of attendance at initial assignment. This is particularly noteworthy, given results we will report below, which suggest that the Conference tries to place higher-quality pastors in larger congregations. In Appendix Table A1, we present regressions showing the correlations between performance after initial placement and initial placement attributes, which – as with Figure 1 – shows no discernible relationship.²¹

We note finally that in discussions with church leadership on the topic of pastor hiring that they of course attempt to recruit higher-quality pastors into their ranks. However, the notion of pastor-church matching at initial placement was not, in general,

²⁰ They do find, however, that principals are able to predict which teachers will perform at the extreme tails of the distribution.

²¹ One further concern might result from a pastor rejecting her initial assignment if assigned to a stagnant church. There are two conditions that would make this problematic for our identification: First, church-level growth rates must be foreseeable in order to make informed rejection choices. Second, unobservable future growth must be correlated with pastor quality. We expect that neither of these conditions hold – pastor ability is to a large extent unknown to both pastor and church at the time of assignment, and we also find that future growth rates of churches are hard to forecast.

a consideration. Overall, both the data as well as qualitative evidence are consistent with “almost” random matching at initial pastor placement.

We therefore take a pastor’s performance at his first church as a plausibly exogenous indication of his ability, and assess whether this estimated ability predicts performance at subsequent placements. We restrict attention to pastors who appear in the data after our sample begins in 1961, to focus on pastors where we observe performance at their first church assignments.²² We first run the following regression:

$$Attendance_Growth_{c,t} = \beta \cdot Controls_{c,t} + \varepsilon_{c,t} \quad (1)$$

for all churches, c , and years, t .²³ Control variables include church fixed effects, year fixed effects, and *County_Population_Growth*. We then calculate a vector of pastor-specific performance residuals, calculated *only* for the set of pastors for whom we observe their first assignments. Following the notation developed previously, our pastor ability measure is thus given by:

$$\hat{\gamma}_i = \frac{1}{K_{1i}} \sum_{k_{1i}=1}^{K_{1i}} \varepsilon_{c,k_{1i}}, \forall i$$

where the time index, k , applies only to the years in each pastor i ’s first assignment, i.e., where $a_i=1$. For example, suppose that a pastor’s initial church placement was in 1975, where he remained for three years. This pastor’s ability measure $\hat{\gamma}$ would simply equal the sample mean of the residual from (1), calculated from 1975-1977. Repeating this

²² We further screen out the four individuals with an “initial” placement in a church with greater than 500 in *Attendance* in the year of the pastor’s arrival. These cases correspond to seasoned pastors who, at some point prior to 1961, accepted an administrative role (e.g., a district superintendent), and have since returned to their roles as individual church pastors.

²³ In Equation (1), it is assumed that each church c only enters the estimation for years it exists in the sample, or following the notation developed earlier, when $\underline{S}_c \leq s_c \leq \overline{S}_c$. For ease of exposition, we thus use a common time index, t .

calculation for every initial placement, we obtain a set of quasi-exogenous ability measures that is less vulnerable to endogenous assignment concerns. We then use the vector $\hat{\gamma}_i$ to predict performance at each pastor's *subsequent* placements,

$$Attendance_Growth_{c,t} | (a_i > 1) = \alpha \hat{\gamma}_i + \beta \cdot Controls_{c,t} + \varepsilon_{c,t}. \quad (2)$$

Note that this estimation includes only non-initial placements, i.e., where $a_i > 1$, while our estimates of pastor ability from initial placements, $\hat{\gamma}_i$, serve as our main covariate.

The results of estimating Equation (2) appear in Table 5. Column (1) indicates a strong, positive relation between estimated first-assignment attendance growth, $\hat{\gamma}_i$, and *Attendance_Growth* in subsequent church placements, which is significant at the one-percent level. (Note that the reduction in sample size relative to earlier tables is due to the $a_i > 1$ restriction.) The estimated coefficient on $\hat{\gamma}_i$ implies an attendance growth elasticity of nearly five percent.

In columns (2) and (3) we divide the sample into those in the pastor's first year versus subsequent years at each church. As previously noted in the variance decomposition shown in Table 4, a disproportionate fraction of variation in attendance growth occurs in pastor-transition years. If this is attributable to pastor quality, we expect pastor ability to be most predictive of performance in these years. Our estimates bear out this prediction: the coefficient on $\hat{\gamma}_i$ doubles to 0.10 for the sample restricted to transition years (column (2)), significant at the one-percent level. This is more than six times larger than the comparable coefficient for the sample of non-transition years in column (3), where the coefficient on $\hat{\gamma}_i$ does not approach statistical significance.

Columns (4) and (5) show the results when we try to reduce measurement error in $\hat{\gamma}_i$, based on the premise that ability may be measured more precisely for pastors with

longer tenures at their initial placements. Accordingly, when we include only pastors where ability $\hat{\gamma}_i$ can be estimated with more than a single year of data (i.e., for all pastor i where $K_{1_i} > 1$), the coefficient of interest increases to 0.205, double the comparable estimate in column (2). Performance in later years remains unrelated to $\hat{\gamma}_i$.

Column (6) indicates that a pastor's early experience is more predictive of performance in small churches, though we note that the point estimate on the interaction of initial performance and church size is significant only at the 10 percent level; a comparison of first-year performances (columns 2 and 7) goes in the same direction, though the effect is smaller in magnitude and does not approach significance. Most initial placements are, of course, at small churches, so the fact that our ability measure provides only weak predictive power of performance at larger churches could result from the different skills required for churches of different sizes. However, given the noisiness of these estimates, any such conclusion needs to be made with caution.

We conclude this section with some robustness and falsification exercises. In the first column of Table 6, we present a placebo test based on $\log(1 + Deaths)$ as a "performance" metric. This is the dimension of church attendance growth that is most plausibly beyond the pastor's control. The results in Column (1) indicate that deaths at a pastor's first placement have no predictive power for deaths at subsequent churches – the coefficient on $\log(1 + Deaths)$ is 0.0135, with a standard error of 0.0175. In the remaining columns, we present analogous results for various alternative performance measures, including *Baptisms*, *Membership_Growth*, and *Revenue_Growth*. We find a significant level of persistence across all performance measures – albeit of varying

magnitudes – aside from *Revenue_Growth*.²⁴ The likely explanation for the lack of impact on revenues is that our proxy for church-level revenues – implied by figures on expenditures and capital growth – is quite noisy, as discussed in Section III.

V. Flow of human capital in the church

If individual pastor quality has a significant impact on church attendance, it is natural to investigate the *flow* of pastoral human capital, both across churches and out of the conference entirely. Whereas the decision to quit largely reflects individual tradeoffs, the conference is responsible for allocating pastors across churches, allowing us to infer the objectives of the conference generally.

a. Pastor exit

We present in Panel A of Table 7 the percent of pastors that exit the sample in year t , as a function of each quartile of the *Attendance_Growth* distribution in year $t-1$. In the first column, we show the results for the full sample. The bottom performance quartile indicates an exit rate of 12.5 percent, over four percentage points higher than the third quartile, which is higher than the exit rate in the second quartile. Exit rates appear similar among pastors in the top half (quartiles 3 and 4) of the distribution.

In Columns (2) and (3) we show the sample split by whether the pastor is a church elder, a higher position in the Methodist hierarchy which affords considerable job security. The elder pastor observations (column 2) are associated with much lower exit rates across all quartiles. However, it is noteworthy that the performance gradient exists in both subsamples, suggesting that there is an important role for self-selected exit by low-performing pastors. In column (4) we show the probability of exit from a

²⁴ The correlations among our various performance measures are all positive, but modest. For example, the correlations between *Attendance_Growth* and *Baptisms*, *Membership_Growth*, and *Revenue_Growth* are 0.15, 0.17, and 0.05 respectively. The correlation with $\log(1 + Deaths)$ is -0.02.

pastor's first church, and in (5), all subsequent churches. Note that in these columns we limit the sample to those pastors that enter the conference on or after 1961, so we can credibly identify their first churches. The exit rate is much higher during initial placements, where up to one fifth of the poorest performers quit, though once again the performance gradient is similar for both subsamples.

We formalize these univariate patterns in a multivariate regression in Table 8, based on a Cox proportional hazard model, stratified by church. In column (1), we predict a pastor's exit in year t as a function of his *Attendance_Growth* in year $t-1$. The reported hazard ratio is not significantly different from one at conventional levels (p-value=0.13), and its value, 0.72, implies that an increase in one-year lagged *Attendance_Growth* of 0.56 – the within-church average standard deviation – would increase the probability of exit by about 1.5 percent.

However, as Table 7 already indicated, the relation between past performance and exit is highly non-linear. Column (2) thus also shows the result of a hazard specification with a discrete indicator for performance, *High Growth*, which denotes whether *Attendance_Growth* in year $t-1$ is above the median, when measured across all churches that year. The effect of *High_Growth* is very large in magnitude; the estimate of 0.788 implies a 21 percent reduction in the probability of exit relative to the baseline hazard rate, and is significantly different from unity at the one-percent level. When we include both *Attendance_Growth* and *High_Growth* together in Column (3), the hazard ratio for *High_Growth* is almost unchanged, while the hazard ratio for *Attendance_Growth* becomes positive though not significantly different from one. (Adding *Elder* as a covariate to control for job security does not affect any of the hazard ratios for attendance measures, though *Elder* itself enters very significantly.)

We now examine how opportunities in the external labor market affect exit decisions, using oil prices to capture economic prospects from employment in secular professions. Oklahoma is home to some of the largest private oil interests in the U.S. The energy sector contributes between 10 and 20 percent of state GDP – an estimate that fluctuates with the price of oil – and thus serves as a plausible external shock to wealth and opportunity in the state. (See, for example, Wolfers (2012) for a discussion of oil price as an instrument for state-level economic shocks.) The simple pairwise correlation between pastor exit rates and oil prices is over 0.3, a relation that is illustrated in the time-series plots in Figure 2. The years 1980 and 2000 are particularly notable, when oil prices spiked sharply, coincident with similarly steep increases in pastoral exit rates.

Column (4) of Table 8 adds the price of oil as a predictor of a pastor’s exit probability. The point estimate of 1.011 (significant at the five-percent level) indicates a one percent increase in a pastor’s exit probability for every dollar increase in the price of oil. Put differently, a one-standard-deviation change in the price of oil (\$10.63) increases the probability that a pastor exits the Oklahoma Conference by eleven percent.

b. Pastor rotations

Next, we consider pastor rotations across churches within the Oklahoma Conference. The church may choose to give underperforming pastors a fresh start at a new parish, and any performance-rotation relationship may be reinforced by the fact that the pastor-parish relations committee at each church may request a new pastor if unsatisfied with the current match. Empirically, we begin by showing rotation rates by *Attendance_Growth* quartiles in the bottom panel of Table 7, to facilitate a comparison

to the exit rate patterns shown in the top panel. The comparison suggests an even stronger relation between performance and rotation, relative to the relation between performance and exit. Moreover, the performance-rotation relation appears monotonic across quartiles – high-performing pastors are less likely to rotate than those with intermediate performance. The patterns are quite similar across the sample splits shown in columns (2) – (5).

As with exit, we apply a hazard model to the rotation decision in the final columns of Table 8, using the same set of covariates as in the first set of columns. The continuous measure of *Attendance_Growth* is a strong predictor of pastor rotation (column (5)), significantly different from unity at the one-percent level. However, we find, as with exit, that rotation decisions are primarily sensitive to above average growth: the hazard ratio for the *High_Growth* dummy is significantly different from one in Column (6), and when both the linear and discrete growth measures are included in Column (7), the hazard ratio for *High_Growth* is unchanged, while the hazard ratio for *Attendance_Growth* becomes indistinguishable from one. We find a borderline significant effect of oil price on the hazard ratio, with an increase in oil price decreasing the likelihood of rotation.

Our results on rotation suggest that the church leaves well-functioning pastor-church matches intact. But given that pastor effects are correlated across church size, and that the importance of pastors appears to some degree scale-invariant, the conference may also choose to promote high-performing pastors to larger congregations where their abilities may be applied over a larger constituency. (To the extent that pastors value the prestige and perquisites that come with larger congregations, this has the added benefit of motivating pastors to grow attendance and membership.)

Table 9 examines this issue directly, linking the performance at a pastor's previous church to the size of the church to which he is currently assigned. We first average *Attendance_Growth* over all of the years for each pastor-church assignment a_i , allowing us to make comparisons between shorter and longer assignments. Then, we examine whether the average growth rate at assignment a_{i-1} predicts the size of a pastor's church a_i . To give a specific example, we relate the size (*Avg_Attendance*) of a pastor's fourth church to his average per-year *Attendance_Growth* in his third church (denoted as *Last Church Performance*). To avoid conflating the effects of a pastor's arrival on the size of his next church, we measure church size using attendance from the year prior to his arrival.

The first column shows a strong relation between a pastor's performance at his last church, and the size of his current assignment. We add fixed effects for pastor church placement number (a_i) and control for the logarithm of the time a pastor has appeared in the sample in columns (2) and (3) respectively, because a pastor is likely to be assigned to larger churches over time. The coefficient on lagged performance increases slightly in magnitude. (Given the relatively small number of observations for each church, if we include church fixed-effects, this saturates the model – the R^2 increases to nearly 0.9, and no variable is statistically significant.) The results suggest a significant role for past performance in predicting congregation size on the next placement. The coefficient on lagged average attendance growth in Column (1), with just year effects, is 0.26, significant at the 5 percent level. Given the standard deviation of lagged average performance of 0.14, this implies that a one standard deviation increase in prior performance increases the attendance at a pastor's next placement by about 3.6 percent.

In addition to improving pastor quality through strategic rotation and the attrition of lower-performing ones, the church may boost the performance of a given pastor through training or better-quality pastor-church matches over time. We examine this possibility by looking at how performance changes across churches, and over a pastor’s time in the conference. Specifically, Table 10 shows the results of a regression of *Attendance_Growth* on the logarithm of a_i , which we label as *Church_Number* in the table, and the logarithm of k_{a_i} (*Years_at_Church*). The first captures the effect of a pastor’s *general* experience through his successive church placements; the second measures the effect of a pastor’s *church-specific* experience. However, as Table 10 shows, in no case is either coefficient significant at conventional levels. These results are consistent with good pastors being “born not made.”

VI. Conclusion

Our analysis of the Methodist Church in Oklahoma reveals that the human capital of pastors is an important determinant of church attendance: replacing a 25th percentile pastor with one at the 75th percentile pastor increases annual attendance growth by three percent, similar to the effect of a 10 percent increase in the surrounding county’s population. We argue that the persistent influence of pastors on attendance growth is causal, based on a pair of tests designed to disentangle the effects of pastor ability from pastor-church matches driven by sources of variation in church prospects that are unobservable (to us): first-church performance is highly predictive of performance at future placements, and variance in attendance growth is concentrated in the first year of pastors’ placements.

The movement of pastors within the Oklahoma ministry is broadly consistent with the conference efficiently allocating labor resources across churches. Low-performing pastors are much more likely to be rotated, consistent with a model of pastor-church matching; further, high-performing pastors tend to be moved to larger congregations, while low-performing pastors are more likely to exit the sample.

Our findings emphasize the importance of considering both the production and consumption side in studying religion. For example, given the critical role of human capital in religious production and sensitivity of pastor exit to economic booms, one must be careful in attributing the effects of income shocks solely to demand-side considerations – as with other markets, analysis of religiosity needs to consider the simultaneous effects of shifts in both supply and demand.

Our inside-the-organization analysis of religious participation also provides micro-foundations for examining the organization of religious enterprise more broadly. This can help to inform analysis of the organization of churches within a religion, and even competition amongst religions, which may be an important input itself into overall participation rates (see, e.g., Finke and Starke (1988), Zaleski and Zech (1995)).

We focus in this study on the Methodist Church, where pastor allocation decisions are made centrally. One direction for further research may be to compare human resource decisions across denominations, comparing the effectiveness of decentralized approaches (e.g., Baptists and Presbyterians) to their centrally organized analogs (e.g., Methodists and Catholics). Such cross-denominational comparisons may thus help to assess whether, in the production of religion, organizational forms contribute to differences in performance. We leave this and similar questions to future work.

References

Abowd, J. M., Kramarz, F., and Margolis, D., 1999, "High Wage Workers and High Wage Firms," *Econometrica*, 67(2), 251-333.

Abowd, J. M., F. Kramarz, P. Lengeremann, and S. Perez-Duarte, 2004, "Are good workers employed by good firms? A test of a simple assortative matching model for France and the United States," working paper.

Azzi, C. and Ehrenberg, R., 1975, "Household Allocation of Time and Church Attendance," *Journal of Political Economy* 83, 27-56.

Becker, Sascha O. and Ludger Woessmann, 2009, "Was Weber Wrong? A Human Capital Theory of Protestant Economic History," *Quarterly Journal of Economics*, Vol. 124, No. 2, 531-596.

Becker, G., 1975, *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*, 2nd Edition; New York: National Bureau of Economic Research.

Bertrand, M. and Schoar, A., 2003, "Managing with Style: The Effect of Managers on Firm Policies," *Quarterly Journal of Economics* 118(4), 1169-1208.

BP Statistical Review of World Energy, June 2010, BP p.l.c., London, UK.

Branch, Gregory F., Eric A. Hanushek, and Steven G. Rivkin, "Estimating the Effect of Leaders on Public Sector Productivity: The Case of School Principals," 2012, NBER Working Paper 17803.

Chetty, R., Friedman, J., Rockoff, J., "The Long-Term Impacts of Teachers: Teacher Value-Added and Student Outcomes in Adulthood," 2012, working paper.

Ellison, Christopher G., 1991, "Religious Involvement and Subjective Well-being," *J. Health & Soc. Behav.*, 32:1, pp. 80-99.

Finke, R., Stark, R., 1988, "Religious Economies and Sacred Canopies: Religious Mobilization in American Cities," 1906, *American Sociological Review* 53:1, pp. 41-49.

----- 2005, *The Churching of America, 1776-2005: Winners and Losers in Our Religious Economy*.

Gillum, R., King, D., Obisesan, O., and Koenig, H., 2008, "Frequency of Attendance at Religious Services and Mortality in a U.S.," National Cohort, *Annals of Epidemiology*, 18(2), 124-129.

Goldberger, A., 1962, "Best Linear Unbiased Prediction in the Generalized Linear Regression Model", *Journal of the American Statistical Association*, 57, 369-375.

Gordon, R., Kane, T., Staiger, D., 2006, "Identifying Effective Teachers Using Performance on the Job," The Hamilton Project white paper 2006-01, the Hamilton Project, Washington, DC.

- Gruber, J. and Hungerman, D., 2008, "The Church vs. the Mall: What Happens When Religion Faces Increased Secular Competition?" *Quarterly Journal of Economics* 123, 831-862.
- Hamberg, E., Pettersson, T., 1994, "The Religious Market: Denominational Competition and Religious Participation in Contemporary Sweden," *Journal for the Scientific Study of Religion* 33:3, pp. 205-16.
- Hanson, G., Xiang, C., 2011, "Exporting Christianity: Governance and Doctrine in the Globalization of US Denominations," working paper.
- Hartzell, J., Parsons, C., Yermack, D., 2010, "Is a Higher Calling Enough? Incentive Compensation in the Church," *Journal of Labor Economics* 28, 509-540.
- Iannaccone, L., 1990, "Religious Practice: A Human Capital Approach," *J. Sci. Study Rel.*, 29:3, pp. 297-314.
- , 1992, "Sacrifice and Stigma: Reducing Free-Riding in Cults, Communes, and Other Collectives," *Journal of Political Economy* 100:2, pp. 271-97.
- , 1998, "Introduction to the Economics of Religion," *Journal of Economic Literature* XXXVI, p. 1465-1496.
- Jovanovic, B., 1979, "Job matching and the theory of turnover," *Journal of Political Economy* 87(5), 972 - 990.
- Kane, T., Rockoff, J., Staiger, D., 2008, "What Does Certification Tell Us About Teacher Effectiveness?: Evidence from New York City," *Economics of Education Review*, 27(6):615-631.
- Lazear, E., Shaw, K., Stanton, C., 2011, "The Value of Bosses," working paper.
- McCleary, R., and Barro, R., 2006, "Religion and Economy," *Journal of Economic Perspectives* 20 (2), 49-72.
- Mincer, Jacob. 1974. *Schooling, Experience and Earnings*. New York: National Bureau of Economic Research.
- Mollick, E. 2011, "People and Process, Suits and Innovators: The Role of Individuals in Firm Performance," forthcoming in *Strategic Management Journal*.
- Morris, C., 1983, "Parametric Empirical Bayes Inference: Theory and Applications," *Journal of the American Statistical Association*, 78:47-55.
- Neuman, Shoshana. 1986, "Religious Observance within a Human Capital Framework: Theory and Application," *App. Econ.*, 18:11, pp. 1193-2022.
- Raudenbush, S., Bryk A., 2002, *Hierarchical Linear Models: Applications and Data Analysis Methods*, Newbury Park, CA: Sage Publications.

- Rockoff, Jonah. 2004. "The Impact of Individual Teachers on Student Achievement: Evidence from Panel Data." *American Economic Review*, 94(2), 247-252.
- Rosen, S., 1981, "The Economics of Superstars," *American Economic Review* 71 (5), 845-858.
- Sawkins, J., Seaman, P., Williams, H., 1997, Church attendance in Great Britain: An ordered logit approach, *Applied Economics* 29, pp. 125-34.
- Simon, H. A., 1971, "Designing Organizations for an Information-Rich World", in Martin Greenberger, *Computers, Communication, and the Public Interest*, Baltimore, MD: The Johns Hopkins Press
- Smith, I., Sawkins, J., Seaman, P., 1998, The Economics of Religious Participation: A Cross-country Study, *Kyklos*, 51 (1), 25-43.
- Stark, R., Bainbridge, W., 1985, *The Future of Religion*, Berkeley, University of California Press.
- Weber, M. 1905. *The Protestant Ethic and the Spirit of Capitalism*. London: Allen & Unwin.
- Wolfers, Justin. 2012. "Are Voters Rational? Evidence from Gubernatorial Elections," working paper.
- Zaleski, P., Zech, C., 1995, The Effect of Religious Market Competition on Church Giving," *Rev. Soc. Econ.* 53:3, pp. 350-67.

Figure 1: Pastor Quality and Initial Placements

The figures groups pastors into terciles based on residual attendance growth, averaged over each pastor's career, excluding the years of his initial placement. The graph shows, for each tercile, the attributes of initial church assignment. Past church growth reflects the growth in attendance in the three years preceding a pastor's arrival at his initial placement. Past population growth is the county-level population growth in the three years preceding a pastor's arrival at his initial placement, while current population growth is country growths during the three years following a pastor's arrival. Church size is $\log(\text{attendance})$ is the logarithm of attendance in the year preceding a pastor's arrival at his initial placement. See text for details.

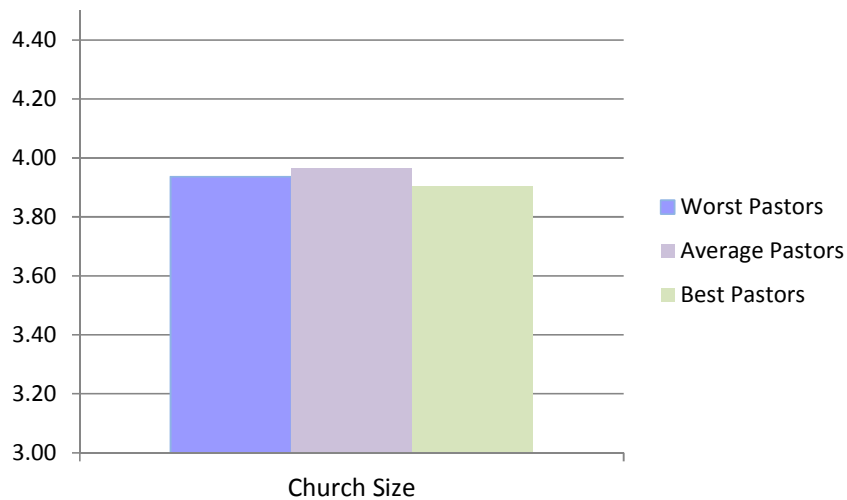
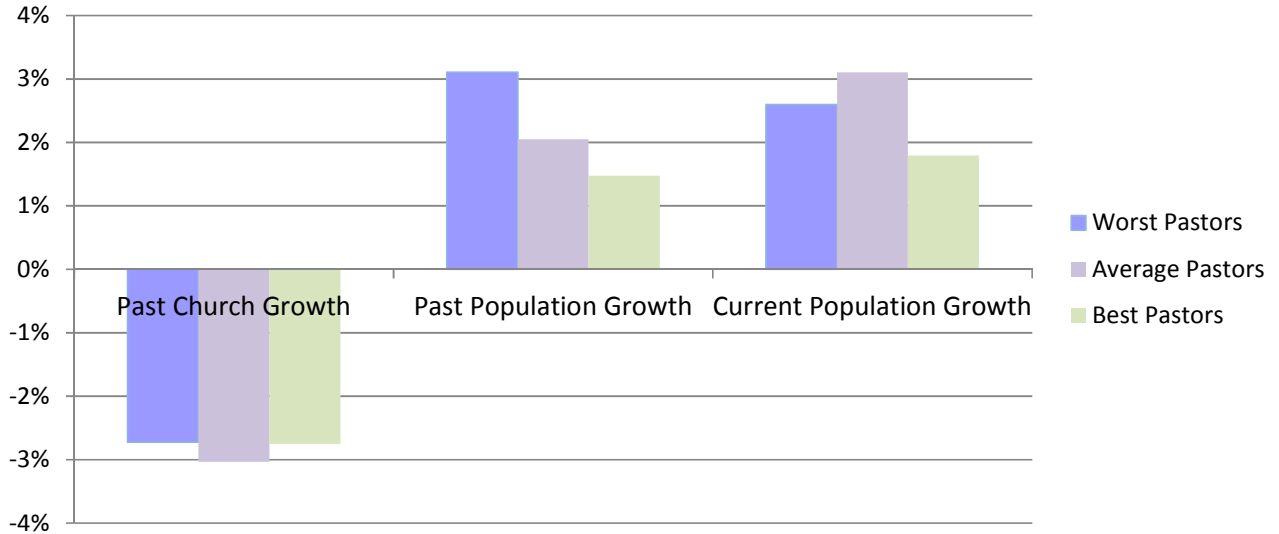


Figure 2: Pastor Exit and Oil Prices

The figure plots crude oil prices in 2008 dollars and the number pastors who leave the United Methodist Church in the state of Oklahoma between 1974 and 2002. Oil prices are from BP Statistical Review of World Energy, June 2010, BP p.l.c., London, UK. Pastor exits in a given year are the number of pastors who were in the dataset in current year but not in the following year.




Table 1: Summary Statistics

In the first panel, data by pastor is summarized. Unique Pastors in a Given Year is a yearly observation equal to the number of unique pastors we observe in the dataset in a given year. Pastor Tenure (Lifetime) is the number of years we observe the pastor in the dataset, given he begins after 1961 and before 1990. Pastor Tenure at Individual Church is the number of years a pastor works at an individual church, given he begins after 1961 and before 1990. Churches Worked by Pastor is the number of churches worked by a pastor who began working after 1961 but before 1990. Pastors Added in a Year is a yearly observation equal to the number of unique pastors we observe in the current year who did not work the previous year. Pastors Dropped out in a Year is a yearly observation equal to the number of unique pastors we observe in the previous year who are not working in the current year. The sample period for our dataset is from 1961 to 2003. In the second panel, church-level performance metrics are summarized for church c in year t . Average attendance is the average level of Sunday morning worship services. Attendance growth is calculated as the difference between the logarithm of the average attendance for a church in the current year and the logarithm of the average attendance for the same church in the previous year ($\Delta \log(Avg_Attendance_{c,t})$). The third row standardizes each church's attendance growth by county-level population growth using data from the US Census. Membership growth is the change in the logarithm of membership between the current year and the previous year ($\Delta \log(Membership_{c,t})$). Baptisms is calculated as $\log(1+Baptisms_{c,t})$. Revenue growth is the change in the logarithm of estimated revenues for church c in year t , $\Delta \log(Revenue_{c,t})$. Deaths is calculated as $\log(1 + Deaths_{c,t})$.

PASTORS

	Mean	St. Dev.	5th	25th	Median	75 th	95 th
Unique Pastors in a Given Year	458.93	18	430	447	460	471	492
Pastor Tenure (Lifetime)	8.77	8.66	1	2	5	14	26
Pastor Tenure at Individual Church	2.64	2.31	1	1	2	3	6
Churches Worked by Pastor (Lifetime)	3.28	2.83	1	1	2	5	10
Pastors Added in a Year	40.21	10.96	25	32	38	47	58
Pastors Dropped out in a Year	40.97	6.89	31	36	40	47	53

CHURCHES

Average Attendance	145.00	207	25	58	92	152	405
Attendance Growth	-1.30%	19%	-26%	-8%	0%	6%	22%
Population-Adjusted Attendance Growth	-1.90%	19%	-27%	-9%	-1%	5%	22%
Membership Growth	-0.20%	11%	-11%	-2%	0%	2%	8%
Baptisms (log)	1.58	1.08	0.00	0.69	1.61	2.30	3.33
Revenue Growth	-1.0%	52%	-67.9%	-16.5%	-0.8%	15.7%	64.6%
Deaths (log)	1.44	0.91	0.00	0.69	1.39	2.08	2.94

Table 2: Pastor Talent and Persistence

The two panels consider the distribution of attendance growth for pastor i at church assignment a_i given as a function of pastor i 's attendance growth at the immediately preceding church assignment, $a_i - 1$. The top panel considers unadjusted attendance growth while the bottom panel considers population adjusted attendance growth. The distribution of attendance growth is broken into quartiles with 1 (4) representing the lowest (highest) quartile.

Attendance Growth – Unadjusted (Panel A)

Church $a_i \setminus$ Church a_{i-1}	1	2	3	4
1 = Lowest	34%	28%	23%	15%
2	23%	33%	26%	18%
3	21%	23%	33%	23%
4 = Highest	16%	25%	29%	30%

Attendance Growth - Population Adjusted (Panel B)

Church $a_i \setminus$ Church a_{i-1}	1	2	3	4
1 = Lowest	33%	28%	23%	16%
2	23%	32%	27%	18%
3	20%	24%	33%	23%
4 = Highest	16%	26%	28%	29%

Table 3: Individual Pastors and Determinants of Attendance Growth

Panel A of this table reports the results from eight regressions where the dependent variable is attendance growth. Attendance growth is defined as the difference in log attendance for consecutive years at a church, $\Delta \log(\text{Avg_Attendance}_{c,t})$ for church c in year t . Local population growth is the difference in log population for the matched county. The first column reports the results when attendance growth is regressed on local population growth. Year fixed effects are added in column 2; church fixed effects are added in column 3, and pastor fixed effects are added in column 4. Columns 5 and 6 (7 and 8) consider the empirical Bayes specification among small (large) churches. Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

Panel B of this table reports the percentile distribution of the empirical Bayes estimates, with the estimated random effects derived using the best linear unbiased predictors. The sample is limited to churches with at least 20 years of data. The first row reports the results for the entire sample, while the second and third rows report results for small and large churches, respectively.

Dependent Variable: Church Attendance Growth (Panel A)								
					Small Churches		Large Churches	
County Population Growth	0.280*** (0.0556)	0.283*** (0.0652)	0.107 (0.0775)	0.0965 (0.0788)	0.117 (0.149)	0.0946 (0.168)	0.0716 (0.0967)	0.0254 (0.100)
Year Fixed Effects	NO	YES	YES	YES	YES	YES	YES	YES
Church Fixed Effects	NO	NO	YES	YES	YES	YES	YES	YES
Pastor Fixed Effects	NO	NO	NO	YES	NO	YES	NO	YES
Observations	15,357	15,357	15,357	15,357	4,770	4,770	6,973	6,973
Adjusted R ²	0.002	0.015	0.102	0.260	0.093	0.349	0.167	0.348

Percentiles of Empirical Bayes Coefficients (Panel B)									
	1st	5th	10th	25th	50th	75th	90th	95th	99 th
All Churches	-5.5%	-3.6%	-2.5%	-1.4%	0.0%	1.3%	2.6%	3.6%	5.1%
Small Churches	-7.9%	-5.5%	-4.2%	-2.3%	-0.5%	1.0%	2.8%	4.0%	6.0%
Large Churches	-5.0%	-3.1%	-2.1%	-0.9%	0.6%	2.1%	3.7%	5.2%	6.1%

Table 4: Pastor Changes and Attendance Growth

The two tables perform a variance decomposition for yearly church attendance changes. Attendance changes are defined as the difference in log attendance for consecutive years (Unadjusted) and the difference in log attendance for consecutive years minus the difference in log population for the zip-code matched area (Population Adjusted). The top panel, calculates the global mean (the average attendance change across all churches in all years) and then calculates a squared deviation from this global mean for each observation. The sum of squared errors (SSE) is calculated for years in which there was a pastor change and for years in which there was no pastor change. The bottom panel, calculates a church-specific mean (the average attendance change for a given church over its years) and then calculates a squared deviation from this church-specific mean for each observation. The sum of squared errors (SSE) is calculated for years in which there was a pastor change and for years in which there was no pastor change.

Variance Decomposition of Attendance Growth

		Population Adjusted	Unadjusted
	# of Obs (%)	SSE (%)	SSE (%)
Years with a Pastor Change	4,789 (28%)	225.72 (38%)	227.96 (38%)
Years without a Pastor Change	12,500 (72%)	373.74 (62%)	376.17 (62%)
TOTAL	17,289 (100%)	599.46 (100%)	604.13 (100%)

Variance Decomposition of Within-Church Attendance Growth

		Population Adjusted	Unadjusted
	# of Obs (%)	SSE (%)	SSE (%)
Years with a Pastor Change	4,789 (28%)	211.505 (38%)	214.10 (38%)
Years without a Pastor Change	12,500 (72%)	341.079 (62%)	344.95 (62%)
TOTAL	17,289 (100%)	552.58 (100%)	559.05 (100%)

Table 5: Initial Placement and Future Performance

The table reports the results from seven regressions where the dependent variable is residual church attendance growth:

$$Attendance_Growth_{c,t} | (a_i > 1) = \alpha \hat{\gamma}_i + \beta \cdot Controls_{c,t} + \varepsilon_{c,t}$$

Residual church attendance growth is defined to be attendance growth in all church assignments (in church c and year t) subsequent to the first placement, i.e., where $a_i > 1$. ($a_i = 1, 2, 3, \dots, A_i$ is the index for pastor i 's successive church assignments.) Pastor i 's performance at his first placement is used as the main regressor. It is given as:

$$\hat{\gamma}_i = \frac{1}{K_{1i}} \sum_{k_{1i}=1}^{K_{1i}} \varepsilon_{c,k_{1i}}, \forall i$$

This is a vector of pastor-specific performance residuals, calculated only for the set of pastors for whom their first assignments are observed. Additionally, the pastor-church specific time index, k , applies only to the years in each pastor i 's first assignment, i.e., where $a_i = 1$.

The sample is restricted to pastors who entered the sample after 1961. In the first column, Residual Church Attendance Growth is regressed on First Assignment Attendance Growth. In columns (2) and (3), the sample is divided into those in the pastor's first year versus subsequent years at each church. In columns (4) and (5), the sample is limited to pastors with tenures longer than one year at their first placement. In columns (1)-(5), log of time in sample, the number of years a pastor appears in the sample, is included as a covariate. In columns (6) and (7), initial placements at large churches are considered. The dummy Large Church is added, as well as the interaction term between First Church Attendance Change and Large Church. Large Church is set to one if the church has an average attendance that is greater than the median attendance that year. Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

Table 5: Initial Placement and Future Performance (continued)

	Dependent Variable: Residual Church Attendance Growth						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
First Assignment Attendance Growth	0.0482*** (0.0155)	0.102*** (0.0293)	0.0137 (0.0165)	0.205*** (0.0678)	0.0376 (0.0359)	0.0631*** (0.0222)	0.116** (0.0452)
Log(Time in Sample)	-0.00837** (0.00337)	-0.0102* (0.00537)	-0.00906** (0.00424)	-0.00467 (0.00871)	-0.00772 (0.00631)	-0.0159*** (0.00357)	-0.0138*** (0.00494)
First Church Attendance Change * Large Church						-0.0538* (0.0302)	-0.0540 (0.0576)
Large Church						0.0339*** (0.00451)	0.0307*** (0.00479)
Tenure Years at First Church		1 Year	> 1 Year	1 Year > 1 Year	> 1 Year > 1 Year		1 Year
Observations	4,429	1,583	2,846	953	1,834	4,429	2,787
R ²	0.016	0.040	0.020	0.056	0.021	0.032	0.036

Table 6: Other Performance Measures and Placebos

Table 6 reports the results of some robustness and falsification exercises. The sample is restricted to pastors who entered the sample after 1961. In column (1), deaths at a pastor's placements subsequent to the first one are regressed on deaths at the pastor's first placement. Deaths is calculated as the logarithm of (1+Deaths). In columns (2)-(5), other performance measures are used: baptisms, membership growth, net transfers, and revenue growth. Net transfer is equal to $\log(1+\text{transfers from other denominations}) - \log(1+\text{transfers to other denomination})$. Logarithm of time in sample, which is the number of years the pastor has appeared in the sample, is present in all regressions as a control. Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable:				
	Deaths	Baptisms	Membership Growth	Net transfer	Revenue Growth
Deaths in 1st Church	0.0135 (0.0175)				
Log(Time in Sample)	0.00682 (0.0147)	0.0158 (0.0192)	0.00186 (0.00156)	-0.00229 (0.0245)	0.00927 (0.00683)
Baptisms in 1st Year		0.0673*** (0.0238)			
Membership Growth in 1st Church			0.0750*** (0.0211)		
Net transfer in 1st Church				0.0706*** (0.0233)	
Revenue Growth in 1st Church					0.000509 (0.0157)
Observations	5,217	5,217	5,213	5,217	4,272
Adjusted R ²	0.005	0.012	0.012	0.009	0.004

Table 7: Univariate Evidence of Performance, Exit and Rotations

Panel A considers the likelihood of a pastor’s exit from the United Methodist Church Oklahoma. It tabulates the likelihood of a pastor’s exit in Year t given the attendance growth of his church in Year t-1. Panel B considers the likelihood of a pastor’s rotating from one United Methodist Church in Oklahoma to another. A pastor is said to have rotated churches if he is at a different church in Year t as he was in Year t-1. Pastor exits from the sample are coded as missing so as to distinguish the Panel B from Panel A. For both panels, attendance changes are defined as the difference in log attendance for consecutive years at a church. The first column of both panels considers all of the observations. The second (third) column considers the subset of observations for which the pastor is (is not) a church elder. The forth (fifth) column considers the subset of observations where the pastor is (is not) in his first church.

Probability of Exit in Year t					
Pastor Rank in Year t - 1	All Obs	Elder	Non-Elder	1 st Church	Not 1 st Church
1 = Lowest	12.5%	10.5%	16.8%	19.0%	11.0%
2	8.1%	6.4%	14.1%	15.1%	7.5%
3	7.5%	5.1%	12.6%	11.7%	6.6%
4 = Highest	7.6%	4.5%	12.9%	13.4%	6.3%

Probability of Rotation in Year t					
Pastor Rank in Year t - 1	All Obs	Elder	Non-Elder	1 st Church	Not 1 st Church
1 = Lowest	26.3%	27.7%	20.3%	19.2%	25.9%
2	25.1%	24.2%	22.3%	23.7%	23.6%
3	20.7%	19.0%	18.5%	18.4%	19.2%
4 = Highest	15.4%	14.4%	15.2%	15.4%	14.5%

Table 8: Regression Evidence of Performance, Exit and Rotations

The table considers the likelihood of pastor exit (first four columns) and rotation (last four columns) in a proportional hazard model. Pastor exit is when the pastor exits the sample; rotation is when a pastor switches churches. Independent variables are attendance growth, a dummy variable for above-median attendance growth (High Growth), and oil prices. Attendance growth is calculated as the difference between the logarithm of the average attendance for a church in the current year and the logarithm of the average attendance for the same church in the previous year. Oil prices are from BP Statistical Review of World Energy, June 2010, BP p.l.c., London, UK. . Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

Dependent Variable:								
	Exit	Exit	Exit	Exit	Rotate	Rotate	Rotate	Rotate
Attendance Growth	0.722 (0.211)		1.407 (0.462)		0.468*** (0.0611)		0.954 (0.160)	
High Growth (dummy)		0.788*** (0.0569)	0.730*** (0.0730)	0.781*** (0.0471)		0.733*** (0.0272)	0.740*** (0.0360)	0.739*** (0.0254)
Oil Price				1.011** (0.00433)				0.992* (0.00430)
Observations	15,476	15,476	15,476	15,476	15,476	15,476	15,476	15,476
Clustering				Year				Year

Table 9: Performance and Future Church Size

Table 9 reports the regressions of church size, the average attendance of the church of the pastor’s current assignment, on pastor *i*’s performance during his previous assignment. The independent variable, Last Church Performance, is attendance growth for each pastor during the final year of his most recent church assignment. In column (1), Church Size is regressed on Last Church Performance. In columns (2) and (3) fixed effects for pastor church placement number and tenure are added. In column (3), log of time in sample, the number of years a pastor is present in the sample, is added as a control. Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

Dependent Variable: Church Size			
Last Church Performance	0.257** (0.117)	0.261** (0.118)	0.258** (0.117)
Log(Time in Sample)			0.555 (0.422)
Church Number FE	NO	YES	YES
Tenure FE	NO	NO	YES
Observations	3,404	3,404	3,404
R ²	0.682	0.694	0.708

Table 10: Church Number and Performance

Table 10 shows the results of regressing attendance growth on $\log(\text{Church_Number})$ and $\log(\text{Years_at_Church})$. Attendance growth is used as the dependent variable; it is calculated as the difference between the logarithm of the average attendance for a church in the current year and the logarithm of the average attendance for the same church in the previous year ($\Delta\log(\text{Avg_Attendance}_{c,t})$). $\log(\text{Church_Number})$ is the logarithm of the number of church assignments the pastor has had, including the current assignment. $\log(\text{Years_at_Church})$ is the pastor-church time index, or the number of years the pastor worked at that assignment. Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

Dependent Variable: Performance			
$\log(\text{Church_Number})$	0.00593 (0.00920)		-0.00257 (0.0115)
$\log(\text{Years_at_Church})$		-0.0152* (0.00917)	-0.00587 (0.00408)
Observations	7,439	7,439	7,439
R ²	0.199	0.523	0.199

Table A1: Pastor Quality and Initial Placements – Not for publication

The table shows the output of regressions predicting attributes of pastors' initial assignments, regressed on pastor performance, defined as residual attendance growth, averaged over each pastor's career, excluding the years of his initial placement. Past church growth reflects the growth in attendance in the three years preceding a pastor's arrival at his initial placement. Past population growth is the county-level population growth in the three years preceding a pastor's arrival at his initial placement, while current population growth is country growths during the three years following a pastor's arrival. Church size is $\log(\text{attendance})$ is the logarithm of attendance in the year preceding a pastor's arrival at his initial placement. See text for details. Robust standard errors are in parentheses. *, **, and *** represent significance at the 10%, 5% and 1% levels, respectively.

	Dependent Variable:			
	Past Church Growth	Past Population Growth	Current Population Growth	Church Size
Pastor Performance After First Church	-0.0279	-0.0453	-0.000277	-0.0396
	-0.114	-0.00793	-0.0115	-0.16
Observations	319	303	377	431
R ²	0.108	0.402	0.484	0.137