Routledge
Taylor & Francis Group

Religiosity and income: a panel cointegration and causality analysis

Dierk Herzer[a] and Holger Strulik[b]

[a]Department of Economics, Helmut-Schmidt-University Hamburg, Hamburg, Germany; [b]Department of Economics, University of Göttingen, Göttingen, Germany

ABSTRACT

In this article, we examine the long-run relationship between religiosity and income using retrospective data on church attendance rates for a panel of countries from 1930 to 1990. We employ panel cointegration and causality techniques to control for omitted variable and endogeneity bias and test for the direction of causality. We show that there exists a negative long-run relationship between the level of religiosity, measured by church attendance, and the level of income, measured by the log of GDP per capita. The result is robust to alternative estimation methods, potential outliers, different samples, different measures of church attendance and alternative specifications of the income variable. Long-run causality runs in both directions, higher income leads to declining religiosity and declining religiosity leads to higher income.

I. Introduction

Secularization, broadly understood, describes the phenomenon that later born generations are less religious. It is revealed, for example, by lower weekly attendance rates at church (see e.g. Wilson 2003; Voas 2009). Secularization appears to be a global phenomenon. According to the latest Win-Gallup (2012) poll, the global average of the religiosity index declined by 9% from 2005 to 2011. Religiosity declined almost everywhere in Europe and the United States (by 13%) and also in some developing countries. Norris and Inglehart (2004) conclude 'During the twentieth century in nearly all post-industrial nations – ranging from Canada and Sweden to France, Britain, and Australia – official church records report that where once the public flocked to Sabbath worship services, the pews are now almost deserted'.

A narrow definition of secularization assigns a cause to this trend, inspired by the observation that secularization occurred in the Western world in conjunction with the Industrial Revolution and the takeoff of economic growth (Azzi and Ehrenberg, 1975; Dobbelaere 1987; Norris and Inglehart 2004; Bruce 2011). According to this so-called secularization hypothesis, improving economic conditions have caused the decline in religiosity and the demand for religious services. The literature proposes different mechanisms

for secularization, which are certainly not mutually exclusive: the elevated sense of self-importance (rising individualism) brought about by progress in science and technology (Bruce 2011), the increasing opportunity costs of leading a puritanical lifestyle (Strulik 2016a), the increasing existential security induced by the takeoff from subsistence towards perpetual income growth (Norris and Inglehart 2004) and the increasing supply of secular consumption activities (Hirschle 2011; Strulik 2016b). (See Bruce (2011) for a detailed analysis.) However, there are also staunch critics of the secularization hypothesis (e.g. Stark 1999; Iannaccone 1998) and many researchers consider the question whether economic prosperity induces religious decline as not yet settled.

Conceptually, the secularization hypothesis describes a process, an inherently dynamic phenomenon where improving income leads to gradually declining religiosity, which may in turn affect income development. So far, however, the dynamic nature of secularization has not been taken fully into account. A rigorous time series analysis was impossible because long-run data on religiosity were, until recently, not available. In this article, we try to improve the state of affairs by applying advanced cointegration and causality analysis on a new long-run data set of income and religiosity for a panel of Christian countries.

CONTACT Dierk Herzer ✉ herzer@hsu-hh.de Department of Economics, Helmut-Schmidt-University Hamburg, Holstenhofweg 85, 22043 Hamburg, Germany

As a measure of religiosity, we use national church attendance rates between 1930 and 1990 from Iannaccone (2003). The novel idea behind Iannaccone's study is to estimate historical church attendance rates using contemporary ISSP questionnaires and replies to inquires on church attendance when the respondents were 11 or 12 years old. Naturally, such a retrospective method is prone to different sorts of bias (e.g. age effects or projection bias). As such, Iannaccone therefore devotes the greater part of his study to carefully demonstrate that there is no reason for concern. The data are robust to numerous tests of internal and external consistency. We combine the retrospective attendance data, which are available at 5 year intervals, with data on GDP per capita from Maddison (2003).

Applying panel cointegration estimation techniques, we find a strong negative association between income and church attendance. According to our benchmark specification, about 89% of the decline in church attendance over time can be attributed to the increase in per capita income. We show that this result is robust to a variety of sensitivity tests, including the use of alternative estimation techniques, the exclusion of potential outliers, the use of different samples, different measures of church attendance and alternative specifications of the income variable. Most importantly, the cointegration results are known to be robust to a variety of estimation problems that often plague empirical work, including omitted variables and endogeneity (as discussed in more detail in Section II).

The second contribution of our article is to address the question of temporal causality: is income an exogenous determinant of church attendance or is the decrease in church attendance both a cause and a consequence of increased income over time? To answer this question, we use Granger causality tests and impulse response functions techniques that were built upon the idea that the cause occurred before the effect. We find that long-run causality runs in both directions: higher income leads to declining religiosity and declining religiosity leads to higher income.[1]

A few quantitative studies have investigated the income – religiosity nexus. McCleary and Barro (2006a) found a negative effect of income on religious participation and beliefs across countries, which is also quantitatively important. A one standard deviation increase in log GDP per capita decreases church attendance by 15%. McCleary and Barro (2006b) arrived at similar results. The latter study also documents a significant negative impact of church attendance on economic growth. Paldam and Gundlach (2013) used the World Value Survey to compile a religiosity index (14 items from 'God is very important in life' to 'Churches answer spiritual needs') and found across countries a negative impact of income on religiosity. On average, religiosity falls by 50% when countries pass through the transition from underdevelopment to becoming a developed country. Lipford and Tollison (2003) documented a bi-causal negative association of income and religious participation across US states. Rupasingha and Chilton (2009) used US county-level data on religious adherence and found a causal negative effect of adherence on economic growth. An increase of religious adherence by one standard deviation would reduce growth by 0.4% per year. In contrast, Franck and Iannaccone (2014) did not find a causal effect of income on church attendance in a panel of 10 industrialized countries over the period between 1925 and 1990, and Becker and Woessmann (2013) did not find a causal effect of (teacher-) income on church attendance in Prussia between 1886 and 1911.

At first sight, it may appear puzzling that our results disagree with those of Franck and Iannaccone (2014) who also used the Iannaccone (2003) data set. One possible explanation is that their panel estimation method implicitly assumes that church attendance and income are stationary in the time-series dimension, although both income and church attendance are dynamic phenomena that exhibit trends. This could be a problem because we know from a large and growing literature on nonstationary panel data that conventional panel data models that ignore the potential non-stationarity of the variables may produce misleading results (see e.g. Entorf 1997; Kao 1999; Eberhardt and Teal 2013). Our results show that the non-finding by Franck and Iannaccone (2014)

[1]The impact of religion on income is a priori unclear. On the one hand, declining religiosity could be harmful for economic growth because it discourages trust (Guiso, Sapienza, and Zingales 2003) or because it reduces incentives to accumulate human capital (Becker and Woessmann 2009). On the other hand, declining religiosity could be conducive to growth because the turn towards material values and the promised pleasure from consumption induces increasing labour supply and capital accumulation (Lipford and Tollison 2003; Strulik 2016a).

disappears when the econometric implications of non-stationarity are properly taken into account.

The article is organized as follows. In Section II, we set up the basic empirical model and discuss some econometric issues (Section 'Basic empirical model and econometric methodology'). We then describe the data and report summary statistics, including pretests for unit roots and cointegration (Section 'Data, unit root and cointegration tests'). In Section III, we present the empirical analysis. In Section 'Panel cointegration estimates', we provide estimates of the cointegrating relationship between religiosity and income, in Section 'Robustness', we test the robustness of the estimates and in Section 'Long-run causality', we investigate the direction of long-run causality between the two variables. We conclude in Section IV.

II. Model and data

Basic empirical model and econometric methodology

Following common practice in (panel) cointegration studies (see e.g. Pedroni 2007; Herzer 2008; Moscone and Tosetti 2010), we consider a static levels model that includes only the two variables of empirical interest: religiosity and income. Specifically, the basic model takes the form

$$CHURCH_{it} = a_i + \beta \log(y_{it}) + \epsilon_{it} \qquad (1)$$

where $i = 1, 2, \ldots, N$ is the country index, $t = 1, 2, \ldots, T$ is the time index and the a_i are country-specific fixed effects. Following previous studies, we use church attendance as our measure of religiosity ($CHURCH_{it}$), while income is measured by real GDP per capita (y_{it}). As discussed in more detail below, the measure of church attendance is the percentage of parents (or children) who attend religious services at least once per week. The income variable is logged, as in most previous studies (see e.g. McCleary and Barro, 2006a, 2006b; Franck and Iannaccone 2014; Becker and Woessmann 2013), but in the Robustness section, we also estimate the model using the non-log-transformed income variable, as in Lipford and Tollison (2003).

Equation (1) assumes that, in the long run, permanent changes in the level of (log) GDP per capita are associated with permanent changes in the level of church attendance. The long-run marginal effect of $\log(y_{it})$ on $CHURCH_{it}$ is the partial derivative of Equation (1) with respect to $\log(y_{it})$: $\partial CHURCH_{it}/\partial \log(y_{it}) = \beta$. Thus, the coefficient β measures the approximate change in church attendance due to one unit change in log per capita GDP, implying that $\beta/100$ is (approximately) the unit change in $CHURCH_{it}$ when y_{it} increases by 1%.

Most scholars of Christian religion consider church attendance to be a good measure of religiosity (Brierley 1999; Bruce 2011; Warner 2010). Attending the Sunday service was obligatory for Protestants of most denominations as well as for members of the Anglican Church and deliberate failure to attend Sunday mass is still considered to be a grave sin according to Catholic canon law. In the economists' language, church attendance is a revealed preference. Since church is the only place where sins can be forgiven, not attending shows a loss of interest in salvation and eternal life, the main goods offered by the church. Attendance is strongly but not perfectly correlated with adherence. In particular, attendance declines somewhat before membership declines (Bruce 2011). On the other hand, church attendance may overestimate religiosity because attending could be driven by secular motives like social capital accumulation. Moreover, there may exist a temporary phase of 'fuzzy fidelity' in which individuals claim that religion is not important in their life while they continue to appear in church occasionally (Voas 2009). In any case, attendance is the most widely available and most frequently used religiosity variable.

We now turn to econometric issues. The first observation is that the underlying variables are trended; they are non-stationary (as shown in Figures 1 and 2). Given that most economic time series are characterized by stochastic rather than deterministic non-stationarity, it is plausible to assume that the trends in $CHURCH_{it}$ and $\log(y_{it})$ are also stochastic – through the presence of a unit root – rather than deterministic – through the presence of polynomial time trends. If this assumption is correct, the linear combination of these integrated (or stochastically trending) variables must be stationary, or, in the terminology of Engle and Granger (1987), $CHURCH_{it}$ and $\log(y_{it})$ must be cointegrated for a linear long-run relationship between
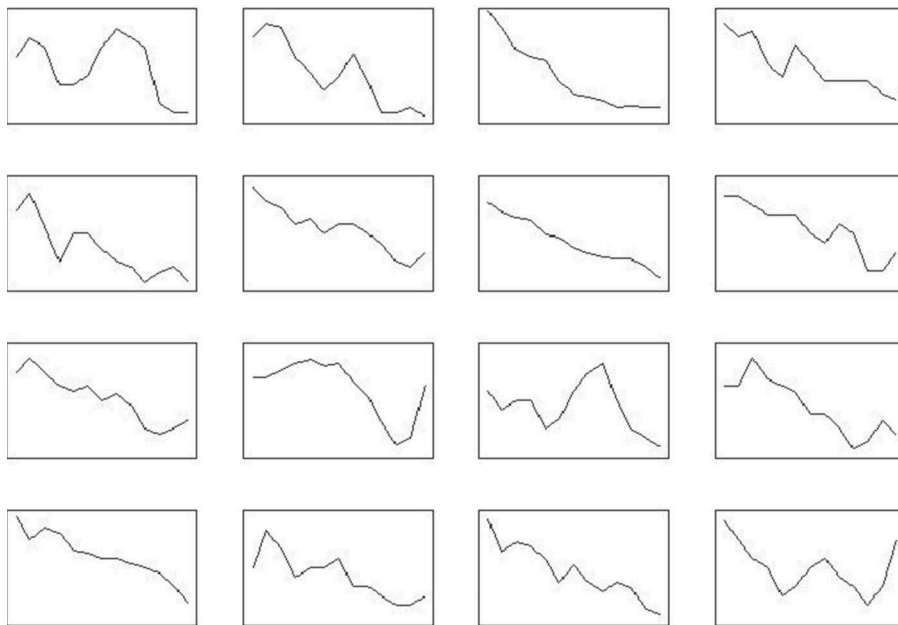
The countries from left to right are: Australia, Austria, Bulgaria, Chile, Denmark, France, Germany, Ireland, Italy, Netherlands, New Zealand, Norway, Spain, Sweden, UK, and the United States

Figure 1. Church attendance of parents over the period 1930–1990.

these two variables (as given in Equation (1)) to exist.[2] If the two variables are not cointegrated, then there is no long-run relationship between religiosity and income, and Equation (1) would be a spurious regression in the sense of Granger and Newbold (1974). Standard regression output must therefore be treated with extreme caution when variables are non-stationary, since the estimated results are potentially spurious (see also Eberhardt and Teal 2013). As shown by Entorf (1997) and Kao (1999), the tendency for spuriously indicating a relationship may even be stronger in panel data regressions than in pure time-series regressions. Thus, the necessary condition for the existence of a non-spurious long-run relationship between $CHURCH_{it}$ and $\log(y_{it})$ is that the two integrated variables cointegrate.[3]

A regression consisting of cointegrated variables has the property of super-consistency such that coefficient estimates converge to the true parameter values at a faster rate than they do in standard regressions with stationary variables, namely at rate T rather than $\sqrt{T}$ (Stock 1987). The important point in this context is that the estimated cointegration coefficients are super-consistent even in the presence of temporal and/or contemporaneous correlation between the stationary error term, ϵ_{it}, and the regressor(s) (Stock 1987), implying that cointegration estimates are not biased by omitted *stationary* variables (see e.g. Bonham and Cohen 2001).

The fact that a regression consisting of cointegrated variables has a stationary error term also implies that no relevant *non-stationary* variables are omitted. Any omitted non-stationary variable that is part of the cointegrating relationship would become part of the error term, thereby producing non-stationary residuals, and thus leading to a failure to detect cointegration (Everaert 2011).

If there is cointegration between a set of variables, then this stationary relationship also exists in extended variable space. In other words, the

[2]As defined by Engle and Granger (1987), two variables are cointegrated [of order (1, 1)] if each variable individually is stationary in first differences (integrated of order 1), but some linear combination of the variables is stationary in levels (integrated of order 0). The conventional concept of cointegration between variables with stochastic trends is thus defined as the existence of a linear relationship between these variables over time that produces stationary residuals. Evidence of cointegration (in the usual linear sense) therefore implies the absence of significant nonlinearities in the estimated relationships, whereas a failure to find (linear) cointegration does not necessarily mean that there is no (nonlinear) long-run relationship among the variables (see e.g. Kanas 2005).

[3]The standard time-series approach is to first difference the variables to remove the non-stationarity in the data and to avoid spurious results. However, this approach precludes the possibility of a long-run or cointegrating relationship in the data and leads to misspecification if a long-run relationship between the levels of the variables exists (see e.g. Granger 1988).
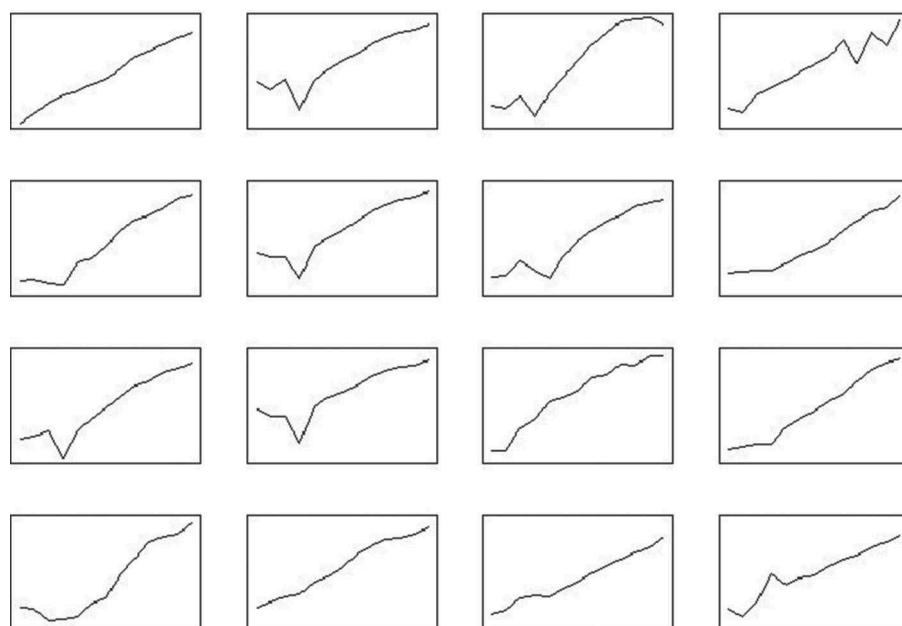
The countries from left to right are: Australia, Austria, Bulgaria, Chile, Denmark, France, Germany, Ireland, Italy, Netherlands, New Zealand, Norway, Spain, Sweden, UK, and the United States

Figure 2. Log of GDP per-capita by country over the period 1930–1990.

cointegration property is invariant to model extensions (see also Lütkepohl 2007), which is in stark contrast to regression analysis where one new variable can alter the existing estimates dramatically (Juselius 2006, 11). Thus, the important implication of finding cointegration is that no additional variables are required to account for the classical omitted variables problem because such a problem does not exist under cointegration; the result for the long-run relationship between religiosity and income would also hold if we included additional independent variables in the model (see also Juselius 1996).

Of course, there are several other factors such as education, fertility and government expenditure that may be associated with religiosity and/or income. Therefore, adding further non-stationary variables to the model may, on the one hand, result in further cointegrating relationships. These, however, would have to be identified and estimated (separately). In particular, the difficulty is that if there is more than one stationary linear combination of the variables, identifying restrictions is required to separate the cointegrating vectors. On the other hand, adding further non-stationary variables to the regression model may result in spurious associations. More specifically, if a non-stationary variable that is not

cointegrated with the other variables is added to the cointegrating regression, the error term will no longer be stationary. As a result, the coefficient of the added variable will not converge to zero, as one would expect of an irrelevant variable in a standard regression (Davidson 1998). This justifies a reduced form model such as Equation (1), given the variables are cointegrated.

The super-consistency of the cointegration estimation also implies that the potential endogeneity of the regressors does not affect the estimated long-run coefficients; the estimated long-run coefficients from reverse regressions should be approximately the inverse of each other due to the super-consistency (Engle and Granger 1987). Nevertheless, there are two problems.

First, although the standard least-squares dummy variable estimator is super-consistent under panel cointegration, it suffers from a second-order asymptotic bias arising from serial correlation and endogeneity in finite samples and as a consequence, its *t*-ratio is not asymptotically standard normal (see e.g. Kao 1999). To deal with this problem, one has to employ an asymptotically efficient (cointegration) estimator. Examples of such estimators include panel versions of the dynamic OLS (DOLS) and fully

modified OLS (FMOLS) methods. As shown by Wagner and Hlouskova (2010), the panel DOLS estimator of Mark and Sul (2003) outperforms other asymptotically efficient estimators. Therefore, this estimator is preferred here, but in the Robustness section, we also present results based on alternative estimation procedures.

Second, although the existence of cointegration implies long-run Granger causality in at least one direction (Granger 1988), cointegration says nothing about the direction of the causal relationship between the variables. A statistically significant cointegrating relationship between $CHURCH_{it}$ and $\log(y_{it})$ does therefore not necessarily imply that, in the long run, changes in income cause changes in religiosity. The causality may run in the opposite direction, from $CHURCH_{it}$ to $\log(y_{it})$, or in both directions. The empirical implication is that it is important not only to employ an asymptotically efficient cointegration estimator (to account for the potential endogeneity of income), but also to explicitly test the direction of long-run causality. As is common practice in testing long-run Granger causality between cointegrated variables, we use a vector error correction model to identify cause and effect in the sense of Granger (1988).

A final econometric issue is the potential cross-sectional dependence (CD) in the regression errors due to common shocks or spillovers among countries at the same time. Standard panel (cointegration) techniques assume cross-sectional independence and may be biased if this assumption does not hold. Therefore, we test for CD in the residuals of the estimated models using the CD test developed by Pesaran (2004). In cases where we find evidence of CD, we employ demeaned variables to control for common effects; that is, in place of $CHURCH_{it}$ and $\log(y_{it})$, we use

$$CHURCH'_{it} = CHURCH_{it} - \overline{CHURCH_t},$$
$$\overline{CHURCH_t} \equiv N^{-1} \sum_{i=1}^{N} CHURCH_{it} \qquad (2a)$$

$$\log(y_{it})' = \log(y_{it}) - \overline{\log(y_t)},$$
$$\overline{\log(y_t)} \equiv N^{-1} \sum_{i=1}^{N} \log(y_{it}), \qquad (2b)$$

which is equivalent to including time dummies in the model. Moreover, we use a battery of panel unit root and cointegration tests, including so-called second-generation panel unit root and cointegration methods that explicitly allow for CD.

Data, unit root and cointegration tests

Data on real per capita GDP are taken from Maddison (2003), available at http://dx.doi.org/10.1787/456125276116. Data on church attendance are from Iannaccone (2003), who uses retrospective survey questions on weekly church attendance rates for respondents and their parents from the International Social Survey Program to construct average weekly church attendance rate of parents and children in 32 countries between 1925 and 1990. Attendance of parents is our preferred measure of religiosity since religious commitments typically develop during adolescent years rather than during early childhood. The attendance rate for children is used in the Robustness section. The data are expressed as a percentage of (the parents of) the respondents.

Given that the data set of Iannaccone (2003) spans a long time period, it is inappropriate to use conventional small T panel data models, which ignore the potential non-stationarity of the variables (see e.g. Phillips and Moon 2000). The appropriate approach is to use panel time-series techniques to account for the time-series properties of the variables and to avoid spurious results by testing for cointegration. Cointegration estimates are not only robust to omitted variables and endogenous regressors (as discussed above) but also robust to non-systematic measurement errors (Stock 1987). The latter is an important advantage for applications such as the present one, because it is likely that church attendance rates based on respondents' self-reports are measured with error.

The data on church attendance are available only every 5 years. That we are forced to use 5-year data for our analysis should not be a serious problem because panel cointegration methods exploit both the time-series and cross-sectional dimensions of the data and can therefore be implemented with a smaller number of time-series observations than their time-series counterparts. The panel cointegration analysis by Madsen, Saxena and Ang (2010), for example, is based on $T = 13$; accurate critical values for panel unit roots and cointegration tests are available even for $T = 10$ (see e.g. Pesaran 2007; Banerjee

and Carrion-i-Silvestre 2011). Moreover, it is well known that the total length of the sample period, rather than the frequency of observation, is the important factor when analysing the integration and cointegration properties of variables (see e.g. Shiller and Perron 1985; Hakkio and Rush 1991; Lahiri and Mamingi 1995). In addition, several studies show that cointegration estimates are remarkably stable across frequencies (see e.g. Chambers 2001; Click and Plummer 2005; Herzer 2013).

However, a potential disadvantage is that the estimators we use are designed for balanced panels, while the underlying data sets are unbalanced in the sense that the number of time-series observations per country varies. To construct a balanced panel, which entails a trade-off between the time span and number of countries in the sample, we select all Christian countries for which complete time-series data are available over the period 1930–1990 (a reasonably long period to conduct cointegration analysis).[4] This sample selection procedure yields a sample of 16 countries and 13 time-series observations per country (208 total observations). In the Robustness section, we also estimate the long-run relationship between $CHURCH_{it}$ and $\log(y_{it})$ using different samples over different time intervals ($N = 11$, $T = 14$ and $N = 19$, $T = 12$). Table 1 lists the countries in our main sample along with the average values for $CHURCH_{it}$ and $\log(y_{it})$ over the period 1930–1990. The United States had the highest GDP per capita, while Bulgaria had the lowest GDP per capita. Church attendance was highest in Ireland and lowest in Denmark.

In Figures 1 and 2, $CHURCH_{it}$ and $\log(y_{it})$ are plotted for each country over the observation period. While GDP per capita increased in all countries, church attendance decreased in all countries between 1930 and 1990. The downward trend in church attendance is especially strong in Austria, Bulgaria, Chile, Denmark, France, Germany, Ireland, Italy, Norway, Spain, Sweden and the UK where church attendance

Table 1. Countries and summary statistics.

	Average of $CHURCH_{it}$	Average of $\log(y_{it})$
Australia	31.85	9.13
Austria	50.46	8.75
Bulgaria	19.54	7.93
Chile	42.38	8.33
Denmark	12.08	9.12
France	31.31	8.93
Germany	26.65	8.94
Ireland	94.54	8.50
Italy	60.69	8.71
Netherlands	53.23	9.02
New Zealand	31.31	9.11
Norway	15.00	8.94
Spain	53.85	8.36
Sweden	13.08	9.10
UK	26.38	9.11
United States	55.92	9.38

declined relatively steadily, compared to Australia, the Netherlands, New Zealand and the United States where church attendance rates exhibit higher volatility. Overall, the time-series evolution is consistent with the possibility that $CHURCH_{it}$ and $\log(y_{it})$ are driven by stochastic trends.

In order to investigate this issue formally, we conduct panel unit root tests. In recent years, a number of panel unit root tests have been developed. The most commonly used are the so-called first-generation panel unit root tests, such as the ADF-Fisher-type test of Madalla and Wu (1999) (MW), the Breitung (2000) test, the Levin, Lin and Chu (2002) test and the Im, Pesaran and Shin (2003) test. Hlouskova and Wagner (2006) find that the Breitung panel unit root test generally has the highest power and smallest size distortions of any of the first-generation panel unit root tests. Therefore, we use the Breitung test. Given, however, that the first-generation tests, which assume cross-sectional independence, exhibit severe size distortions in the presence of CD, we also use second-generation panel unit root tests to allow for CD. More specifically, we use the panel unit root tests developed by Breitung and Das (2005) and Pesaran (2007). The Breitung and Das test is an extension of the Breitung test and is based on modified standard errors that are robust to CD. A potential disadvantage of the cross-sectionally robust

---

[4]Overall, time-series data are available for the period 1930–1990 (1925–1990) for 16 (11) Christian countries and one non-Christian country, Japan. Non-Christian countries, especially Buddhist countries (such as Japan), tend to have systematically lower church attendance rates than Christian nations (see e.g. Iannaccone 2003), primarily due to the fact that Christians attend worship services as an integral part of their religious practice, whereas other religious traditions, especially Buddhism and Hinduism, are less, or not at all, organized around regular worship services. The implication is that church attendance is not a consistent measure of religiosity across countries with different religious traditions. Although this is less of a problem for our time-series panel analysis than for studies focusing mainly on the cross-sectional variation in the relationship between religiosity and income, we therefore decided to include only Christian countries and, hence, to exclude Japan. As shown in the working paper version of this study (Herzer and Strulik 2013), the results do not change qualitatively when Japan is included.

test is that it requires $N < T$, which prevents us from using our main sample (with $N = 16$ and $T = 13$). Therefore, we apply the Breitung–Das procedure to an alternative sample of 11 countries over 14 time periods (a sample that also will be used in the Robustness section), while the Pesaran panel unit root test is again applied to the main sample. Like the Breitung test and the Breitung–Das test, the Pesaran test is an ADF-type test. It is based on an average of the individual country ADF t-statistics and filters out the CD by augmenting the individual ADF regressions with the cross-sectional averages of lagged levels and first differences of the individual series.

As can be seen from Table 2, the Breitung test and the Breitung and Das test fail to reject the null hypothesis of a unit root at the 10% level for both variables. The Pesaran test does not reject the unit root null hypothesis for $\log(y_{it})$ at the 10% level, while the unit root null for $CHURCH_{it}$ is rejected at the 10% level but not the 5% level. Since none of the tests reject the unit root null at the 5% level (despite the fact that panel unit root tests have high power), we conclude that both $CHURCH_{it}$ and $\log(y_{it})$ are driven by stochastic trends.

In order to ensure that the relationship between $CHURCH_{it}$ and $\log(y_{it})$ is not spurious, we test for cointegration using the standard panel and group ADF and PP test statistics suggested by Pedroni (1999, 2004). However, these tests do not take account of potential CD in the errors, which could bias the results. To test for cointegration in the presence of possible CD, we also employ the panel cointegration test recently proposed by Banerjee and Carrion-i-Silvestre (2011). This test involves four steps. The first is to estimate the parameters of the cointegrating regression using the pooled common correlated effects (CCE) estimation technique

Table 2. Panel unit root tests.

Variable	Breitung (2000)	Breitung and Das (2005)	Pesaran (2007)
$CHURCH_{it}$	1.772	0.769	−2.837*
$\log(y_{it})$	−1.000	−0.742	−2.606

The panel unit-root tests are ADF-type tests. The models allow for individual-specific intercepts and time trends. Given the relatively small number of time series observations, only one lag was used to adjust for autocorrelation. The Breitung and the Breitung and Das (2005) statistics are asymptotically distributed as a standard normal. The relevant 10% (5%) [1%] critical value for the Pesaran statistic with an intercept and a linear trend is −2.71 (−2.85) [−3.14] (when $N = 16$ and $T = 13$). The critical values were calculated from the response-surface estimates in Otero and Smith (2013). *Indicates significance at the 10% level.

Table 3. Panel cointegration tests.

Pedroni (1999, 2004) panel ADF statistic	−5.306***
Pedroni (1999, 2004) group ADF statistic	−7.799***
Pedroni (1999, 2004) panel PP statistic	−4.478***
Pedroni (1999, 2004) group PP statistic	−6.777***
Banerjee and Carrion-i-Silvestre (2011) statistic	−3.065**

Given the relatively small number of time series observations, only one lag was used. Large negative values lead to rejection of the null hypothesis of no cointegration. The ADF statistics and the PP statistics are asymptotically distributed as a standard normal. The 5% critical value for the Banerjee and Carrion-i-Silvestre (2011) statistic (for $T = 15$ and $N = 15$) is −2.54. Banerjee and Carrion-i-Silvestre (2011) do not report 1% critical values. ***, **Indicate a rejection of the null hypothesis of no cointegration at the 1% and 5% level, respectively.

advanced by Pesaran (2006). The pooled CCE estimator accounts for unobserved common factors by augmenting the cointegrating regression with the cross-sectional averages of the dependent and independent variables; these averages are interacted with country dummies to allow for country-specific parameters. In the second step, the estimated parameters are used to construct the residuals from the long-run relationship, $\hat{\mu} = CHURCH_{it} - \hat{\beta}\log(y_{it})$. Then, these long-run residual series are regressed on country dummies D_i to compute OLS residuals from this regression as $\hat{e}_{it} = \hat{\mu} - \hat{D}_i$. Finally, the Pesaran (2007) unit root test is computed for the estimated OLS residuals. If the presence of a unit root in $\hat{e}_{it}$ can be rejected, it can be concluded that there is a cointegrating relationship between the variables.

The results of these tests are presented in Table 3. The ADF and the PP statistics reject the null hypothesis of no cointegration at the 1% level, and the Banerjee and Carrion-i-Silvestre test rejects the null of no cointegration at least at the 5% level (1% critical values are not available in Banerjee and Carrion-i-Silvestre (2011)). Overall, these results strongly support the existence of a linear long-run relationship between religiosity and income over the period 1930–1990. As discussed above, the finding of cointegration also implies that there are no missing (trending) variables and that therefore no additional variables are required in the model given by Equation (1).

III. Empirical analysis

Panel cointegration estimates

We use the panel DOLS estimator suggested by Mark and Sul (2003) to estimate the long-run relationship between religiosity and income. The DOLS

estimator is super-consistent, asymptotically unbiased and normally distributed, even in the presence of endogenous regressors. Moreover, recent Monte Carlo evidence by Wagner and Hlouskova (2010) suggests that Mark and Sul's panel DOLS estimator outperforms other estimators, particularly when the number of time-series observations is small.[5] The idea behind this estimator is to account for possible serial correlation and endogeneity of the regressors by augmenting the cointegrating regression with lead, lag and current values of the first differences of the I(1) regressors. Accordingly, in our case, the DOLS regression is given by

$$CHURCH_{it} = a_i + \beta \log(y_{it})$$
$$+ \sum_{j=-k}^{k} \theta_{ij} \Delta \log(y_{it-j}) + e_{it}. \quad (3)$$

To ensure that our results are not affected by CD (due to common shocks or spillovers among countries at the same time), we compute the CD test suggested by Pesaran (2004). The CD test statistic is defined as

$$CD = \sqrt{\frac{2T}{N(N-1)}} \left(\sum_{i=1}^{N-1} \sum_{j=i+1}^{N} \hat{\rho}_{ij}, \right) \quad (4)$$

where

$$\hat{\rho}_{ij} = \hat{\rho}_{ji} = \frac{\sum_{t=1}^{T} \hat{\mu}_{it} \hat{\mu}_{jt}}{\left(\sum_{t=1}^{T} \hat{\mu}_{it}^2 \right)^{1/2} \left(\sum_{t=1}^{T} \hat{\mu}_{jt}^2 \right)^{1/2}} \quad (5)$$

is the sample estimate of the pair-wise correlation of the residuals of the estimated models, $\hat{\mu}_{it}$. The CD test statistic is normally distributed under the null hypothesis of no CD. As the CD test in the first row of Table 4 shows, the panel DOLS model does not appear to suffer from cross-sectionally dependent residuals; thus, valid inferences can be drawn from the regression results.

As reported in the first row of Table 4, the DOLS regression provides a highly significant negative relationship between religiosity and income. The point estimate implies (if viewed causally) that, in the long run, an increase in per capita GDP by 1% decreases the weekly church attendance rate by about 0.10 percentage points; thus, a doubling of GDP explains a

Table 4. Estimates of the long-run relationship between religiosity and income.

Estimation method	Coeff. on $\log(y_{it})$	CD stat.	Demeaned	Obs.
(1) Pooled panel DOLS estimator (Mark and Sul 2003)	−10.124*** (−11.93)	0.97	No	160
(2) Pooled panel DOLS estimator (Mark and Sul 2003)	−11.721*** (−4.57)	−1.70*	Yes	160
(3) Group-mean panel DOLS estimator (Pedroni 2001)	−8.605*** (−30.42)	−1.00	No	160
(4) Pooled panel FMOLS estimator (Kao and Chiang 2000)	−9.762*** (−13.78)	2.46**	No	192
(5) Pooled panel FMOLS estimator (Kao and Chiang 2000)	−10.428*** (−5.14)	−2.33**	Yes	192

The dependent variable is $CHURCH_{it}$. The DOLS regressions were estimated with one lead and one lag ($k = 1$), given the relatively small number of time-series observations. t-Statistics are in parenthesis. The CD test statistic is normally distributed under the null hypothesis of no cross-sectional dependence. ***, **, *Indicate significance at the 1%, 5% and 10% level, respectively.
DOLS: Dynamic OLS; FMOLS: fully modified OLS.

10% decline in church attendance. This effect can be considered to be (economically) large. To see this, consider the standardized regression coefficient for $\log(y_{it})$, calculated by multiplying the unstandardized coefficient by the ratio of the standard deviations of the independent and dependent variables. The standard deviation of $Church_{it}$ is 22.33 and the standard deviation of $\log(y_{it})$ is 0.64. This means that in the long run, a one standard deviation increase in $\log(y_{it})$ is associated with a decline in church attendance of about 30% of a standard deviation, that is by about 19 percentage points. This result is somewhat larger but in the same ballpark as the earlier estimate of McCleary and Barro (2006a, 160) who found in cross panel regression that a one standard deviation increase in log GDP per capita (i.e. by 0.87) reduces monthly church attendance by 15 percentage points.

Multiplying the estimated coefficient (−0.10) with the average 5-year growth rate of GDP per capita in our sample (11.1%) provides a value of −1.13. On average, the church attendance rate declined by 1.26 percentage points every 5 years, implying that the increase in per capita income accounts for 89% of the decrease in church attendance in the period from 1930 to 1990. Of course, this quantitative estimate must be interpreted with caution given the relatively small number of observations. However, what can be

[5]An extension of the panel DOLS estimator is the dynamic seemingly unrelated regression (DSUR) estimator proposed by Mark, Ogaki and Sul (2005), which was not included in the simulation study by Wagner and Hlouskova (2010). Unfortunately, the DSUR estimator is seriously biased or not applicable when, as in our case, the number of time periods is small relative to the number of cross-sectional units (Mark, Ogaki, and Sul 2005; Di Iorio and Fachin 2012).

safely concluded is that the estimated effect is not only statistically but also economically significant.

During the observation period, log GDP per capita increased in Denmark by 1.6 units from 7.9 (the log of 2666) to 9.7 (the log of 16,866) and the estimated long-run relationship predicts a decline of church attendance by $1.6 \times 10.0 = 16$ percentage points while actually attendance declined by 13 percentage points form 19% to 6%. In Spain, a country of comparable initial GDP, log GDP per capita rose from 7.9 (the log of 2620) to 9.4 (the log of 12,055) such that the long-run relationship motivates a decline of church attendance by 15 percentage points. Actually, however, attendance declined by 31 percentage points from 68% to 37%. These calculations indicate that the long-run relationship obtained from the whole sample underestimates the initial decline of religiosity for countries of initially high attendance rates. This seems to be a natural outcome since it is easier to reduce high attendance rates than to reduce attendance rates that are already low initially. Below we try to accommodate this phenomenon by splitting the sample with respect to initial attendance rates.[6]

In the second row of Table 4, we also present DOLS results based on demeaned data. The coefficient on the income variable remains negative and significant and gets somewhat larger (in absolute value) than its counterpart in row 1. However, the CD test indicates the presence of cross sectional error dependence (at the 10% level), which could have biased the results. The finding of CD for the demeaned data is consistent with studies showing that the demeaning procedure may introduce cross-sectional correlation among the error terms when it is not already present (see e.g. Caporale and Cerrato 2006). In the following, we therefore use demeaned variables only when the CD test indicates the presence of CD.

Robustness

We perform several sensitivity exercises. First, we examine whether the negative relationship between religiosity and income is robust to alternative estimation techniques. A potential problem with the pooled results (reported in row 1) could be that they are based on the implicit assumption of homogeneity of the long-run parameters. It is well known that, while efficiency gains from the pooling of observations over the cross-sectional units can be achieved when the individual slope coefficients are the same, pooled estimators may yield inconsistent and potentially misleading estimates of the sample mean of the individual coefficients when the true slope coefficients are heterogeneous. Although a comparative study by Baltagi and Griffin (1997) concludes that the efficiency gains from pooling more than offset the biases due to individual country heterogeneity, we nonetheless allow the long-run coefficients to vary across countries by using the group-mean panel DOLS estimator suggested by Pedroni (2001). This estimator involves estimating separate DOLS regressions for each country and averaging the long-run coefficients, $\hat{\beta} = N^{-1} \sum_{i=1}^{N}$. The corresponding t-statistic is computed as the sum of the individual t-statistics (calculated using heteroscedasticity and autocorrelation consistent SEs) divided by the root of the number of cross-sectional units, $t_\beta = \sum_{i=1}^{N} t_{\hat{\beta}_i} / \sqrt{N}$. In addition, we use the panel FMOLS estimator suggested by Kao and Chiang (2000). Like the time-series FMOLS estimator, the panel FMOLS estimator incorporates a semi-parametric correction to the OLS estimator, which eliminates the second-order bias induced by the endogeneity of the regressors. We report the results of these estimation methods in rows 3–5 of Table 4.

The results show a negative and significant relationship between religiosity and income. The group-mean DOLS estimator produces a somewhat smaller coefficient than the panel DOLS estimator. However, given the relatively small number of time-series observations, and since especially in this case, the efficiency gains from pooling are likely to more than offset the potential biases due to individual heterogeneity, we prefer the pooled approach. The pooled DOLS estimator is also preferred over the pooled FMOLS estimator, since the results of the pooled FMOLS procedure (reported in row 4 and 5) appear to be affected by CD, even when demeaned data are used, as the CD test statistics show.

---

[6]An interesting case is Ireland where attendance rates remained very high during the whole observation period (98% in 1930 and 92% in 1990). In the years following our observation period, however, Ireland experienced a spectacular economic take off. Log GDP increased from 9.4 (log of 11,818) to 10.2 (log of 26,643). As an out of sample prediction, the estimated long-run relationship suggests that this should lead to a decline of church attendance by 7.2 percentage points. Actually, according to Hirschle (2010), attendance declined by about 22 percentage points.

Given the relatively small number of countries in our sample, we also need to ensure that the estimated effect is not due to individual outliers. To this end, we re-estimate the DOLS regression, excluding one country at a time from the sample. The sequentially estimated coefficients and their t-statistics are presented in Figure 3. Each number on the horizontal axes represents the country omitted from DOLS regression; on the vertical axes, we plot the respective coefficients and t-statistics in the remaining sample. As can be seen, the estimated coefficients are relatively stable and always significant at the 1% level, suggesting that our results are not due to potential outliers.

Next, we re-estimate the DOLS regression for four subsamples: countries with incomes above the sample average, countries with incomes below the sample average, countries with church attendance rates above the sample average and countries with church attendance rates below the sample average.

The resulting coefficients are listed in Table 5. Regardless of which subsample is used, the null hypothesis of no CD cannot be rejected (as the CD statistics show), and the long-run relationship between religiosity and income is negative and significant at the 1% level. From this, it can be concluded that the negative coefficient on $\log(y_{it})$ is driven neither by high-income or low-income countries nor by highly religious or highly secular countries. Clearly, it would be desirable to also assess whether there are systematic differences in the effects of income on religiosity between low-income and high-income countries or between countries with high levels of church attendance and those

Table 5. DOLS estimates for subsamples.

Estimation method	Coeff. on $\log(y_{it})$	CD stat.	No. of countries	Obs.
(1) Countries with incomes above the sample average	−9.108*** (−7.90)	0.19	10	100
(2) Countries with incomes below the sample average	−11.226*** (−10.79)	0.77	6	60
(3) Countries with church attendance rates above the sample average	−10.588*** (−8.66)	1.49	7	70
(4) Countries with church attendance rates below the sample average	−9.704*** (−8.93)	−1.19	9	90

The dependent variable is $CHURCH_{it}$. The DOLS regressions were estimated with one lead and one lag ($k = 1$), given the relatively small number of time-series observations. t-Statistics are in parenthesis. The CD test statistic is normally distributed under the null hypothesis of no cross-sectional dependence. ***, **, *Indicate significance at the 1%, 5% and 10% level, respectively.
DOLS: Dynamic OLS.

with low levels of church attendance. However, the small and varying subsample sizes make it difficult or impossible to obtain comparable estimates of these effects. Our overall judgment is that although the estimated effects vary across the subsamples, the differences are not large enough to be considered conclusive. Thus, we cannot conclude, based on our results, that the effects differ systematically between high- and low-income countries or between high- and low-attendance countries.

We also estimate the long-run relationship between $CHURCH_{it}$ and $\log(y_{it})$ using different samples and time periods. Specifically, the DOLS regression is estimated on a sample of 11 countries over the period 1925–1990 ($N = 11$, $T = 14$) and on a sample of 19 countries over the period 1935–1990
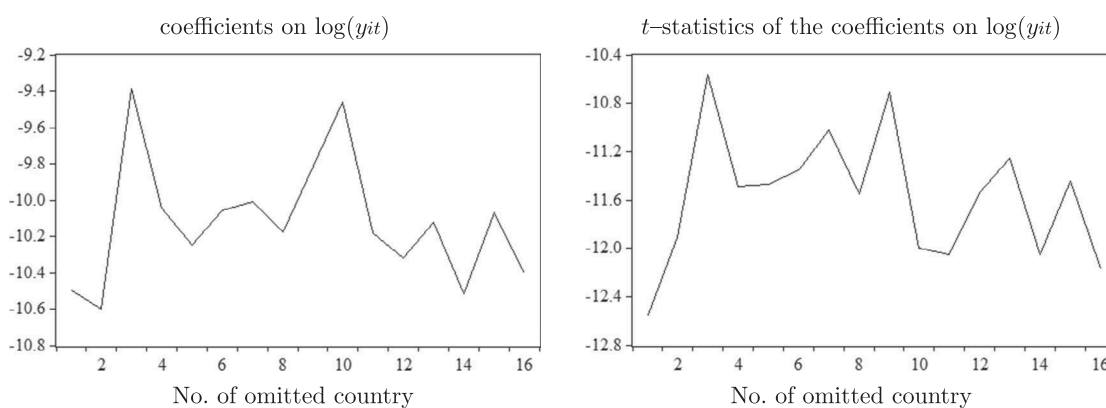

Figure 3. DOLS estimation with single country excluded from the sample coefficients on $\log(y_{it})$ t-statistics of the coefficients on $\log(y_{it})$.

$(N = 19, T = 12)$.[7] As can be seen from Table 6, the null hypothesis of no CD cannot be rejected, and the estimated $\log(y_{it})$ coefficients remain negative and highly significant.

Finally, we examine whether the results are robust to alternative measures of religiosity and alternative specifications of the income variable. Franck and Iannaccone (2014) measure religiosity not only by the parental rate but also by the church attendance rate of children. In the study by Lipford and Tollison (2003), income enters in non-logarithmic form. Table 7 presents the results of the DOLS regressions using these two different variables, labelled $CHURCHchildren_{it}$ and y_{it}, both separately and jointly. As can be seen, all coefficients are negative and statistically significant. Summarizing, the negative effect of income on religiosity is robust to different estimation techniques, potential outliers, different samples, different measures of church

Table 6. DOLS estimates for different samples.

	Coeff. on $\log(y_{it})$	CD stat.	No. countries	Obs.
(1) Smaller sample over a longer time period (1925–1990)	−9.146*** (−8.83)	−0.01	11	121
(2) Larger sample over a shorter time period (1935–1990)	−8.831*** (−8.45)	−0.43	19	171

The dependent variable is $CHURCH_{it}$. The DOLS regressions were estimated with one lead and one lag ($k = 1$), given the relatively small number of time-series observations. t-Statistics are in parenthesis. The CD test statistic is normally distributed under the null hypothesis of no cross-sectional dependence. ***Indicates significance at the 1% level.
DOLS: Dynamic OLS.

Table 7. DOLS estimates using different measures of religiosity and alternative specifications of the income variable.

	Coefficient of the income variable	CD stat.	Obs.
(1) Regressand: $CHURCH\ children_{it}$ Regressor: $\log(y_{it})$	−15.621*** (−11.29)	0.34	160
(2) Regressand: $CHURCH_{it}$ Regressor: y_{it}	−0.001*** (−7.44)	0.89	160
(3) Regressand: $CHURCH\ children_{it}$ Regressor: y_{it}	−0.002*** (−11.30)	−1.27	160

The dependent variable is $CHURCH_{it}$. The DOLS regressions were estimated with one lead and one lag ($k = 1$), given the relatively small number of time-series observations. t-Statistics are in parenthesis. The CD test statistic is normally distributed under the null hypothesis of no cross-sectional dependence. *** indicates significance at the 1%, levels, respectively.
DOLS: Dynamic OLS.

attendance and alternative specifications of the income variable.

Long-run causality

The above interpretation of the estimation results is based on the assumption that long-run causality runs from income to religiosity. However, while cointegration implies causality in at least one direction, it says nothing about the direction of the causal relationship between the variables, as discussed above. Causality may run in either direction, from income to religiosity or from religiosity to income, or in both directions. To test the direction of long-run causality, we follow common practice in the applied panel cointegration literature and employ a two-step procedure. In the first step, we use the (DOLS) estimate of the long-run relationship (from the first row of Table 4) to construct the disequilibrium term

$$ec_{it} = CHURCH_{it} - [\hat{a}_i - 10.124 \log(y_{it})]. \quad (6)$$

In the second step, we estimate the error correction model

$$\Delta CHURCH_{it} = c_{1i} + a_1 ec_{it-1}$$
$$+ \sum_{j=1}^{k} \varphi_{11j} \Delta CHURCH_{it-j}$$
$$+ \sum_{j=1}^{k} \varphi_{12j} \Delta \log(y_{it-j})$$
$$+ e_{it}^{CHURCH} \quad (7a)$$

$$\Delta \log(y_{it}) = c_{2i} + a_2 ec_{it-1}$$
$$+ \sum_{j=1}^{k} \varphi_{11j} \Delta CHURCH_{it-j}$$
$$+ \sum_{j=1}^{k} \varphi_{12j} \Delta \log(y_{it-j}) + e_{it}^{\log(y)} \quad (7b)$$

In this type of dynamic model, the dynamics of both short-run (changes) and long-run (levels) adjustment processes are included; the lagged changes capture the short-run dynamics, while the error correction term captures the long-run relationship. More specifically, the error correction term represents the error in, or deviation from, the long-

---

[7]The countries in the 11-country sample are Australia, Austria, Denmark, Germany, Ireland, the Netherlands, New Zealand, Norway, Spain, the UK and the United States. The countries in the 19-country sample are Australia, Austria, Bulgaria, Canada, Chile, Denmark, France, Germany, Ireland, Italy, the Netherlands, New Zealand, Norway, Portugal, Spain, Sweden, Switzerland, the UK and the United States.

run relationship, and the adjustment coefficients a_1 and a_2 capture how $CHURCH_{it}$ and $\log(y_{it})$ respond to deviations from the long-run relationship. From the Granger representation theorem (Engle and Granger 1987), it follows that at least one of the adjustment coefficients must be nonzero if a long-run relationship between the variables is to hold. A statistically significant error correction term also implies long-run Granger causality from the explanatory variables to the dependent variables (Granger 1988), and thus that the dependent variables are endogenous in the long run. An insignificant error correction term implies long-run Granger non-causality, and thus that explanatory variables are weakly exogenous (Hall and Milne 1994). Given that all variables in the model, including ec_{it-1}, are stationary (because the level variables are cointegrated), a conventional likelihood ratio chi-square test can be used to test the null hypothesis of weak exogeneity, $H_0 : a_{1,2} = 0$.

Table 8 reports the results. As can be seen from row 1, the null hypothesis of weak exogeneity is rejected for $CHURCH_{it}$ at the 1% level, and the CD test suggests that this inference is not affected by CD. Since the null hypothesis of no CD is decisively rejected in row 2 for the residuals from the $\Delta \log(y_{it})$ equation, row 3 uses the demeaned data to account for the CD through common time effects. Row 3 indicates that we can also reject the null hypothesis of weak exogeneity of $\log(y_{it})$. From this, it can be concluded that the statistical long-run causality is bidirectional, implying that declining religiosity is both a consequence and cause of economic growth.

To check the robustness of this conclusion, we perform a standard panel Granger causality test based on a (dynamic) VAR model in levels with fixed effects. While it is well known from the time series literature that in general, the asymptotic distributions of the Wald (or likelihood ratio) test for Granger causality in levels VARs with integrated variables are nonstandard (see e.g. Toda and Phillips 1993), Lütkepohl and Reimers (1992) show that the conventional Wald test for a bivariate cointegrated VAR model is asymptotically distributed as chi-square and therefore valid as a test for Granger causality.

We report the p-values of the Granger causality chi-square statistics using one and two lags,

Table 8. Tests for long-run causality /weak exogeneity.

Weak exogeneity of	$\chi^2(1)$ (p-values)	CD stat.	Demeaned	Obs.
(1) $CHURCH_{it}$	54.97	0.27	No	176
Dependent variable: $\Delta CHURCH_{it}$	(0.000)			
(2) $\log(y_{it})$	13.94	6.09***	No	176
Dependent variable: $\Delta \log(y_{it})$	(0.000)			
(3) $\log(y_{it})$	8.42	−0.54	Yes	176
Dependent variable: $\Delta \log(y_{it})$	(0.004)			

The number of degrees of freedom v in the standard $\chi^2(v)$ tests correspond to the number of zero restrictions. Given the relatively small number of time-series observations, the lag length was set to $k = 1$. The CD test statistic is normally distributed under the null hypothesis of no cross-sectional dependence. ***Indicate a rejection of the null of no cross-sectional dependence at the 1% level.

respectively, of each variable in Table 9. The number of lags was determined by the Schwarz criterion (with a maximum of two lags), as is common practice in testing for Granger causality in a VAR model. As can be seen from the first row, the null hypothesis of no Granger causality from $\log(y_{it})$ to $CHURCH_{it}$ is decisively rejected, and the sum of the coefficients on lagged income is negative in the church attendance equation. This confirms our result that increasing income leads to declining religiosity.

Rows 2 and 3 of Table 9 show that the Granger causality test rejects the null hypothesis that one lag of $CHURCH_{it}$ does not help predict $\log(y_{it})$ at the 1% level and that the coefficient on lagged church attendance is also negative. While the CD test rejects the null hypothesis of no CD in row 2, the results in row 3, using the demeaned data, do not appear to suffer from CD. These results support the conclusion that increasing religiosity leads to declining income.

Finally, Figure 4 presents generalized impulse response functions based on a one lag panel vector error correction model using the raw data over a 50-year horizon (10 5-year periods).[8] As can be seen in the left panel of Figure 4, a one standard deviation shock in income results in a gradual and permanent decrease in church attendance and reaches its full impact after 20 years (four periods). The right panel shows that income gradually and permanently decreases in response to a one standard deviation innovation in church attendance and that the full impact is reached after 25 years (five periods). The impulse response functions are thus consistent with the Granger causality tests reported above.

---

[8]The demeaned data produce qualitatively similar results.

Table 9. Panel VAR causality tests.

Null hypothesis	Lags	p-Value	Sum of coeff. of causal variable	CD stat.	Demeaned	Obs.
(1) $\log(y_{it})$ does not cause $CHURCH_{it}$ Dependent variable: $CHURCH_{it}$	2	0.000	−3.397	0.02	No	176
(2) $CHURCH_{it}$ does not cause $\log(y_{it})$ Dependent variable: $\log(y_{it})$	1	0.000	−0.015	7.85***	No	192
(3) $CHURCH_{it}$ does not cause $\log(y_{it})$ Dependent variable: $\log(y_{it})$	1	0.000	−0.009	−1.44	Yes	192

This table reports the p-values of Granger-causality VAR tests. The null hypothesis is that one or two lags of the (demeaned) series of $\log(y_{it})$ ($CHURCH_{it}$) do not help predict the series of $CHURCH_{it}$ ($\log(y_{it})$). The number of lags was determined by the Schwarz criterion with a maximum of two lags. The CD test statistic is normally distributed under the null hypothesis of no cross-sectional dependence. ***Indicate a rejection of the null of no cross-sectional dependence at the 1% level.
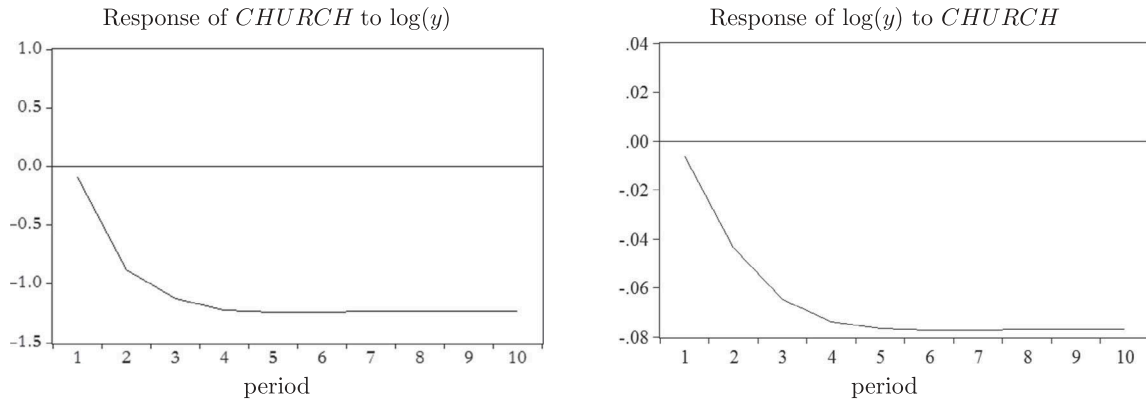
Figure 4. Impulse-responses.

IV. Conclusion

In this article, we took into account that secularization is understood as a dynamic process and investigated the time-series properties of church attendance rates and income per capita. We showed that there exists a cointegrating relationship between both variables and found that according to our preferred specification, a doubling of income is associated with a 10 percentage point decline in church attendance. The presence of cointegration between two or more variables not only implies the existence of a non-spurious long-run relationship between these variables but also that the resulting estimates are robust to a variety of estimation problems that often plague empirical work, including endogeneity and omitted variables. The latter implies that our results are not biased by missing measures of religious supply or other drivers of the demand for religion. This is an important feature with respect to the demand–supply debate in the theory of religion.

We have documented that the negative relationship between income and religiosity is robust to alternative estimation methods, potential outliers, different samples, different measures of church attendance and alternative specifications of the income variable. We also found that long-run causality runs in both directions, higher income leads to declining religiosity and declining religiosity leads to higher income. Secularization appears to be both cause and consequence of economic development.

Our results are consistent with those of Lipford and Tollison (2003), McCleary and Barro (2006a, 2006b) and Paldam and Gundlach (2013), who found that economic development tends to reduce religiosity, but differ from those of Becker and Woessmann (2013) and Franck and Iannaccone (2014), who found insignificant effects of per capita income on religiosity. These studies differ in terms of how they measure religiosity and income, country coverage, whether they use cross-sectional or panel data, the time period covered and/or the empirical methods used. It is therefore difficult to determine why some studies found that economic development has a negative effect on religiosity while others found no effect. One potential reason why our results differ from those of Franck and Iannaccone (2014), who used the same data on church attendance, is that we accounted for the non-stationarity of the data by using panel cointegration procedures. More specifically, Franck and

Iannaccone (2014) employed standard regression techniques (such as OLS), which do not account for the non-stationarity of the underlying data and may be seriously biased (even when cointegration is present; Kao and Chiang 2000). Another reason for the different results could be that our per capita GDP data are in constant 1990 international dollars (Geary–Khamis dollars), while Franck and Iannaccone (2014) used GDP per capita in national currency at current prices.

A limitation of our analysis is the relatively small sample of countries, comprising only relatively rich Western countries that are all historically Christian. Whether the results can be transferred to developing countries or non-Christian countries is an intriguing question, which we are looking forward to addressing when time-series data on religiosity becomes available for these countries as well.

Eventually, with respect to future developments, the strong linear relationship between church attendance and the log of income has to be nonlinear if income continues to grow because church attendance is bounded by zero from below. Our analysis, comprising most of the 20th century, however, has shown that a fading trend of secularization is not yet visible in the data.

Acknowledgements

Disclosure statement

References

Azzi, C., and R. Ehrenberg. 1975. "Household Allocation of Time and Church Attendance." *Journal of Political Economy* 83: 27–56. doi:10.1086/260305.

Baltagi, B. H., and J. M. Griffin. 1997. "Pooled Estimators Vs. Their Heterogeneous Counterparts in the Context of Dynamic Demand for Gasoline." *Journal of Econometrics* 77 (2): 303–327. doi:10.1016/S0304-4076(96)01802-7.

Banerjee, A., and J. J. Carrion-i-Silvestre. 2011. "Testing for Panel Cointegration Using Common Correlated Effects Estimators." *Department of Economics Discussion Paper No. 11-16*. Birmingham: University of Birmingham.

Becker, S. O., and L. Woessmann. 2009. "Was Weber Wrong? A Human Capital Theory of Protestant Economic History." *Quarterly Journal of Economics* 124: 531–596. doi:10.1162/qjec.2009.124.2.531.

Becker, S. O., and L. Woessmann. 2013. "Not the Opium of the People: Income and Secularization in a Panel of Prussian Counties, *American Economic Review*." *Papers and Proceedings* 103: 539–544.

Bonham, C. S., and R. H. Cohen. 2001. "To Aggregate, Pool, or Neither. Testing the Rational-Expectations Hypothesis Using Survey Data." *Journal of Business & Economic Statistics* 19 (3): 278–291. doi:10.1198/073500101681019936.

Breitung, J. 2000. "The Local Power of Some Unit Root Tests for Panel Data." *Advances in Econometrics* 15: 161–177.

Breitung, J., and S. Das. 2005. "Panel Unit Root Tests under Cross-Sectional Dependence." *Statistica Neerlandica* 59 (4): 414–433. doi:10.1111/stan.2005.59.issue-4.

Brierley, P. 1999. *Religious Trends 1999/2000 No. 2*. London: Christian Research Association.

Bruce, S. 2011. *Secularization*. Oxford: Oxford University Press.

Caporale, G. M., and M. Cerrato. 2006. "Panel Data Tests of PPP: A Critical Overview." *Applied Financial Economics* 16 (1–2): 73–91. doi:10.1080/09603100500389143.

Chambers, M. J. 2001. "Cointegration and Sampling Frequency." *Department of Economics Discussion Paper No. 531*. Essex: University of Essex.

Click, R. W., and M. G. Plummer. 2005. Stock market integration in ASEAN after the Asian financial crisis. *Journal of Asian Economics* 16(1): 5–28.

Davidson, J. 1998. "Structural Relations, Cointegration and Identification: Some Simple Results and Their Application." *Journal of Econometrics* 87 (1): 87–113. doi:10.1016/S0304-4076(98)00007-4.

Di Iorio, F., and S. Fachin. 2012. "A Note on the Estimation of Long-Run Relationships in Panel Equations with Cross-Section Linkages." *Economics: The Open-Access, Open-Assessment E-Journal* 6 (20): 1–18.

Dobbelaere, K. 1987. "Some Trends in the European Sociology of Religion: The Secularization Debate." *Sociological Analysis* 48: 107–137. doi:10.2307/3711197.

Eberhardt, M., and F. Teal. 2013. "Structural Change and Cross-Country Growth Empirics." *The World Bank Economic Review* 27 (2): 229–271. doi:10.1093/wber/lhs020.

Engle, R. F., and C. W. J. Granger. 1987. "Cointegration and Error-Correction: Representation, Estimation, and Testing." *Econometrica* 55 (2): 251–276. doi:10.2307/1913236.

Entorf, H. 1997. "Random Walks with Drifts: Nonsense Regression and Spurious Fixed-Effect Estimation." *Journal of Econometrics* 80 (2): 287–296. doi:10.1016/S0304-4076(97)00041-9.

Everaert, G. 2011. "Estimation and Inference in Time Series with Omitted I(1) Variables." *Journal of Time Series Econometrics* 2 (2): 1–26. doi:10.2202/1941-1928.1054.

Franck, R., and L. R. Iannaccone. 2014. "Religious Decline in the 20th Century West: Testing Alternative Explanations." *Public Choice* 159 (3–4): 385–414. doi:10.1007/s11127-013-0103-9.

Granger, C. W. J. 1988. "Some Recent Development in a Concept of Causality." *Journal of Econometrics* 39 (1–2): 199–211. doi:10.1016/0304-4076(88)90045-0.

Granger, C. W. J., and P. Newbold. 1974. "Spurious Regressions in Econometrics." *Journal of Econometrics* 2 (2): 111–120. doi:10.1016/0304-4076(74)90034-7.

Guiso, L., P. Sapienza, and L. Zingales. 2003. "People's Opium? Religion and Economic Attitudes." *Journal of Monetary Economics* 50: 225–282. doi:10.1016/S0304-3932(02)00202-7.

Hakkio, C. S., and M. Rush. 1991. "Cointegration: How Short Is the Long Run?" *Journal of International Money and Finance* 10 (4): 571–581. doi:10.1016/0261-5606(91)90008-8.

Hall, S. G., and A. Milne. 1994. "The Relevance of P-Star Analysis to UK Monetary Policy." *The Economic Journal* 104 (424): 597–604. doi:10.2307/2234634.

Herzer, D. 2008. "The Long-Run Relationship between Outward FDI and Domestic Output: Evidence from Panel Data." *Economics Letters* 100: 147–149. doi:10.1016/j.econlet.2007.12.004.

Herzer, D. 2013. "Cross-Country Heterogeneity and the Trade-Income Relationship." *World Development* 44 (4): 194–211. doi:10.1016/j.worlddev.2012.09.014.

Herzer, D., and H. Strulik 2013. "Religiosity and Income: A Panel Cointegration and Causality Analysis." *CEGE Discussion Paper No. 168*. Gttingen: Center for European, Governance and Economic Development Research (CEGE).

Hirschle, J. 2010. "From Religious to Consumption-Related Routine Activities? Analyzing Ireland's Economic Boom and the Decline in Church Attendance." *Journal for the Scientific Study of Religion* 49: 673–687. doi:10.1111/jssr.2010.49.issue-4.

Hirschle, J. 2011. "The Affluent Society and Its Religious Consequences: An Empirical Investigation of 20 European Countries." *Socio-Economic Review* 9: 261–285. doi:10.1093/ser/mwq025.

Hlouskova, L., and M. Wagner. 2006. "The Performance of Panel Unit Root and Stationary Tests: Results from a Large Scale Simulation Study." *Econometric Reviews* 25 (1): 85–116. doi:10.1080/07474930500545504.

Iannaccone, L. 2003. *Looking Backward: A Cross-National Study of Religious Trends*. Fairfax: George Mason University.

Iannaccone, L. R. 1998. "Introduction to the Economics of Religion." *Journal of Economic Literature* 36 (3): 1465–1495.

Im, K. S., M. H. Pesaran, and Y. Shin. 2003. "Testing for Unit Roots in Heterogeneous Panels." *Journal of Econometrics* 115 (1): 53–74. doi:10.1016/S0304-4076(03)00092-7.

Juselius, K. 1996. "An Empirical Analysis of the Changing Role of the German Bundesbank after 1983." *Oxford Bulletin of Economics and Statistics* 58 (4): 791–819. doi:10.1111/j.1468-0084.1996.mp58004010.x.

Juselius, K. 2006. *The cointegrated VAR Model: Methodology and Applications*. Oxford: University Press.

Kanas, A. 2005. "Nonlinearity in the Stock Price-Dividend Relation." *Journal of International Money & Finance* 24 (4): 583–606. doi:10.1016/j.jimonfin.2005.03.003.

Kao, C. 1999. "Spurious Regression and Residual-Based Tests for Cointegration in Panel Data." *Journal of Econometrics* 90 (1): 1–44. doi:10.1016/S0304-4076(98)00023-2.

Kao, C., and M. Chiang. 2000. "On the Estimation and Inference of a Cointegrated Regression in Panel Data." *Advances in Econometrics* 15: 179–222.

Lahiri, K., and N. Mamingi. 1995. "Power versus Frequency of Observation Another View." *Economics Letters* 49 (2): 121–124. doi:10.1016/0165-1765(95)00668-6.

Levin, A., C.-F. Lin, and C.-S. J. Chu. 2002. "Unit Root Test in Panel Data: Asymptotic and Finite-Sample Properties." *Journal of Econometrics* 108 (1): 1–24. doi:10.1016/S0304-4076(01)00098-7.

Lipford, J. W., and R. D. Tollison. 2003. "Religious Participation and Income." *Journal of Economic Behavior & Organization* 51 (2): 249–260. doi:10.1016/S0167-2681(02)00096-3.

Lütkepohl, H. 2007. "General-To-Specific or Specific-To-General Modelling? an Opinion on Current Econometric Terminology." *Journal of Econometrics* 136 (1): 319–324. doi:10.1016/j.jeconom.2005.11.014.

Lütkepohl, H., and H.-E. Reimers. 1992. "Granger-Causality in Cointegrated VAR Processes. the Case of the Term Structure." *Economics Letters* 40 (3): 263–268. doi:10.1016/0165-1765(92)90002-G.

Madalla, G. S., and S. Wu. 1999. "A Comparative Study of Unit Root Tests with Panel Data and A New Simple Test." *Oxford Bulletin of Economics and Statistics* 61 (Special Issue): 631–652. doi:10.1111/1468-0084.61.s1.13.

Maddison, A. 2003. *The World Economy: Historical statistics*. Paris: OECD.

Madsen, J. B., S. Saxena, and J. B. Ang. 2010. "The Indian Growth Miracle and Endogenous Growth." *Journal of Development Economics* 93 (1): 37–48. doi:10.1016/j.jdeveco.2009.06.002.

Mark, N. C., M. Ogaki, and D. Sul. 2005. "Dynamic Seemingly Unrelated Cointegrating Regressions." *Review of Economic Studies* 72 (3): 797–820. doi:10.1111/roes.2005.72.issue-3.

Mark, N. C., and D. Sul. 2003. "Cointegration Vector Estimation by Panel DOLS and Long-run Money Demand." *Oxford Bulletin of Economics and Statistics* 65 (5): 655–680. doi:10.1111/obes.2003.65.issue-5.

McCleary, R. M., and R. J. Barro. 2006a. "Religion and Political Economy in an International Panel." *Journal for the Scientific Study of Religion* 45 (2): 149–175. doi:10.1111/jssr.2006.45.issue-2.

McCleary, R. M., and R. J. Barro. 2006b. "Religion and Economy." *Journal of Economic Perspectives* 20 (2): 49–72. doi:10.1257/jep.20.2.49.

Moscone, F., and E. Tosetti. 2010. "Health Expenditure and Income in the United States." *Health Economics* 19 (12): 1385–1403. doi:10.1002/hec.v19:12.

Norris, P., and R. Inglehart. 2004. *Sacred and Secular*. Cambridge: Cambridge University Press.

Otero, J., and J. Smith. 2013. "Response Surface Estimates of the Cross-Sectionally Augmented IPS Tests for Panel Unit Roots." *Computational Economics* 41 (1): 1–9. doi:10.1007/s10614-011-9309-4.

Paldam, M., and E. Gundlach. 2013. "The Religious Transition." *Public Choice* 156: 105–123. doi:10.1007/s11127-012-9934-z.

Pedroni, P. 1999. "Critical Values for Cointegration Tests in Heterogeneous Panels with Multiple Regressors." *Oxford Bulletin of Economics and Statistics* 61 (s1): 653–670. doi:10.1111/obes.1999.61.issue-S1.

Pedroni, P. 2001. "Purchasing Power Parity Tests in Cointegrated Panels." *Review of Economics and Statistics* 83 (4): 727–731. doi:10.1162/003465301753237803.

Pedroni, P. 2004. "Panel Cointegration: Asymptotic and Finite Sample Properties of Pooled Time Series Tests with an Application to the PPP Hypothesis." *Econometric Theory* 20 (3): 597–625. doi:10.1017/S0266466604203073.

Pedroni, P. 2007. "Social Capital, Barriers to Production and Capital Shares: Implications for the Importance of Parameter Heterogeneity from a Nonstationary Panel Approach." *Journal of Applied Econometrics* 22 (2): 429–451. doi:10.1002/(ISSN)1099-1255.

Pesaran, M. H. 2004. "General Diagnostic Tests for Cross Section Dependence in Panels." *CESifo Working Paper Series 1229*. Munich: CESifo Group Munich.

Pesaran, M. H. 2006. "Estimation and Inference in Large Heterogeneous Panels with a Multifactor Error Structure." *Econometrica* 74 (4): 967–1012. doi:10.1111/ecta.2006.74.issue-4.

Pesaran, M. H. 2007. "A Simple Panel Unit Root Test in the Presence of Cross-Section Dependence." *Journal of Applied Econometrics* 22 (2): 265–312. doi:10.1002/(ISSN)1099-1255.

Phillips, P., and H. Moon. 2000. "Nonstationary Panel Data Analysis: An Overview of Some Recent Developments." *Econometric Reviews* 19 (3): 263–286. doi:10.1080/07474930008800473.

Rupasingha, A., and B. Chilton. 2009. "Religious Adherence and County Economic Growth in the US." *Journal of Economic Behavior & Organization* 72: 438–450. doi:10.1016/j.jebo.2009.05.020.

Shiller, R. J., and P. Perron. 1985. "Testing the Random Walk Hypothesis: Power versus Frequency of Observation." *Economics Letters* 18 (4): 381–386. doi:10.1016/0165-1765(85)90058-8.

Stark, R. 1999. "Secularization, Rip." *Sociology of Religion* 60 (3): 249–273. doi:10.2307/3711936.

Stock, J. H. 1987. "Asymptotic Properties of Least Squares Estimators of Cointegrating Vectors." *Econometrica* 55 (5): 1035–1056. doi:10.2307/1911260.

Strulik, H. 2016a. "An Economic Theory of Religious Belief." *Journal of Economic Behavior & Organization* 128: 35–46. doi:10.1016/j.jebo.2016.04.007.

Strulik, H. 2016b. "Secularization and Long-Run Economic Growth." *Economic Inquiry* 54: 177–200. doi:10.1111/ecin.2016.54.issue-1.

Toda, H. Y., and P. C. B. Phillips. 1993. "Vector Autoregressions and Causality." *Econometrica* 61 (6): 1367–1393. doi:10.2307/2951647.

Voas, D. 2009. "The Rise and Fall of Fuzzy Fidelity in Europe." *European Sociological Review* 25: 155–168. doi:10.1093/esr/jcn044.

Wagner, M., and J. Hlouskova. 2010. "The Performance of Panel Cointegration Methods: Results from a Large Scale Simulation Study." *Econometric Reviews* 29 (2): 182–223. doi:10.1080/07474930903382182.

Warner, R. 2010. *Secularization and Its Discontents*. New York: Continuum Books.

Wilson, B. 2003. "Secularisation." In *Key Ideas in Sociology*, edited by Slattery, M. Cheltenham: Nelson Thornes.

WIN-Gallup. 2012. *Global Index of Religiosity and Atheism*. http://www.wingia.com/