Churches as Social Insurance:
Oil Risk and Religion in the U.S. South[*]

Andreas Ferrara[†] Patrick A. Testa[‡]

January 5, 2022

Abstract

Religious communities are important providers of social insurance. We show that risk associated with *oil dependence* facilitated the proliferation of religious communities throughout the U.S. South during the 20th century. Known oil abundance predicts higher rates of church membership, which are not driven by selective migration or explained by changes in income, education, or population. Consistent with social insurance, greater oil price volatility increases effects, while greater access to credit, state-level social insurance, and private insurance crowds out effects. Religious communities limit spillovers of oil price shocks across sectors, reducing increases in unemployment following a negative shock by about 30%.

Keywords: Social insurance, religious participation, resource abundance, oil, risk.
JEL Codes: N31, N32, H41, G52

---

[*] We are grateful to Guy Michaels for sharing with us various data from Michaels (2011), as well as Philipp Ager for pointing us to the charitable spending data in the original 1936 U.S. Census of Religious Bodies. Thanks to Sascha O. Becker, Jeanet Bentzen, Dan Bogart, Jean-Paul Carvalho, Lewis Davis, Klaus Desmet, Arcangelo Dimico, Kerice Doten-Snitker, Price Fishback, Tony Gill, Dan Hungerman, Dan Keniston, Igor Kopylov, Sara LaLumia, Ralf Meisenzahl, Jonathan Moreno, Ömer Özak, Nancy Qian, Jon Pritchett, Gary Richardson, Jared Rubin, Michael Shin, Stergios Skaperdas, and Anthony Wray as well as participants in talks at LSU, Tulane, UC Irvine, ASREC 24h, the 90th SEA Meeting, and ASSA 2021 for helpful feedback. A previous version of this paper was circulated in September 2020 as CAGE WP No. 513 under the title, “Resource Blessing? Oil, Risk, and Religious Communities as Social Insurance in the U.S. South.” All errors are our own.

[†] University of Pittsburgh, Department of Economics. Email: a.ferrara@pitt.edu.

[‡] Tulane University, Department of Economics. Email: ptesta@tulane.edu.

1 Introduction

In a world with economic uncertainty, individuals seek to insure themselves against negative shocks and rely substantially on the state for help. For instance, $2.7 trillion, or about two thirds of U.S. federal spending, was devoted to social insurance and other public assistance expenditures in 2016 alone (Pew Research Center, 2017). Prior to the New Deal era, however, provision of social insurance was largely eschewed by the federal government (Fishback, 2020).

Instead, early providers of social insurance were often *religious communities*.[1] In 1926, U.S. churches spent $150 million on supporting members and their communities, compared to $23 and $37 million in social expenditures by state and local governments, respectively (Gruber and Hungerman, 2007). Churches continue to be important in settings lacking strong formal insurance mechanisms, in which they serve as networks for providing charity and aid (McCleary and Barro, 2006). A growing literature exists showing how environments characterized by risk, such as those with greater rainfall variability (Ager and Ciccone, 2018) or earthquake risk (Bentzen, 2019), tend to foster higher levels of religious participation and belief.

This paper documents the development of religious communities in the face of another source of considerable economic volatility: natural resource dependence.[2] In resource-rich settings, managing risk associated with such volatility is key to enjoying their economic benefits. Indeed, volatility is a key factor driving the "resource curse," in which resource dependence is linked to poor economic outcomes (van der Ploeg and Poelhekke, 2009; Venables, 2016).

To study the relationship between natural resource risk and religious participation, we consider the setting of the U.S. South, where oil abundance has contributed to both increased wealth and economic risk.[3] Beginning in the late 19th century, discoveries of large oilfields throughout Texas, Oklahoma, and Louisiana sparked the births of new urban areas, or "boomtowns," built around oil and marked by rapid economic growth (Michaels, 2011). These oil communities simultaneously endured great risk, as a petroleum industry characterized by price volatility came to dominate their economies over the long-run (Brown and Yücel, 2004). To the extent that strong formal insurance mechanisms were often limited in oil boomtowns, such risk might in turn have given rise to increased religious participation throughout the oil-rich South, where Christian church membership remains highly concentrated today.

To test this hypothesis, we consider five questions. First, we examine whether oil endowments predict greater religious participation in the U.S. South over the short- and long-run. Second, we explore whether

---

[1] Religious communities continue to provide such services today, we discussed in Section 2.

[2] The quantities of resources extracted and used tend to be highly persistent in the face of changes in price. Low price elasticities mean that even small shocks can generate large changes in price (Ross, 2012).

[3] This paper adopts a nomenclatural interchangeability of resource "dependence" and "abundance" from that literature, as is appropriate in our setting (Michaels, 2011), although we explore variation in dependence later on.

the relationship between oil abundance and religious participation is driven by selective migration or urban growth following oil discoveries. Third, we evaluate to what extent effects are driven by oil's economic risk, using volatility in world oil prices and variation in non-oil economic activity. Fourth, we examine whether religious communities emerged as a form of social insurance against such risk, using variation in the availability of public and private insurance substitutes. Lastly, we explore the precise channels through which religious communities smooth consumption in the face of oil price shocks. We answer these questions primarily using a difference-in-differences (DD) design, exploiting variation in the locations and timing of major oilfield discoveries from 1890 to 1990, paired with county-level data on church membership.[4] Following Michaels (2011), our dataset covers counties in the three states at the epicenter of the United States' oil boom—Texas, Oklahoma, and Louisiana—as well as nearby counties in surrounding states. Counties are considered to be "oil abundant" if they lie above at least one oilfield with over 100 million barrels of oil pre-extraction. Church membership data are available via the Association of Religion Data Archives (ARDA) for 15 major Christian denominations in the form of 9 religious "censuses," spanning 1890 to 1990. Importantly, such major denominations are likely to span not just oil areas but multiple regions, with idiosyncratic risk.[5]

First, our main results confirm that Southern counties with known oil abundance exhibit higher rates of religious participation over the sample period. DD estimates show a 6.6 to 8.1 percentage point increase in membership as a share of population among these major Christian churches following major oil discoveries, compared to a sample mean of about 33%,[6] with effects persisting over time in event-study specifications. Overall, these effects explain around 30% of the growth in church membership that occurred among these denominations in Southern oil counties from 1890 to 1990.[7] To provide evidence for the exogeneity of oil discoveries and rule out anticipatory effects, we perform several balance tests, showing that oil discoveries cannot be predicted by relevant pre-discovery county characteristics, and fail to estimate statistically significant pre-treatment effects. We also adopt leading tools for DD designs with heterogeneous treatment timing to rule out contamination of DD and event-study estimates stemming from heterogeneous treatment effects (Goodman-Bacon, 2021; Sun and Abraham, 2021). Results are furthermore robust to dropping individual years and states, using different treatment year definitions, and

---

[4] Religious membership is a stricter definition of participation than adherence or belief, as it frequently entails some form of baptism or confirmation. While some religious censuses report adherence, all report membership.

[5] Some denominations, most notably several Black Baptist and Methodist groups, are excluded from our sample as they lack data for much of the oil discovery period. Despite these omissions, based on the 1926 data our measure of Christian membership still accounts for nearly 80% of a measure that adds Black Baptist and Methodist groups.

[6] This rate may seem low for three reasons: (i) membership is a strict measure of religious participation that entails rites such as baptism or confirmation; (ii) membership rates were somewhat low in the early 20th century before rising after WWII, as discussed under Figure 3; and (iii) not all groups are included in our sample, as discussed in the previous footnote.

[7] The unconditional sample mean rose from 21.5% in 1890 to 43.9% in 1990. Hence our lower-bound estimate of 6.6 percentage points accounts for $\frac{6.63}{43.9-21.5} \times 100 = 29.6\%$ of this increase.

correcting for spatial autocorrelation.

Second, we show that this relationship is driven neither by selective migration among more religious types following major oil discoveries, nor by differential changes in income, education, or population size and composition across oil and non-oil counties. First, using the first name patterns of migrants' children in the U.S. Census to infer their religious characteristics, we test whether Biblical names predict migration by household heads into oil counties following major oil discoveries (Anderson and Bentzen, 2021; Raz, 2021). We find no evidence that the relationship between oil abundance and religious participation is driven by selection. Second, we show that treatment effects cannot be explained away by changes in income, education, or population. Third, we show that results are robust to comparing oil and non-oil counties with similar pre-discovery population growth trends, implying that results are not driven by oil counties that were already on a higher population growth trajectory. To do this, we generate a sample of matched oil and non-oil counties based on their pre-discovery population growth trends using propensity score matching. Lastly, we generate another sample matching oil and non-oil counties based on population trends over the entire sample period. Oil counties exhibit higher levels of religious participation even when compared to non-oil counties growing similarly in terms of population and population density.

Third, we evaluate to what extent such effects are driven by oil's economic risk. To measure volatility in oil prices, we use the running standard deviation of real world crude oil prices over time. A one sample standard deviation increase in oil price volatility coincides with treatment effects about 24% larger, even after controlling for the price level of oil. Effects are also smaller among counties with greater population density, manufacturing output, and cotton activity at the beginning of the event window, suggesting that they stem from higher-risk boomtowns that emerged around or are largely dependent on oil.

Fourth, we provide evidence for religious communities as a social insurance mechanism against oil's risk. We begin by examining whether the existence of formal insurance mechanisms with more complete contracts "crowds out" increases in church membership in response to oil discoveries (Hungerman, 2005; Chen, 2010). To do this, we exploit state-level variation in public social insurance, which emerged after the New Deal. The availability of unemployment insurance and workers compensation at poverty-line levels reduces the size of treatment effects by about two thirds. Effects are also more than halved for counties with access to private insurance prior to treatment, and credit availability as measured by county-level access to savings and loans associations and bank tellers reduces effects almost entirely. We also examine whether demand for such substitutes similarly increases with oil abundance. We find savings, time deposits, and the number of insurance agents and brokers per capita to all be increasing following major oil discoveries, even after controlling for urban growth.

Lastly, we examine whether economic fluctuations driven by oil price shocks are smaller in oil-abundant counties when a larger religious community is present. Using a subsample of already-treated and never-treated counties, we compare how local labor composition responds to oil price shocks when religious participation levels are above and below the sample median. We find that increases in relative unemployment following a negative oil price shock are about 30% smaller in oil-abundant counties conditional upon a large religious community being present. Consistent with historical and sociological evidence that churches are key providers of economic aid as well as various social services, we find evidence that religious communities help reduce out-migration following negative shocks, while also limiting the spread of such shocks across sectors, particularly agriculture. Corresponding to this, we find that increases in church membership are larger among conservative Christian denominations, which tend to be more efficient providers of such services to members (Iannaccone, 1992), with oil discoveries also driving increased competition among denominations overall.

This paper makes several contributions to our understanding of natural resources and the roots of contemporary U.S. culture and identity. We highlight a novel and important connection between fundamental endowments of the natural environment and the evolution of local cultural attributes, namely the distinctly Christian religious expression of the U.S. South. We show how places with known oil abundance historically saw heightened levels of religious participation, even when comparing places with otherwise similar trajectories in urbanization and population growth. Through careful analysis, we provide empirical support for a key channel through which this relationship took form–social insurance–exploiting a rich and multifarious set of evidence based on heterogeneity in oil prices over time, differential local public insurance provision across states, spatial and temporal variation in private insurance substitutes, and more. To our knowledge, this is the first paper to connect religious intensity to natural resource dependence and its associated economic risk. Our work nonetheless advances a burgeoning body of research that connects variability in the natural environment to religious participation and belief (Ager, Hansen and Lonstrup, 2016; Ager and Ciccone, 2018; Bentzen, 2019) and other sociocultural attributes (Raz, 2021; Fiszbein, Jung and Vollrath, 2021). In the U.S. setting, these findings add to a literature on the historical roots of American culture and identity (Bazzi, Fiszbein and Gebresilasse, 2020; Fouka, Mazumder and Tabellini, 2021), including the unique evangelicalism that characterizes the U.S. South (Dochuk, 2012, 2019), as well as the origins and persistence of its salient geographic sociocultural divides (Bazzi, Ferrara, Fiszbein, Pearson and Testa, 2021; Desmet and Wacziarg, 2021).

These results also offer new insight into the role of religious communities as social insurance mechanisms in history. This relates to a large literature on informal insurance mechanisms (Townsend, 1995), which include not only churches but agricultural guilds and familial networks (Richardson, 2005; Fafchamps and Lund, 2003). Similar to our paper is Ager et al. (2016), who document increases in religious participation as a form of mutual insurance in the aftermath of the Great Mississippi Flood of 1927. Consistent with our findings, they show that affected counties lacking in formal mechanisms for managing risk turned to religion to smooth consumption, while increased access to credit crowded out such effects. In addition to this, our paper finds that the availability of insurance, both private and public forms, sizably reduces increases in the demand for religion in oil-abundant counties. In this way, we also complement existing work on the "crowding out" effects of both government welfare and lending institutions on religious participation (Hungerman, 2005; Gruber and Hungerman, 2007; Chen, 2010). The importance of this mechanism may help explain declines in religious participation in developed regions of the world, where formal insurance mechanisms tend to be stronger.

Our findings furthermore represent an important contribution to the broader debate in economics on natural resources. Recent work by van der Ploeg and Poelhekke (2009), van der Ploeg (2011), Cavalcanti, Mohaddes and Raissi (2014), Venables (2016) and others have emphasized the role of volatility in driving the "resource curse." Yet less work has been done to study how resource-rich economies actually deal with volatility. We show that religious communities emerged in oil-abundant counties in the U.S. South as a form of social insurance, mitigating the impact of oil price shocks. Moreover, we find that demand for religion is greater in the absence of formal insurance and lending institutions. That churches mitigate the effects of shocks as a substitute for other, more formal mechanisms for smoothing consumption is consistent with existing evidence that resource-rich countries with poor financial institutions tend to be more volatile (van der Ploeg and Poelhekke, 2010), while shedding new light on how countries that *do* lack mature financial and formal insurance mechanisms may deal with resource risk.

Lastly, our findings add to a large literature on urban growth and persistence. Why and to what extent "boomtowns" and other settlements persist over the long-run are questions of active inquiry among urban and development economists (Bleakley and Lin, 2012; Fafchamps, Koelle and Shilpi, 2016). This is especially the case for settlements established around natural factors, where the resource curse threatens to undermine long-run economic outcomes. U.S. oil settlements, on the other hand, have proven quite successful over the long-run, which may offer lessons for elsewhere (Michaels, 2011; Allcott and Keniston, 2018). In particular, our findings suggest that institutions for managing adverse economic shocks may be crucial for transforming mono-industrial boomtowns into high-income agglomerations.

Before turning to background, we offer a few notes on interpretation. While we provide evidence linking a place's reliance on oil and its religious activity through a social insurance channel, we are by no means arguing that this is the *only* channel through which oil discoveries affect religious participation. Rather, this is a key mechanism for which we find strong empirical support, relative to others for which we do not. Nor is oil necessarily the only factor driving the rise of religious participation in the U.S. South during the 20th century. Finally, we are not claiming external validity with respect to the relationship between natural resource dependence and religious intensity. This should be investigated in subsequent research, including across a variety of natural resources and religious denominations.

2 Background

Religion as Social Insurance

Historians and social scientists have long noted the role of churches in providing various forms of social and economic support. In the United States, the 19th century Social Gospel movement coincided with an increase in church activity dedicated to addressing poverty, inequality, and education (Cnaan, Boddie, Handy, Yancey and Schneider, 2002). Prior to the advent of federal government welfare and social insurance programs in the 1930s,[8] religious communities had become a leading provider of financial assistance and other aid, such as food and clothing, as well as employment matching services, job and vocational training, and other social services (Gruber and Hungerman, 2007; Chen, 2010). In 1926, for example, churches provided a total of $150 million in charitable spending according to the U.S. Census of Religious Bodies. In comparison, social and charitable spending by state governments amounted to $23 million and $37 million by local governments in the same year (Gruber and Hungerman, 2007).

Even today, churches play an active role in providing important support to their communities, particularly where formal mechanisms for dealing with economic uncertainty are lacking (Bartkowski and Regis, 2003; Hungerman, 2005; McCleary and Barro, 2006; Scheve and Stasavage, 2006). One U.S. study by Cnaan et al. (2002) finds that about 75% of congregations have some form of financial support mechanism for the poor, while Glaeser and Sacerdote (2008) note that 3 in 5 Catholics and 3 in 4 Baptists in the U.S. General Social Survey trust their congregation to help them in times of hardship. The sources of such support tend to be the voluntary contributions of other members.[9] This is consistent with club models of religion and subsequent empirical work documenting how religious communities provide local public goods to their members, derived from the religious "investments" of members such as tithing and service-based contributions (Iannaccone, 1992, 1998). To the extent that major religious organizations span networks of individuals and communities with idiosyncratic risk and state realizations, churches

---

[8] The 1935 Social Security bill established federal unemployment insurance. Savings and loan banks became widespread after the 1932 Federal Home Loan Bank Act.

[9] For example, Wilson and Janoski (1995) find that 71% of conservative Protestants who attend church weekly engage in church-related volunteer work, relative to 38% that do not actively attend a congregation.

thus entail an important source of mutual insurance and consumption smoothing for religious participants (Berman, 2000; Chen, 2010; Ager and Ciccone, 2018).[10]

Oil Abundance as a Source of Risk

One source of risk against which religious communities might insure is generated by natural resource dependence. Natural resource quantities are generally inelastic in price in the short-run, leading to large fluctuations in world prices (van der Ploeg and Poelhekke, 2010; Ross, 2012). High volatility in the world prices of natural resources can in turn translate into considerable volatility in real income for the economies endowed with them, which may have ramifications for income *growth*, in the form of the so-called "resource curse." For instance, van der Ploeg and Poelhekke (2009) find that natural resource abundance is associated with positive growth after controlling for volatility, although they leave it relatively open as to how one might manage the risk associated with such volatility.[11]

This paper focuses on one of the most price-volatile of all natural resources: oil. According to van der Ploeg and Poelhekke (2009), oil prices are more volatile than the prices of agricultural raw materials, food products, and ores and metals. By one estimate, world oil prices experience more volatility than 95% of all other products sold in the U.S. (Ross, 2012). By another, oil-rich countries experience over 100% greater volatility (e.g. in revenues) than non-resource-rich countries, compared with 50% for mineral-rich countries (Venables, 2016).

Petroleum has been a major industry in the U.S. since the discovery at Oil Creek in Pennsylvania in 1859, but America's "oil age" arguably began in Texas at the turn of the century with the strike at Spindletop. The oil boom that followed gave rise to hundreds of new settlements, or "boomtowns," throughout the Southern United States, especially in the states of Texas, Louisiana, and Oklahoma. The economies that emerged were highly oil-centric and remained heavily dependent on oil through the late 1980s (Brown and Yücel, 2004). Yet despite oil's booms and busts (see Figure 1), they oversaw considerable economic growth throughout the 20th century (Michaels, 2011).

Oil and Religion in the U.S. South

Southern oil communities experienced oil's volatility not only via the market's booms and busts but also in the form of accidents such as fires and explosions. Though religion has featured prominently in the U.S. South since at least the Civil War period,[12] recent work by Dochuk (2012, 2019) documents how these

---

[10] Consequently, lower levels of religious participation tend to be observed when formal mechanisms for smoothing consumption are strong (Gill and Lundsgaarde, 2004; Gruber and Hungerman, 2007; Ager et al., 2016).

[11] Their analysis does provide some insight relevant for this paper, in that resource-rich countries experience greater income volatility when they are more resource *dependent* (i.e. % GDP) and when financial systems are less developed.

[12] Of course, there is also considerable spatial variation in religiosity within the South, with the epicenter of the "Bible Belt" moving westward over time (Brunn, Webster and Archer, 2011).

challenges inspired an evangelical fervor and reliance on religion in the South's oil-patch boomtowns. Weak state governments meant that oil activities went largely unregulated, and the church emerged as a key institution serving oil communities.

Thereafter, petroleum and the Christian culture of the American South and Southwest "collaborated to construct a shared ideology and system of institutions" that persisted through the interwar era (Dochuk, 2012, p. 55). Today, they share a space in American political culture,[13] with oil and evangelicals overlapping on the ideological right in the latter's opposition to climate change legislation and offshore drilling bans (Pew Research Center, 2015).

3 Data and Empirical Strategy

We now describe the county-level dataset compiled for this paper. Our dataset combines information on county-specific oil discoveries with data from population and religious censuses for the U.S. We discuss each in turn before describing how this data are used to estimate the effect of oil abundance on religious participation in the U.S. South.

Our main sample is based on the one in Michaels (2011), which covers all counties with major oil discoveries in Texas, Oklahoma, Louisiana, and surrounding states, including any county within 200 miles of an oil-abundant county in those three states (see Figure 2).[14] This leaves a relatively large number of control counties while limiting the geographic heterogeneity of the sample. Oilfields with 100 million barrels of oil or more are defined as a "major." The main treatment variable is an indicator for when at least one major oilfield is discovered in a county. The first county to discover a major oilfield in this sample did so in 1893, while the most recent discovery was made in 1982. Figure 3 shows the evolution of the outcome and treatment over time.

Data on the locations of major U.S. oilfields come from the Oil and Gas Journal Data Book (2000), which lists all $\geq$ 100 million barrel oilfields in the United States, their locations by state, and their discovery years. We link major oilfields with data for all county-oilfields from the Oil and Gas Field Code Master List (U.S. Department of Energy, 2004), which lists all oil and gas fields in the United States, their counties, and each field's discovery year within each county.

We combine the oil discovery data with information on county-level religious membership, from nine religious "censuses," beginning in 1890 and ending in 1990. These correspond to the red lines in Figure 3. All censuses are collectively available through the Association of Religion Data Archives (ARDA).[15] All

---

[13] Also see Bentzen and Sperling (2020) for the interrelation between politics and religiosity today.

[14] We adopt the same sample as Michaels (2011), although our replication moves 9 counties in Texas from oil-abundant to non-abundant, as only natural gas portions of major oilfields are present in those, leaving 162 oil-abundant and 613 non-oil.

[15] For 1906-36, we opt to use to use the versions from Haines (2010), for reasons of data cleanliness and formatting.

religious censuses measure participation using counts of church "membership," which may entail baptism or confirmation, while only some also report a less restrictive count of "adherents." We therefore adopt the former. As far as possible, we harmonize church membership and denominations across years for major Christian groups. Many smaller denominations, as well as predominately-Black Baptist and Methodist groups, are excluded as their memberships were not reported for much of the oil-discovery period. We then aggregate membership data across groups to construct our measure of religious participation. County boundaries are harmonized to boundaries from the nearest religious census year, following the procedure used in Hornbeck (2010) and described in Ferrara, Testa and Zhou (2021), in order to consistently match them with census data and to avoid issues of the merging or splitting of counties over time. More detailed information on variable construction can be found in the Data Appendix.

Other data come from the U.S. Census of Population and Housing and the Census of Agriculture, compiled by Haines (2010) and Haines, Fishback and Rhode (2018), respectively, as well as manufacturing data from Matheis (2016) and data from the individual full-count Census provided by Ruggles, Flood, Goeken, Grover, Meyer, Pacas and Sobek (2018). These include data on county population, sector composition, wages and output, savings and loan associations, bank deposits, and the number of insurers and tellers. To measure educational attainment, we use the percentage of adults (aged 25 plus) with a high school degree. These are not available for 1926 and prior, for which we use the literacy rate. To measure income, we use median family income in 2018 dollars. As these data are also not available prior for 1926 and prior, we rely on another commonly-used metric for those years based on occupational income scores. This score is computed based on the median income in a given occupation as of 1950, assigned to individuals' occupations in previous years to proxy for their income. We then average those income scores by county based on the full-count census files. Once these census data are combined with religious data, county boundaries are harmonized again to 2000 boundaries to create a unified panel. Detailed information on this data can be found in the Data Appendix. Data on names of migrants comes from the 1940 Census and are matched with the list of Biblical names at behindthename.com to determine their religious content. Lastly, data on public social insurance, including unemployment insurance and workers compensation at the state level for 1940 through 1990, are merged from Fishback (2020).

3.1 Estimation

We estimate the effect of oil abundance on church membership using the following generalized difference-in-differences (DD) framework with two-way fixed effects:

$$\mathrm{y}_{ct} = \mu_c + \theta_t + \beta \cdot \mathrm{D}(t \geq E_c) + \varepsilon_{ct}, \tag{1}$$

where y_{ct} is the share of church members in the population of county c in year t, E_c is a county's first major oil discovery event, μ_c are county fixed effects, and θ_t are year fixed effects. The coefficient of interest is β, which captures average differences in religious participation between oil-abundant and non-oil counties relative to such differences prior to a major oil discovery.

Several conditions must be satisfied for β to be interpreted as the causal effect. First, the locations of major oil discoveries should be exogenous to relevant factors. That is, there must not be time-varying confounding factors that correlate both with the discovery of oil and a county's religious participation. Second, there must be no changes in the outcome across treated oil and non-oil control counties in the absence of treatment (i.e. pre-trends). Third, given the heterogeneity in treatment timing, treatment effects must be constant over time or β will be biased (Goodman-Bacon, 2021). We discuss these concerns now.

Plausible Exogeneity of Oil Discoveries

One identification concern is potential selection into treatment. We thus examine whether observable county characteristics can predict major oil discoveries. Due to heterogeneous treatment timing, control counties lack a natural pre-oil counterfactual period. To deal with this, we define local *clusters* of counties around each treated county that discovers oil in period t.[16] We then regress an indicator for being an oil-abundant county on the observable characteristics of all counties in a given cluster in period $t-1$. Based on the fact that all counties in the vicinity of the treated county should have been similarly likely *ex ante* to also discover the same major oilfield, we can test whether certain characteristics in fact predict such discoveries locally, indicating non-random treatment assignment.[17] Figure A1 plots the coefficients from these balance tests, using a variety of specifications to isolate potential local differences. We find no statistically significant differences in observable pre-oil characteristics between non-oil and eventual oil-abundant counties, corroborating the identifying assumption that major oilfield discoveries are as good as randomly assigned in space.

This finding is consistent with the historical narrative. Early methods of oil discovery were unreliable as they relied on first discovering surface saps or random drilling. For example, the strike at Spindletop that began the Texas oil boom came when Patillo Higgins, a Sunday school teacher in Beaumont, noticed clouds of gaseous liquid in spring water on a class field trip (Dochuk, 2012). Much of the speculation that followed Spindletop subsequently also took place in Southeast Texas. Yet the most significant Texan oil discoveries would not be made until 1930, 170 miles north of Spindletop, when Columbus Joiner

---

[16] These local clusters are built by first computing the distance of each control county to all oil counties. All control counties that have the same oil county as their closest treated county are then grouped together.

[17] The pre-oil discovery characteristics we consider are the shares of Christians, Blacks, French, Italian, and German, log population size, employment shares in manufacturing and agriculture, the number of manufacturing firms, farms, and area dedicated to agriculture, the share of agricultural land used for cotton, total land area, and latitude and longitude. Except for latitude and longitude, all variables are standardized to have mean zero and variance one for the purpose of comparability.

discovered the 134,000 acre, five-county East Texas Oilfield northeast of Henderson.[18] Because none had yet sought oil in that region, tracts of land were sold at little cost to incoming prospectors (Olien and Olien, 2000). Technologies to detect oil continued to be rather unreliable in later years. In 1950, only 20% of performed onshore oil explorations were successful (Gorelick, 2009). Even in 1978, the exploratory success rate was only at 27.5% (Forbes and Zampelli, 2002).[19] This suggests that wealthier counties, for instance, were not more likely to discover oil, corroborating the balance tests. Note that this does not rule out post-treatment migratory selection, which we discuss in Section 4.1 below.

Estimating Dynamic Treatment Effects

To estimate changes in treatment effects over time, including possible pre-trends, we utilize an event-study framework in addition to the standard DD. This event-study framework entails the following:

$$\mathrm{y}_{ct} = \mu_c + \theta_t + \sum_{l=\underline{l}}^{-2} \gamma_l \cdot \mathrm{D}(t - E_c = l)_{ct} + \sum_{l=0}^{\bar{l}} \gamma_l \cdot \mathrm{D}(t - E_c = l)_{ct} + \varepsilon_{ct}, \tag{2}$$

Under this approach, treatment effects are expressed over an *effect window* $l \in [\underline{l}, \bar{l}]$, which we set to be $[-3, +3]$, and are estimated relative to the omitted period before the observed event (i.e., $l = -1$). For $l < -1$, γ_l estimates pre-trends and for $l \geq 0$, γ_l estimates the dynamic treatment effects of the event. The effect window length was chosen to create a mostly balanced panel. Counties with observations outside of the effect window are binned into -4 and $+4$ groups. These dummies are used in estimation only to identify dynamic treatment effects, with such observations serving as controls (Schmidheiny and Siegloch, 2019). As such, the estimates γ_{-4} and γ_4 are themselves not of interest and not reported.

4 Main Results

To study the connection between oil abundance and religious participation in the U.S. South, we exploit variation in the timing and locations of major oil discoveries in Texas, Oklahoma, Louisiana, and surrounding states between 1890 and 1990. Following a major oil discovery, a county is considered to be treated and of known "oil abundance."

We begin by estimating the difference-in-differences (DD) framework in equation (1). In our preferred specification, we exclude 255 counties directly adjacent to oil-abundant counties to avoid biases arising from spillover effects.[20] This minimizes the probability that the control group is itself partially

---

[18] Given this, we consider specifications that cluster standard errors beyond the county level or allow for spatial correlation to account for the tendency of oil endowments and other relevant characteristics to span beyond the county level.

[19] In a later sensitivity exercise we show that no particular year is driving our estimates, suggesting that changes in oil discovery technologies are unlikely to be influencing our results (see Figure A2).

[20] Table A1 in the Appendix shows that excluding adjacent control counties is sufficient to remove most spillover effects, as compared to excluding counties within 50, 100, and 150 km of the nearest oil county, while *not* excluding counties in the

treated, which would bias estimated effects towards zero, while nonetheless retaining the comparability of treated and control counties on observable characteristics due to geographic proximity.

Estimates using this "donut" subsample can be found in column (1) of Table 1, which shows a 6.6 percentage points (pp) increase in the rate of church membership associated with oil abundance among the 15 major Christian churches sampled—about 20% above the mean. Note that the sample mean indicates average *membership* among the denominations sampled, which typically requires baptism or confirmation and thus often excludes children and recent converts.[21] Secondary specifications corroborate these findings. Estimates from the full sample reduce estimates, as expected, to about 4.8 pp.

Table 1 also reports the δ statistic by Oster (2019) to test for the sensitivity of our results with respect to selection on unobservables. The test uses the coefficient and R^2 movements from the controlled and uncontrolled estimates to develop a bounding exercise. The δ statistic then reports how much more influence the unobservables, relative to the observables, should need to have in the relationship between oil and religion in order to explain away our oil abundance effect if those observables could be included in the regression. The value of δ in column (1) implies that the unobservables would need to be 1.7 times more important than the observable controls in order to potentially explain away our treatment effect. A value of $\delta \geq 1$ is typically considered a robust result.

Columns (3-6) consider how treatment effects change with the inclusion of various covariates. These include controls for population size, manufacturing employment, and cotton production. Estimates are not particularly sensitive to any particular set of covariates. That being said, we do caution that these controls could potentially be outcomes of the oil treatment themselves, and our preferred specification is therefore the one reported in column (1).

Event-study Estimates, Dynamic Effects, and Pre-trends

A large literature on difference-in-differences (DD) estimation illustrates several challenges underlying both the validity of such designs and the unbiasedness of their estimates (Goodman-Bacon, 2021; Sun and Abraham, 2021). In this section, we address two main concerns: (i) pre-treatment effects, in which changes in the outcome potentially occur across treated oil and non-oil counties prior to the former's treatment switching on; and (ii) heterogeneous treatment effects in the presence of heterogeneous treatment timing, which can bias DD estimates due to the use of different timing groups (i.e. early- versus late-treated) as controls alongside never-treated counties. To do this, we use the event-study framework defined in equation (2), which interacts the treatment indicator with event time dummies for census periods

---

proximity reduces our treatment effect as predicted.

[21] As an exception, Roman Catholic participation is only available after 1970 in terms of adherence. That being said, this is the case across all counties, and our results are robust to dropping these years.

leading up to and following the treatment year. This approach allows us to estimate possible pre-trends as well as evaluate changes in the treatment effect over time. Because different timing groups can also contaminate estimates in *event studies* under heterogeneous treatment timing, we also use the alternative event-study estimator proposed by Sun and Abraham (2021).

We begin by plotting our baseline event-study estimates with 3 relatively balanced pre-periods in panel (a) of Figure 4, which corresponds to the DD estimate in column (1) of Table 1. Two things stand out. First, post-treatment effects appear somewhat small at first, likely due in part to heterogeneity in the amount of time since the treatment year as of period 0, and then continue increasing over time thereafter. In other words, our setting indeed has heterogeneous treatment effects in the presence of heterogeneous treatment timing. As we later show, the increasing nature of our effects in our setting likely stems from higher levels of oil price volatility in the later years of the sample. Second, there is a slight downward-sloping trend, albeit a statistically insignificant one, in our pre-treatment estimates.

We first turn our attention to the increasing nature of our treatment effects over time. Following Goodman-Bacon (2021), this pattern raises the possibility that DD estimates are biased toward zero. This is more likely to be the case if estimation relies greatly on time variation, i.e. comparisons between different timing groups. Fortunately, our sample contains a large number of never-treated counties, with the Bacon decomposition showing nearly 90% of variation used in estimating equation (1) as coming from comparisons between treated and never-treated units, and DD estimates using *only* this variation to be 7.75.

The relative unimportance of time-variation in our setting is corroborated by robustness exercises showing DD estimates to be largely unchanged when alternative treatment years are used, changing treatment timing. Besides our preferred treatment definition, in which the treatment "switches on" the year a county discovers its first major oilfield, we construct four alternative measures. The first defines as the treatment year the year *any oil at all* is discovered in an oil-abundant county. The second considers an oil-abundant county to be treated the year it *or an adjacent oil-abundant county* discovers a major oilfield, whichever is first. The third considers an oil-abundant county to be treated the first year any of its major oilfields are discovered *in any county*, even if not in that county. The fourth considers an oil-abundant county to be treated the first year any county in the set of oil-abundant counties with which it is *contiguous* discovers a major oilfield. Despite these increasingly conservative definitions, Table A4 shows estimates to be highly stable, varying by no more than half of a percentage point.

Our estimation therefore benefits from the existence of a large number of never-treated counties. Earlier, we established that these serve as useful controls for to-be-treated counties in our sample, to the extent that they are locally comparable along relevant dimensions prior to the latter's treatment. Yet for equation (1) to identify causal effects, it must also be the case that we observe no pre-trends, i.e. effects prior to the treatment year. While we do not estimate statistically significant pre-trends, the slight downward-sloping trend in panel (a) of Figure 4 may indicate the presence of subtle anticipatory affects, with oil discoveries inspiring searches for others nearby and some changes in the outcome in turn occurring prior to any discovery having been made. If true, this would invalidate our research design.

We perform three supplementary event-study exercises to show that our results are unlikely to be driven by pre-trends and that the downward-sloping trend in panel (a) of Figure 4 is likely the result of contamination from the use of different timing groups as controls. The first two exercises entail basic extensions of equation (2). Although one common strategy for dealing with possible pre-treatment "effects" is to control for county-specific linear time trends, this is likely to dilute estimates of dynamic treatment effects in the presence of increasing effects over time. Goodman-Bacon (2021) proposes an alternative strategy, in which pre-treatment trends are estimated directly and then subtracted from the outcome for all periods. Adjusting for pre-trends accordingly in our main specification decreases estimates somewhat to 4.9 (.9) pp, which remains highly significant. The associated event-study plot can be seen in panel (b) of Figure 4.

A second approach involves modifying the estimation strategy to better compare treated counties with control counties that are likely to experience similar pre-treatment activity. This involves matching on pre-treatment trends in log population or log population density, via propensity score matching, and the use of matched-pair fixed effects. These exercises tend to flatten pre-trends whilst generating moderate decreases in DD estimates relative to our main specification. One such example can seen in panel (c) of Figure 4, which matches on pre-treatment growth trends in log population density. This particular example yields a smaller DD estimate, of 4.0 (1.3) pp. All DD estimates from pre-treatment matching exercises can be found in Table A2. This offers further reassurance that possible pre-trends are not a threat to identification.

The third exercise employs an alternative event-study estimator that is robust to heterogeneous treatment timing. Sun and Abraham (2021) show that when heterogeneous treatment effects exist in settings with heterogeneous treatment timing, pre-trend tests based on pre-treatment estimates are invalid. As in Goodman-Bacon (2021), this is driven by contamination from the use of different timing groups as controls alongside never-treated units. They propose an alternative estimator that relies on never-treated counties as controls while excluding different timing groups. Estimates from this approach can be found in panel (d) of Figure 4. Post-treatment estimates are highly similar to those from OLS, while pre-treatment estimates are smaller and statistically insignificant, with minimal downward trend.

Robustness: Spatial Correlation and Spillovers

A secondary identification issue concerns *spatial interdependencies* across counties. We consider such factors now. One such issue involves spatial correlation in the disturbance ε_{ct}, which is likely to span beyond individual counties c. We adopt several approaches to account for these. First, we estimate Conley standard errors, which assume that unobservables may be correlated across contemporaneous counties up to distances of 25, 50, 100 and 150 km. In general, this generates smaller standard errors than specifications that cluster standard errors, as shown in Table A5. Second, we estimate effects using standard errors clustered at trans-county levels of aggregation, at which the treatment is likely to have occurred in practice. We adopt two approaches: (i) clustering at the oilfield level, where a county is assigned to whichever of its major oilfields was discovered first in any county, and (ii) clustering by the set of contiguous oil-abundant counties of which a county is a part, even if they share no common oilfields, as treatment in practice may span even beyond the major oilfield level. These standard errors are slightly larger but estimates remain highly significant, as shown in Table A6.

Another issue is that effects from oil discoveries are not necessarily confined to oil-abundant counties but may have positive spillovers for neighboring counties, biasing estimates. Indeed, Table 1 considers both "donut" and full sample specifications, where including counties adjacent to oil-abundant counties decreases treatment effects by nearly a fifth. To further examine the importance of spillover effects, we consider several additional donut exercises, excluding counties within 50, 100, and 150 km of a treated county. Increasingly extreme measures to limit spillovers increase estimates, as expected, with the most restrictive approach producing a DD estimate of 7.5 (1.0), as shown in Table A1.

Robustness: Sample Issues

This subsection considers two robustness exercises concerning the nature of the sample. The first examines potential heterogeneity of effects across space and sample year. For instance, effects may be driven entirely by Texas, which makes up nearly a third of the sample and contains more than half of the treated counties, or by certain census periods, given the vast event window. To probe for the sensitivity of our results to particular states or years, we re-estimate equation (1) and exclude each state and year at a time and observe how the oil abundance effect changes from our baseline result. Figure A2 plots the coefficients from this exercise. The omitted group is reported on the left scale and the coefficients represent the estimated effect of oil abundance on the share of Christian membership under a given sample restriction. The baseline result is represented by the vertical red line. In none of the cases does excluding any given state or year significantly change the estimated coefficient relative to the baseline result.

The second exercise deals with the construction of the outcome variable. Because our measure of religious participation requires the same set of Christian denominations across all years, certain groups are inevitably excluded, such as various Pentecostal churches, Black Methodist churches such as the African Methodist Episcopal Church, and Black Baptist churches, namely the National Baptist Convention. This latter group in particular did not participate in the last four religious censuses in our sample.

Unfortunately, it is not enough to simply exclude National Baptists, as the 1906 Religious Census reports National Baptists *in combination* with the non-Black Northern and/or Southern Baptists for each county. To remedy this, our preferred outcome variable sums Northern, Southern, and National Baptists in 1916 and uses the county-level ratio of non-Black to Black Baptists in that year to impute non-Black Baptist counts for 1906. As the actual ratio may differ over time, however, we also consider several alternative approaches. The first simply drops 1906; the second uses the same ratio but from 1890, which lacks some counties for Oklahoma; and the third excludes Baptists entirely. Effect sizes are remarkably stable and statistically significant across all specifications, as shown in Table A7.

4.1 Are Effects Driven by Selective Migration?

We have established a robust positive relationship in the data between oil abundance and religious participation. Prior to interpreting this as a causal effect, however, we must address the relative importance of *selective migration* in driving this relationship. An ideal experiment would take two similar populations of fixed size and composition, randomly endow one with oil and related industry, and measure how patterns of religious participation change in the former relative to the latter. In reality, many of those living in oil-abundant counties are likely to have migrated there *following* oil's discovery and subsequent increases in labor demand. This raises the possibility that observed increases in religious participation are not a response to living in oil-abundant areas but rather the outcome of compositional change among the population following the discovery of oil, driven by migratory sorting among more religious types.

To test this concern, we examine whether an individual's religious characteristics predict their migration to oil-abundant counties following the discovery of oil, in both the short- and long-run. To measure the religious characteristics of migrants, we follow Anderson and Bentzen (2021) and use the given names of their born children. Names are considered "Biblical names" if they match or are phonetically related to names featured in the Old or New Testament, as listed on behindthename.com. Table 2 shows the most common Biblical and non-Biblical names in our dataset. Our sample of migrants to Southern oil and non-oil counties is derived from the "residence in 1935" question in the 1940 U.S. Census, which associates census respondents with their migration status between 1935 and 1940, when the Southern oil boom was at its peak (see Figure 3).[22] We therefore consider the names of unmarried children of household heads, between the ages of 5 and 18 as of 1940, as they would have been named prior to moving and thus would have had names indicative of the pre-migration religious characteristics of migrating household heads, similar to Raz (2021).

We first regress an indicator for whether a child's household head moved between 1935 and 1940 to a county in which a major oil discovery was made during that period on an indicator for whether that child had a Biblical name. We exclude from the set of counties those in which oil was discovered prior to 1935, such that migrants to oil-abundant counties are compared only to individuals who never selected into an oil-abundant county. To minimize spillover effects (e.g. from commuters), we also exclude from the set of non-oil destination counties those that were nonetheless adjacent to known oil-abundant counties as of 1940. Since we are comparing migrants across oil- and non-oil counties, we include state fixed effects. We also control for dummies for a child's age, race, sex, and place of birth. Standard errors are clustered at the family level. In secondary specifications, we also include a dummy for whether the household head worked in the oil sector.

In nearly all cases, we find evidence consistent with *negative* selection, in which household heads with Biblically-named children are *less* likely to move to newly oil-abundant counties. This is consistent with the notion that correlates of urban growth are often negatively correlated with religious participation, as discussed in the next subsection. In contrast, household heads working in oil appear to select positively into newly oil-abundant counties. These estimates are shown in columns (1-4) of Table 3.

Even if we observe no evidence of migratory sorting by religious types in our baseline regression, there may still be selective migration occurring among later-movers. For instance, this could account for our increasing treatment effects in event studies. Hence, we also examine whether religious characteristics predict migration between 1935 and 1940 to counties in which major oilfields were discovered *prior* to 1935. To do this, we regress an indicator for whether a child's household head moved between 1935 and 1940 to a county in which major oil was discovered prior to 1935 on an indicator for whether that child had a Biblical name. We exclude migrants to counties that were treated between 1935 and 1940, as well as migrants to non-oil counties that were adjacent to known oil-abundant counties as of 1940 to reduce spillover effects. These estimates, shown in columns (5-6) of Table 3, are similar in sign and size to those associated with movers to newly oil-abundant counties, suggesting that selective migration by religious types is not driving our treatment effects over the short- or long-run.

We consider three extensions of this exercise. The first excludes children of non-migrants from the sample, such that the religious characteristics of movers are being compared only to those of movers to

---

[22] This avoids sample limitations and matching problems associated with using the linked census over our long event window.

(non-adjacent) never or non-yet-treated counties. In contrast, the baseline exercise also included in the comparison group (i) non-migrants living in the same never or non-yet-treated county in both 1935 and 1940, as well as (ii) non-migrants already living in an not-yet-treated oil-abundant county as of 1935 that became treated between 1935 and 1940. For the "long-run" exercise, this means that some individuals who migrated to treated counties prior to 1935, i.e. closer to their treatment year, were included in the comparison group. In general, estimates are not sensitive to who we include. These results can be found in Table A8. Second, we consider a version that limits the set of names to those coded as either Biblical or non-Biblical under *both* the NYSIIS and Soundex phonetic coding algorithms. Results are relatively unchanged in this smaller sample. These results can be found in Table A9. Third, we consider a sample that sets the maximum age of children to 10, following Bazzi et al. (2020). These estimates, shown in Table A10, are unchanged as well.

4.2 Are Effects Explained by Changes in Education, Income, or Population?

Our results thus far suggest that local religious communities indeed grew in the response to major oil discoveries. However, we cannot yet speak to whether oil abundance *per se* directly drove such growth. For instance, economic development associated with oil discoveries may be tied to a number of important changes locally, such as changes in population size and density, income growth, and increased education levels. These are potentially relevant factors, of which changes may explain some or even all of oil's affect on religious participation. For instance, a substantial body of literature has studied how religious participation changes alongside education and income.[23]

To address this, we first examine whether changes in education, income, population size, or population composition capture the relevant variation in our treatment variable, by including them as controls in our main regression. Table 4 confirms many of these variables to be relevant. We find a negative and significant relationship between educational attainment and religious participation, alongside a small albeit significant, negative relationship for income, similar to Becker and Woessmann (2013) in the Prussian context. At the same time, neither variable explains away or meaningfully alters the treatment effects. Columns (3-5) consider alongside these the inclusion of various population size and compositional controls, again with little consequence to our treatment effects. These are instructive, as they suggest our effects are not driven by agglomeration or compositional effects from oil discoveries. Column (3) for instance takes into account the extensive margin of population change. Especially in the earlier sample periods, increases in population size were mainly driven by the inflow of new people rather than new

---

[23] The relationship between education and religiosity has been studied prominently by Marx (1844). For a more modern treatise, see for instance Glaeser and Sacerdote (2008) and Becker, Nagler and Woessmann (2017). The relationship between income and religion has been studied by Bettendorf and Dijkgraaf (2010), Becker and Woessmann (2013), Herzer and Strulik (2017), and others. We also refer to reader to the overviews by Iannaccone (1998) and McCleary and Barro (2006).

births. Column (4) takes into account the intensive margin. Reaffirming the selective migration exercise from the previous section, changes in compositional factors result in little change in estimates.

We further probe for the robustness of our results with respect to issues of population growth and urbanization, using a series of matching exercises in order to better compare treated counties with control counties that have similar urbanization trajectories, to which our results are highly robust. This exercise involves matching and comparing treated and control countries based on either similar (i) pre-treatment or (ii) full sample trends in log population or log population density, using propensity score matching and estimation with matched-pair fixed effects. While estimates vary somewhat in terms of size, these matching exercises universally preserve the sign and statistical significance in our estimates, again suggesting that oil abundance is affecting religious participation through channels independent of initial or post-oil urban growth. These can be found in Tables A2 and A3.

5 Religious Communities as Social Insurance

Having established in the previous section that major oil discoveries predict large and persistent increases in the rate of church membership in the oil-rich U.S. South over the 20th century, which are explained by neither selective migration nor changes in income, education, and population size, we now turn to our favored interpretation. In this section, we examine to what extent religious communities emerged specifically as a form of *social insurance*. We begin by first examining simply whether treatment effects are rooted in the *risk* associated with oil abundance. We then examine whether demand for religious communities stems from their provision of actual economic support.

5.1 Oil Risk Drives Religious Participation

Downturns in the oil market may generate substantial economic distress for oil communities.[24] Uncertainty in the presence of such shocks may trigger increased religious participation as a form of *social insurance*, to the extent that the major denominations we consider span multiple regions with idiosyncratic risk.[25] To test this, we first establish the importance of oil *risk* in driving our treatment effects, using volatility in oil prices and variation in non-oil economic activity.

Oil Price Volatility and Treatment Effects

To capture volatility in world oil prices at a given point in time, we use running standard deviations of real crude oil prices (in 2018 USD) over preceding 5, 10, and 25 year periods. A larger standard deviation implies greater price volatility. Because crude oil prices are determined by the world market,

---

[24] Brown and Yücel (2004) show this in the context of 20th century Texas, while our own county panel data show effects of oil price shocks on unemployment, income, and various sectoral outcomes in oil counties, as shown in Table A11.

[25] See the Theory Appendix for a model illustrating this.

their fluctuations are largely exogenous to the conditions in any individual oil county. We then interact this measure of volatility with the treatment indicator in equation (1) to capture the interaction effects of oil volatility on religious participation in oil-abundant counties relative to non-oil counties.

Table 5 shows these interaction effects using the preferred "donut" specification with two-way fixed effects and no other controls. Estimates reveal that a one standard deviation increase in oil price volatility coincides with treatment effects that are about 24% larger, consistent with our hypothesis, with larger and more precise interaction effects coming from longer-run measures of volatility.

Figure 1 illustrates the temporal sources of variation in oil price volatility in our sample, which is moderate in the early decades of the sample, low in the middle years, and high from 1970. To the extent that oil price volatility is an important driver of the treatment effect, this explains why dynamic treatment effects are increasing over time in the event-study, as later event years are more likely to correspond to the 1971, 1980 or 1990 religious censuses.

One concern is that oil prices tend to be higher on average at times of high volatility. Although we previously uncovered little association between income and religious participation, it remains possible that such short-term positive shocks to oil price *levels* may be driving this effect, rather than the *volatility* in prices per se. To test this, we include in Table A12 lagged demeaned real price level interactions alongside the volatility interactions, which largely sustains our estimates of the latter's effects.

Oil Dependence and Treatment Effects

Crude oil prices are not the only source of risk faced by oil communities. When the Texas oil boom began at the turn of the 20th century, much of Texas was sparsely populated and neither Oklahoma nor New Mexico were U.S. states. And although major agglomerations often sprung up around oil boomtowns (Michaels, 2011), many oil communities in the sample remained highly oil-dependent throughout the 20th century—and arguably would have remained sparsely populated in the absence of oil's riches. To the extent that oil is the primary source of wealth and urban growth in many parts of the U.S. South and Southwest, the local economic risk entailed by oil abundance may be amplified even more in those places. This in turn may have contributed to the treatment effects observed above, driving even greater reliance on religious communities for economic support during oil downturns.

To capture heterogeneous effects from greater dependence on oil, we develop several time-invariant variables with which to interact the treatment dummy. These variables capture non-oil economic activity at the start of the event window, when to-be-treated and never-treated counties were similar on average. The first two capture the *extent of urbanization* at the start of the Oil Age. One of these is a continuous variable, measuring log population density in 1900. Another is a dummy capturing whether a county contained part of an urban agglomeration with a population of at least 25,000 in 1890. To the extent that a county had high levels of urbanization prior to discovering oil, its economic performance would likely be less dependent on oil's bounty.

The second set of variables captures a county's *manufacturing activity* prior to oil's discovery. The first measures log manufacturing output per capita in 1900. The second is a dummy indicating above-sample-median manufacturing output in 1900. Having a larger manufacturing presence implies lower aggregate risk: not only does it mean a place has historically had economic activity independent of oil, but manufacturing serves as a key substitute sector for mining labor during oil downturns.[26]

The third set of variables captures a county's *cotton activity* prior to oil's discovery, using a continuous measure of the share of land used for cotton in 1900 and a corresponding above-sample-median cotton dummy. Cotton in itself was a source of economic risk throughout the South in the early 20th century, but this risk was driven by different sources than oil's, such as the boll weevil. Hence, cotton-rich places that discovered major oil might have smaller treatment effects, due to risk diversification.

Consistent with oil dependence driving economic risk, treatment effects are significantly larger in more oil-dependent counties, with smaller pre-discovery populations, manufacturing sectors, and cotton density. Estimates produced from these interactions, as shown in Table A13, reveal that treatment effects are driven entirely by counties without an urban presence prior to oil's discovery—with a one standard deviation increase from the sample mean in log population density in 1900 reducing treatment effects to zero—and reduced significantly by the pre-oil presence of manufacturing and cotton. This suggests effects are being driven largely by higher-risk "boomtowns" whose economic activity stems from oil.

5.2 Crowd Out

Although religious participation appears to increase in response to oil's economic risk, it is possible such effects are rooted in psychological forms of support rather than economic ones. Indeed, evidence for "religious coping" has been documented in recent work by Bentzen (2019). Though we do not rule out religious coping as a possible additional channel, we now provide evidence that religious communities indeed provide economically-meaningful forms of social insurance demanded in response to oil risk.

To show this, we first examine key substitutes in the provision of insurance. If religious participation emerges in oil counties in response to increased demand for social insurance, one would expect a greater supply of private and public substitutes for smoothing consumption, with more complete contracts, to *crowd out* relative increases in religious participation in oil counties (Hungerman, 2005; Chen, 2010).

---

[26] Whereas a decrease in oil prices tends to decrease mining wages, mining output, and even agricultural employment in oil versus non-oil counties, the same decrease tends to *increase* manufacturing employment and slightly manufacturing wages, consistent with a shift to the right in the manufacturing labor supply curve. For these estimates, see Appendix Table A11.

We begin by examining public insurance substitutes, in the form of state-provided social insurance.

Oil and Religion in the Public Social Insurance Era

Public social insurance first became widespread in the U.S. in the 1930s but continued to vary by state thereafter, with Southern states generally providing less. Hence, to study interaction effects from public social insurance, we interact our county-level oil abundance indicator with time- and state-varying measures of maximum weekly (i) unemployment insurance and (ii) workers compensation benefits from Fishback (2020). These measures are each adjusted to reflect relative value, using the ratio of benefits to either (i) the national poverty line weekly income equivalent for a 4 person family, (ii) average weekly earnings plus benefits for manufacturing workers in the state, or (iii) state weekly per capita income.

As shown in Table 6, we find that the introduction of public social insurance benefits comparable to national poverty line income standards reduces effects by nearly two thirds. When it is available at levels comparable to even higher standards, such as average state incomes or manufacturing earnings, positive effects of oil on religious participation disappear entirely, although availability to this extent is seldom the case in our sample. This mirrors previous findings by Gruber and Hungerman (2007) showing a reduction in charitable giving by churches following the New Deal.

Access to Credit and Private Insurance

Do oil counties that had greater access to credit and private insurance prior to oil's discovery exhibit similar "crowd out" effects to that observed with public social insurance? To answer this, we use data on historical credit availability, as measured by the number of savings and loan associations and bank tellers in each county, as well as on private insurance, as measured by the number of insurance agents and brokers in each county. As a large number of counties report having zero of these, we generate a dummy variable for each, indicating whether a county had any prior to treatment.

Due to limitations in how far back credit availability data go, we drop some treated counties in our preferred specifications to reduce "bad control" concerns. In particular, savings and loan associations became prominent throughout the U.S. following the Federal Home Loan Bank Act of 1932, and county-level reporting on them first corresponds with our religious census dataset in 1950. County-level data on bank tellers and insurance agents, meanwhile, are available for the full set of sample counties starting in 1910. We thus feature specifications that drop counties treated prior to 1950 when examining heterogeneous effects from savings and loan banks and pre-1910 for bank tellers and private insurance.

Estimates for these are featured in Table 7 and are consistent with the notion that greater pre-treatment access to credit and private insurance crowds out increases in religious participation from oil dependence. In particular, having a savings and loan association or access to bank tellers reduces effects almost entirely in specifications that drop previously treated observations, as shown in columns (1b) and (2b), respectively. Similarly, effects are more than halved for counties with access to private insurance, as shown in column (3b). Effects are robust to instead including all observations, as shown in the remaining columns. Overall, this suggests religious communities are indeed engaged in the provision of social insurance, primarily in the absence of private substitutes.

Given the attenuation of effects from the supply of private substitutes, one might also expect *demand* for them to increase with oil abundance, all else fixed. To measure demand for private banking as used for smoothing consumption, we use two continuous measures of savings as outcomes: log savings capital in savings and loan associations per capita and log time deposits per capita. To measure demand for private insurance, we use the logged number of insurance agents and brokers per capita.

These results are again consistent with the notion of religious communities and private banking or insurance as substitutes. Table A15 shows oil abundance increases demand for savings capital and time deposits by nearly 10%, while it increases the number of insurance agents and brokers per capita by nearly 50%. Note that, as many counties lack banks, savings and loan associations, or insurance agencies, results for these are somewhat sensitive to including zero-valued counties, with statistical significance for log savings capital appearing only in their absence.

One concern is that the results in Tables 7 and A15 do not reflect demand for these private substitutes but rather are merely corollaries of urban growth. For instance, a place with no insurance agents is likely to be rural and therefore more religious, while oil abundance is likely to increase population size and in turn the probability that an insurance branch locates in that county. We replicate Table 7 but control for pre-treatment log population density in 1950 and 1910 to probe for robustness, for columns (1a-b) and (2a-3b) respectively. The results are reported in the Appendix in Table A14. Though interaction effects are reduced, neither main nor interaction effects lose statistical significance. Hence, while initial population size does seem to matter for the relationship between private substitutes for consumption smoothing and religious participation, it does not explain away our main finding.

Appendix Table A16 performs a similar exercise for the regressions in Table A15, controlling for time-varying log population density to test whether post-oil urban growth captures the relevant variation driving increases in outcomes in Table A15. The only effect that loses statistical significance is for the log number of insurance agents per capita when zero-valued counties are included. In general, increases in log population density and urbanization during the sample period do not explain away our findings.

5.3 Do Religious Communities Smooth Economic Volatility?

Now that we have shown that religious communities exhibit substitutability with other insurance mechanisms in oil-abundant settings, we consider whether relative fluctuations in economic outcomes following oil prices shocks are indeed smaller in oil counties when a large religious community is present. To do this, we compare how relative labor outcomes in oil counties respond to oil price shocks when religious participation is above versus below the sample median, using our measure of Christian membership in 1936 and a subsample of already- and never-treated counties spanning 1940 to 1990.[27]

Consistent with the historical evidence in Section 2, in which churches have historically been involved in providing economic aid and services such as employment matching and job training to their members, relative increases in unemployment in oil counties after a negative oil price shock are significantly smaller when a large religious community is present. Table 8 shows that, whereas a negative oil price shock of $30 2018 USD—about a sample standard deviation—increases the unemployment rate in an oil-abundant county relative to non-oil county by about 1 pp among below-median Christian counties, the same shock increases the relative unemployment by only about 0.7 pp among above-median Christian counties, with the difference between these coefficients being highly statistically significant.

To what extent is this rooted in reduced consumption volatility? The rest of Table 8 suggests that religious communities indeed achieve this by limiting the spread of demand shocks beyond the oil sector. Following a negative oil price shock, oil counties with relatively small religious communities tend to experience out-migration and the local agricultural sector suffers, while manufacturing becomes relatively more important for local labor demand. In oil counties with relatively *large* religious communities, however, migratory patterns remain flat relative to non-oil counties, and relative fluctuations in agricultural and manufacturing employment are both reduced by about two thirds. This is consistent with the notion that material aid and provision of labor market matching services by religious communities may smooth the local effects of income shocks in oil counties—not just for oil workers but for their neighbors.

5.4 Evidence from Denominational Responses

Further evidence that effects are driven by increased demand for social insurance can be found by looking at different denominational responses to oil discoveries, such as (i) changes in competition among denominations and (ii) relative changes in participation among denominations that are more versus less likely to efficiently provide social insurance. First, column (1) of Table A18 shows that oil abundance decreases the concentration of different Christian denominations in our sample of denominations by about

---

[27] This year is chosen because most outcomes are reported in the U.S. Census County Data Books beginning in 1940. We therefore exclude counties treated after 1936, which as such may see increases in Christian participation later in the panel thus biasing estimates. Results are largely robust to not dropping these counties, as shown in Table A17.

3.8%. This makes sense: a positive demand shock for the goods and services provided by religious communities, such as social insurance, may trigger new entrants into the market for religion and spur all denominations, including pre-existing ones, to increase their level of effort, increasing religious participation across multiple denominations (Iyer, 2016).

Which denominations are driving this? We examine the three most prominent groups in our sample of the U.S. South: Roman Catholics, Southern Baptists, and United Methodists. We find sizable and significant membership increases following major oil discoveries for only the first two, as shown in Table A18. This is consistent with club models of religion, as described in Section 2, in which joining a religious community entails access to certain goods and services (e.g. social insurance) in return for the costs entailed by membership. More liberal denominations with fewer such membership costs (e.g. less sacrifice), particularly mainline Protestant denominations such as Methodists, in turn tend to entail greater "free-riding" with respect to such goods, at the expense of their provision to members (Iannaccone, 1992). Provision of social support and volunteering among mainline Protestants also tends to be more community- than congregation-focused, relative to evangelical groups (Wilson and Janoski, 1995). A demand shock for social insurance would therefore not necessarily generate an increase in membership among Methodists, relative to Southern Baptists. Catholics, meanwhile, are ideologically mixed, having become more liberal overall since the Vatican II reforms while remaining conservative in the Gulf Coast South, i.e. in the heavily Catholic parts of our sample (Iannaccone, 1994).

There are nonetheless key differences between Catholics and Southern Baptists. Catholics tended to focus their aid on urban areas, in which 4% of total church expenditures in 1936 went toward local relief and charity, versus about 2.1% in rural areas (U.S. Census Bureau, 1936). Meanwhile, Southern Baptists invested more in rural areas, with 3.2% of expenditures going to local relief and charity in such places, versus 1.9% in urban areas. These comparative advantages reflect their membership: only 19.4% of Catholics in 1936 were rural versus 62.1% of Southern Baptists. The heatmaps in Figure 5 show the spatial differentiation of Catholics and Southern Baptists early in the event window.

After WWII, Catholics continued cluster along the Gulf Coast and near Mexico, while Southern Baptists became increasingly dominant throughout the remaining Southern counties in our sample, as shown in subsequent heatmaps in Figures A5 and A6. Indeed, most of the increase in overall membership observed in our data stems from the latter's growth.[28] Of this growth, a disproportionate amount occurred in oil counties: Southern Baptist membership grew 437% in oil-abundant counties between 1906 and 1990, compared to 244% in (non-adjacent) non-oil counties.[29] The heatmaps also show the relative

---

[28] Southern Baptists grew from 4.8% of sample population in 1906 to 20% in 1990; Catholics grew from 6.8% to 13.5%; United Methodists grew from 5.7% to 7.4%.

[29] For comparison, sample Catholic participation grew 143% in oil counties versus 60% in non-oil counties over the same

stagnation of the Methodists, whose historically conservative Southern groups joined the more liberal Methodist Episcopals in 1939.

6 Conclusion

Social insurance and public assistance make up much of federal spending in the U.S. today. Yet before the advent of such programs, churches were a key provider of social and economic support (Gruber and Hungerman, 2007). This continues to be the case in settings lacking strong formal insurance mechanisms, where churches may serve as networks for providing charity and aid (McCleary and Barro, 2006).

In this paper, we provide evidence that local religious communities grew throughout the U.S. South in part as social insurance mechanisms in the face of oil and its associated volatility. Using county-level data from 1890 to 1990 on major oil discoveries and church participation in Texas, Oklahoma, Louisiana, and neighboring states, we document a strong and persistent relationship between oil discoveries and religious participation. Difference-in-differences regressions show that Christian church membership increases by 6-8 percentage points in oil abundant counties in the aftermath of their first major oil discovery. We find no evidence that effects are driven by selective migration or urban growth in oil-rich places. These effects continue to grow over time, with the volatile later years of our sample feeding into heightened demand for church support. Part of this persistence post-1940 is rooted in a lack of "crowd out": Southern states often lack strong state social insurance programs, even today, and we find that such states continue to exhibit stronger effects. Another part likely boils down to path dependence, with Christian faith becoming an increasingly core part of Southern identity in the U.S. Indeed, our findings explain 30% of the overall increase in Christian church memberships from 1890 to 1990 in our sample of Southern oil counties.

This is no small finding. Beyond illustrating how individuals may cope with economic uncertainty in the absence of strong formal insurance mechanisms, this result highlights the historical importance of religious communities in solving such problems in American history (Ager et al., 2016; Ager and Ciccone, 2018; Bentzen, 2019). In doing so, it provides new insight into the economic roots of the unique religious identity of the U.S. South, and the importance of its abundant natural resources in shaping it. Given the importance of natural resources for economies around the globe, as well as the vast heterogeneity in religious engagement across space, future work should examine the extent to which this connection extends beyond our setting, to other periods and places in history.

---

period, while United Methodists grew just 29% in oil counties and 24% in non-oil counties.

References

Ager, Philipp and Antonio Ciccone, “Agricultural risk and the spread of religious communities,” *Journal of the European Economic Association*, 2018, *16* (4), 1021–1068.

_ , Casper W. Hansen, and Lars Lonstrup, “Church Membership and Social Insurance: Evidence from the Great Mississippi Flood of 1927,” *University of Southern Denmark Discussion Papers on Business and Economics No. 7/2016*, 2016.

Allcott, Hunt and Daniel Keniston, “Dutch Disease or Agglomeration? The Local Economic Effects of Natural Resource Booms in Modern America,” *Review of Economic Studies*, 2018, *85* (2), 695–731.

Anderson, Lars Harhoff and Jeanet S. Bentzen, “In the Name of God! Religiosity and the Transition to Modern Economic Growth,” *University of Copenhagen Working Paper 5/2021*, 2021.

Bartkowski, John P. and Helen A. Regis, *Charitable Choices: Religion, Race, and Poverty in the Post-welfare Era*, New York University Press, New York, NY, 2003.

Bazzi, Samuel, Andreas Ferrara, Martin Fiszbein, Thomas P. Pearson, and Patrick A. Testa, “The Other Great Migration: Southern Whites and the New Right,” *NBER Working Paper No. 29506*, 2021.

_ , Martin Fiszbein, and Mesay Gebresilasse, “Frontier Culture: The Roots and Persistence of “Rugged Individualism” in the United States,” *Econometrica*, 2020, *88* (6), 2329–2368.

Becker, Sascha O. and Ludger Woessmann, “Not the Opium of the People: Income and Secularization in a Panel of Prussian Counties,” *American Economic Review: Papers and Proceedings*, 2013, *103* (3), 539–544.

_ , Markus Nagler, and Ludger Woessmann, “Education and religious participation: city-level evidence from Germany’s secularization period 1890–1930,” *Journal of Economic Growth*, 2017, *22*, 273–311.

Bentzen, Jeanet S., “Acts of God? Religiosity and Natural Disasters Across Subnational World Districts,” *Economic Journal*, 2019, *129* (622), 2295–2321.

_ and L. Sperling, “God Politics,” *CEPR Discussion Paper No. DP14380*, 2020.

Berman, Eli, “Sect, Subsidy, and Sacrifice: An Economist’s View of Ultra-Orthodox Jews,” *Quarterly Journal of Economics*, 2000, *115* (3), 905–953.

Bettendorf, Leon J.H. and Elbert Dijkgraaf, “Religion and income: Heterogeneity between countries,” *Journal of Economic Behavior and Organization*, 2010, *74* (1-2), 12–29.

Bleakley, Hoyt and Jeffrey Lin, “Portage and Path Dependence,” *Quarterly Journal of Economics*, 2012, *127* (2), 587–644.

Brown, Stephen P.A. and Mine Yücel, “Do Higher Oil Prices Still Benefit Texas?,” *Federal Reserve Bank of Dallas, Southwest Economy*, 2004, *5*, 1–6.

Brunn, Stanley D., Gerald R. Webster, and Clark J. Archer, “The Bible Belt in a Changing South: Shrinking, Relocating, and Multiple Buckles,” *Southeastern Geographer*, 2011, *51* (4), 513–549.

Cavalcanti, Tiago V.de V., Kamiar Mohaddes, and Mehdi Raissi, “Commodity Price Volatility and the Sources of Growth,” *Journal of Applied Econometrics*, 2014, *30* (6), 857–873.

Chen, Daniel L., “Club Goods and Group Identity: Evidence from Islamic Resurgence during the Indonesian Financial Crisis,” *Journal of Political Economy*, 2010, *118* (2), 300–354.

Cnaan, Ram A., Stephanie C. Boddie, Femida Handy, Gaynor Yancey, and R. Schneider, *The Invisible Caring Hand: American Congregations and the Provision of Welfare*, New York University Press, New York, NY, 2002.

Desmet, Klaus and Romain Wacziarg, “The Cultural Divide,” *The Economic Journal*, 2021, *131* (637), 2058–88.

Dochuk, Darren, “Blessed by Oil, Cursed with Crude: God and Black Gold in the American Southwest,” *Journal of American History*, 2012, *99* (1), 51–61.

_ , *Anointed with Oil: How Christianity and Crude Made Modern America*, Basic Books, New York, NY, 2019.

Fafchamps, Marcel and Susan Lund, “Risk-Sharing Networks in Rural Philippines,” *Journal of Development Economics*, 2003, *71* (2), 261–287.

_ , Michael Koelle, and Forhad Shilpi, “Gold Mining and Proto-Urbanization: Recent Evidence from Ghana,” *Journal of Economic Geography*, 2016, *17* (5), 975–1008.

Ferrara, Andreas, Patrick A. Testa, and Liyang Zhou, “New Area- and Population-based Geographic Crosswalks for U.S. Counties and Congressional Districts, 1790-2020,” *CAGE Working Paper No. 588*, 2021.

Fishback, Price, “Social Insurance and Public Assistance in the Twentieth-Century United States,” *Journal of Economic History*, 2020, *80* (2), 311–350.

Fiszbein, Martin, Yeonha Jung, and Dietrich Vollrath, “Agrarian Origins of Individualism and Collectivism,” *Working Paper*, 2021.

Forbes, Kevin F. and Ernest M. Zampelli, “Technology and the exploratory success rate in the U.S. onshore,” *Quarterly Review of Economics and Finance*, 2002, *42* (2), 319–334.

Fouka, Vasiliki, Soumyajit Mazumder, and Marco Tabellini, “From Immigrants to Americans: Race and Assimilation During the Great Migration,” *Review of Economic Studies, forthcoming*, 2021.

Gill, Anthony and Erik Lundsgaarde, “State Welfare Spending and Religiosity: A Cross-National Analysis,” *Rationality and Society*, 2004, *16* (4), 399–436.

Glaeser, Edward L. and Bruce I. Sacerdote, “Education and Religion,” *Journal of Human Capital*, 2008, *2* (2), 188–215.

Goodman-Bacon, Andrew, “Difference-in-Differences with Variation in Treatment Timing,” *Journal of Econometrics*, 2021, *Forthcoming.*

Gorelick, Steven M., *Oil Panic and the Global Crisis - Predictions and Myths*, 1st Edition, Wiley Publishing, Oxford, UK, 2009.

Gruber, Jonathan and Daniel L. Hungerman, “Faith-Based Charity and Crowd-Out During the Great Depression,” *Journal of Public Economics*, 2007, *91* (5-6), 1043–1069.

Haines, Michael, “Historical, Demographic, Economic, and Social Data: The United States, 1790-2002,” *Inter-university Consortium for Political and Social Research [distributor], Ann Arbor, MI, 2010-05-21. https://doi.org/10.3886/ICPSR02896.v3*, 2010.

_ , Price Fishback, and Paul Rhode, “United States Agriculture Data, 1840 - 2012,” *Inter-university Consortium for Political and Social Research [distributor], Ann Arbor, MI, 2018-08-20. https://doi.org/10.3886/ICPSR35206.v4*, 2018.

Herzer, Dierk and Holger Strulik, “Religiosity and income: a panel cointegration and causality analysis,” *Applied Economics*, 2017, *49* (30), 2922–2938.

Hornbeck, Richard, “Barbed Wire: Property Rights and Agricultural Development,” *Quarterly Journal of Economics*, 2010, *125* (2), 767–810.

Hungerman, Daniel L., “Are Church and State Substitutes? Evidence from the 1996 Welfare Reform,” *Journal of Public Economics*, 2005, *989* (1-2), 2245–2267.

Iannaccone, Laurence, “Sacrifice and Stigma: Reducing Free-Riding in Cults, Communes, and Other Collectives,” *Journal of Political Economy*, 1992, *100* (2), 271–291.

_ , “Why Strict Churches Are Strong,” *American Journal of Sociology*, 1994, *99* (5), 1180–1211.

_ , “Introduction to the Economics of Religion,” *Journal of Economic Literature*, 1998, *36* (3), 1465–1495.

Iyer, Sriya, “The New Economics of Religion,” *Journal of Economic Literature*, 2016, *54* (2), 395–441.

Marx, Karl, *Zur Kritik der Hegelschen Rechtsphilosophie: Einleitung*, in Deutsch Französische Jahrbücher, edited by Arnold Ruge and Karl Marx, pp. 71–85. Paris: Bureau der Jahrbücher, 1844.

Matheis, Michael, “Local Economic Impacts of Coal Mining in the United States 1870 to 1970,” *Journal of Economic History*, 2016, *76* (4), 1152–1181.

McCleary, Rachel M. and Robert J. Barro, “Religion and Economy,” *Journal of Economic Perspectives*, 2006, *20* (2), 49–72.

Michaels, Guy, “The Long Term Consequences of Resource-Based Specialisation,” *Economic Journal*, 2011, *121* (551), 31–57.

Oil and Gas Journal Data Book, Pennwell Publishing Co, Tulsa, OK, 2000.

Olien, Roger M. and Diana Davids Olien, *Oil and Ideology: The Cultural Creation of the American Petroleum Industry*, University of North Carolina Press, Chapel Hill, NC, 2000.

Oster, Emily, “Unobservable Selection and Coefficient Stability: Theory and Evidence,” *Journal of Business & Economic Statistics*, 2019, *37* (2), 187–204.

Pew Research Center, “Religion and Views on Climate and Energy Issues,” *October 22, 2015: https://www.pewresearch.org/science/2015/10/22/religion-and-views-on-climateand-energy-issues/. Accessed on March*

10, 2020, 2015.

—, "What does the federal government spend your tax dollars on? Social insurance programs, mostly," *April 4, 2017: https://www.pewresearch.org/fact-tank/2017/04/04/what-does-the-federal-government-spend-your-tax-dollars-on-social-insurance-programs-mostly/. Accessed on July 23, 2020*, 2017.

—, "The Age Gap in Religion Around the World," *June 13, 2018: https://www.pewforum.org/2018/06/13/why-do-levels-of-religious-observance-vary-by-age-and-country/. Accessed on August 27, 2020*, 2018.

Raz, Itzchak Tzachi, "Learning is Caring: Soil Heterogeneity, Social Learning and the Formation of Close-knit Communities," *Hebrew University of Jeruselum Working Paper*, 2021.

Richardson, Gary, "The Prudent Village: Risk Pooling Institutions in Medieval English Agriculture," *Journal of Economic History*, 2005, *65* (2), 386–413.

Ross, Michael L., *The Oil Curse: How Petroleum Wealth Shapes the Development of Nations*, Princeton University Press, Princeton, NJ, 2012.

Ruggles, Steve, Sarah Flood, Ronald Goeken, Josiah Grover, Erin Meyer, Jose Pacas, and Matthew Sobek, "IPUMS USA: Version 10.0 [dataset]," *IPUMS [distributor], Minneapolis, MN. https://doi.org/10.18128/D010.V10.0*, 2018.

Scheve, Kenneth and David Stasavage, "Religion and Preferences for Social Insurance," *Quarterly Journal of Political Science*, 2006, *1* (3), 255–286.

Schmidheiny, Kurt and Sebastian Siegloch, "On Event Study Designs and Distributed-Lag Models: Equivalence, Generalization and Practical Implications," *IZA DP No. 12079*, 2019.

Sun, Liyang and Sarah Abraham, "Estimating Dynamic Treatment Effects in Event Studies with Heterogeneous Treatment Effects," *Journal of Econometrics*, 2021, *225* (2), 175–199.

Townsend, Robert M., "Consumption Insurance: An Evaluation of Risk-Bearing Systems in Low-Income Economies," *Journal of Economic Perspectives*, 1995, *9* (3), 83–102.

U.S. Census Bureau, *Census of Religious Bodies*, U.S. Government Printing Office, Washington, DC, 1936.

U.S. Department of Energy, *Oil and Gas Field Code Master List*, Energy Information Administration, Washington, DC, 2004.

van der Ploeg, Frederick, "Natural Resources: Curse or Blessing?," *Journal of Economic Literature*, 2011, *49* (2), 366–420.

— **and Steven Poelhekke**, "Volatility and the Natural Resource Curse," *Oxford Economic Papers New Series*, 2009, *61* (4), 727–760.

— **and** —, "The Pungent Smell of "Red Herrings": Subsoil Assets, Rents, Volatility and the Resource Curse," *Journal of Environmental Economics and Management*, 2010, *60* (1), 44–55.

Venables, Anthony J., "Using Natural Resources for Development: Why Has It Proven So Difficult?," *Journal of Economic Perspectives*, 2016, *30* (1), 161–184.

Wilson, John and Thomas Janoski, "The Contribution of Religion to Volunteer Work," *Sociology of Religion*, 1995, *56* (2), 137–152.

Tables

Table 1: Oil Abundance and Religious Participation

	Outcome: Membership in major Christian churches (% population)					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	6.627^{***}	4.801^{***}	8.126^{***}	6.679^{***}	8.055^{***}	7.282^{***}
	(0.921)	(0.892)	(1.023)	(0.908)	(1.011)	(0.953)
Sample	Donut	Full	Donut	Donut	Donut	Donut
Log pop. control	No	No	Yes	No	Yes	Yes
Compositional controls	No	No	No	Yes	Yes	Yes
Other controls	No	No	No	No	No	Yes
Observations	4574	6808	4574	4574	4574	4563
Counties	520	774	520	520	520	520
Adj. R^2	0.761	0.750	0.765	0.764	0.767	0.772
Outcome mean	32.88	33.43	32.88	32.88	32.88	32.87
Oster's δ	1.703	1.544	2.157	1.727	2.107	1.753

Note: Estimates are from difference-in-differences regressions of membership in 15 major, mainstream Christian denominations (% population) in county c in year t on an indicator called "oil abundance," which equals one for a county in years following a major oil discovery and is zero otherwise. Major oil discoveries are defined as oilfields holding 100 million barrels of oil or more. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, covering the nine church and religious censuses held between 1890 and 1990. Counties that are adjacent to counties which eventually discover oil are excluded in the *donut* sample to limit spillover effects that might dilute the treatment, while the full sample includes those neighboring counties. All regressions include county and sample year fixed effects. Column (3) includes the county log population control. Column (4) includes compositional population controls, including the percent Black, percent French, percent Italian, and percent German population. Column (5) includes all five population controls. Column (6) includes population controls as well as the share of land in agriculture used for cotton production, the share of agricultural and manufacturing employment. The bottom row of the table reports the δ statistic by Oster (2019) which indicates how much selection on unobservables than on observables would have to play a role in order to explain away the oil abundance effect. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 2: Top Biblical and Non-Biblical Names in 1940

Biblical names		Non-Biblical names	
Name	Share	Name	Share
1. John	2.5%	1. William (Bill)	2.5%
2. Mary	2.2%	2. Charles (Carl)	1.6%
3. James (Jim)	2.0%	3. Robert (Bob)	1.4%
4. Elizabeth (Betty)	1.7%	4. George	0.9%
5. Joseph	1.1%	5. Wilhelmina (Wilma)	0.8%
6. Anna	1.0%	6. Frank	0.7%
7. Thomas	0.8%	7. Edward (Ned)	0.7%
8. Sarah (Sadie)	0.5%	8. Lillian	0.6%
9. Ruth	0.5%	9. Catherine (Kate, Kitty)	0.5%
10. Lois (Louis, Lewis)	0.5%	10. Frances (Fanny)	0.5%
Share of sample in top ten	12.7%	Share of sample in top ten	10.1%

Note: Tables shows the rank and frequency of the top ten most common Biblical and non-Biblical names in the 1940 names sample from the U.S. Census, using the list of names in the Bible from behindthename.com. Names are based on groupings of phonetically similar names as determined by both the NYSIIS and Soundex phonetic coding algorithms. A few popular names—such as Margaret, Henry, and Dorothy, which each make up just under a percent of the names sample—are designated as religious names by Soundex but not by NYSIIS. The reverse is not true. None of these have Biblical or otherwise Christian roots. On this basis, we treat NYSIIS as our preferred phonetic device but report estimates from both for completeness.

Table 3: Does Selective Migration Drive Effects?

	Outcome: Household head moved to oil-abundant county, 1935-40					
	(1)	(2)	(3)	(4)	(5)	(6)
Biblical name	-0.001***	-0.001***	-0.001**	-0.001*	-0.001*	-0.001
	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)	(0.000)
HH head works in oil		0.199***		0.199***		0.170***
		(0.007)		(0.007)		(0.003)
Observations	868740	868740	868740	868740	1567559	1567559
Families	399,502	399,502	399,502	399,502	735,904	735,904
Adj. R^2	0.010	0.029	0.010	0.029	0.030	0.045
Outcome mean	0.022	0.022	0.022	0.022	0.082	0.082
Oil discovery period	1935-40	1935-40	1935-40	1935-40	Before 1935	Before 1935
Age in 1940	5-18	5-18	5-18	5-18	5-18	5-18
Phonetic algorithm	NYSIIS	NYSIIS	Soundex	Soundex	NYSIIS	NYSIIS

Note: Estimates are from regressions of an indicator for whether a child's household head moved between 1935 and 1940 to an oil-abundant county on an indicator for whether that child had a Biblical name. A county is considered oil abundant if it lies above a known major oil field, holding 100 million barrels of oil or more. The sample consists of all unmarried children between the ages of 5 and 18 in 1940 who reside as of that year in counties in Louisiana, Oklahoma, and Texas. We exclude counties that are adjacent to known oil counties to limit spillover effects that might dilute the treatment. All regressions include state fixed effects. We also control for dummies for the child's age, race, sex, and place of birth. Standard errors are clustered at the family level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 4: Do Changes in Education, Income, or Population Explain Treatment Effects?

	Outcome: Membership in major Christian churches (% population)				
	Education (1)	Income (2)	Population size (3)	Population composition (4)	All (5)
Oil abundance	6.022^{***}	6.414^{***}	8.126^{***}	6.679^{***}	7.494^{***}
	(0.959)	(1.022)	(1.023)	(0.908)	(1.033)
Education	-0.163^{**}				-0.092
	(0.069)				(0.066)
Income		-0.000^{**}			0.000
		(0.000)			(0.000)
Log population			-2.679^{***}		-3.478^{***}
			(0.622)		(0.803)
% Black				0.010	0.071
				(0.043)	(0.055)
% French				10.022^{***}	18.026^{***}
				(3.325)	(6.777)
% Italian				1.547^{**}	0.895
				(0.770)	(1.031)
% German				-0.835^{***}	0.184
				(0.269)	(0.479)
Observations	4122	3939	4574	4574	3939
Counties	520	520	520	520	520
Adj. R^2	0.773	0.777	0.765	0.764	0.785
Outcome mean	34.12	34.24	32.88	32.88	34.24

Note: Estimates are from difference-in-differences regressions of membership in 15 major, mainstream Christian denominations (% population) in county c in year t on an indicator called "oil abundance," which equals one for a county in years following a major oil discovery and is zero otherwise. Major oil discoveries are defined as oilfields holding 100 million barrels of oil or more. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, covering the nine church and religious censuses held between 1890 and 1990. We exclude counties that are adjacent to oil counties to limit spillover effects that might dilute the treatment. All regressions include county and sample year fixed effects. To measure educational attainment, we use the percentage of adults (aged 25 plus) with a high school degree. As these not available for 1926 and prior, we use the literacy rate for those years. To measure income, we use median family income in 2018 dollars. As income data are not available prior for 1926 and prior, we rely on another commonly used metric for those years based on occupational income scores. This score is computed based on the median income in a given occupation in 1950, assigned to individuals' occupations in previous years to proxy for their income. For the 1906, 1916, and 1926 religious census years, we average those income scores by county based on the full-count Census files. Estimates using only the period from 1936 are qualitatively similar and available upon request. Columns (3) and (5) control for log population to capture factors associated with urbanization and market size. Columns (4) and (5) control for compositional factors associated with migration. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 5: Heterogeneous Effects: Oil Price Volatility

	Outcome: Membership in major Christian churches (% population)					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	5.499***	5.182***	5.015***	5.674***	5.578***	4.618***
	(0.884)	(0.885)	(0.890)	(1.031)	(0.920)	(0.946)
Oil × 5 yr price s.d.	0.130***					
	(0.044)					
Oil × 10 yr price s.d.		0.106***				
		(0.032)				
Oil × 25 yr price s.d.			0.119***			
			(0.035)			
Oil × 5 yr log price s.d.				4.966		
				(3.781)		
Oil × 10 yr log price s.d.					3.066**	
					(1.445)	
Oil × 25 yr log price s.d.						4.949***
						(1.592)
Observations	4574	4574	4574	4574	4574	4574
Counties	520	520	520	520	520	520
Adj. R^2	0.761	0.761	0.761	0.761	0.761	0.761
Outcome mean	32.88	32.88	32.88	32.88	32.88	32.88
Interaction sample st. dev.	6.716	10.247	10.202	.085	.205	.212

Note: Estimates are from difference-in-differences regressions of membership in 15 major, mainstream Christian denominations (% population) in county c in year t on an indicator called “oil abundance,” which equals one for a county in years following a major oil discovery and is zero otherwise. Major oil discoveries are defined as oilfields holding 100 million barrels of oil or more. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, covering the nine church and religious censuses held between 1890 and 1990. We exclude counties that are adjacent to oil counties to limit spillover effects that might dilute the treatment. All regressions include county and sample year fixed effects. The additional regressors include interactions of the oil abundance indicator with the standard deviation of world per barrel real (2018 USD) oil prices (columns 1-3) and of the log world oil price (columns 4-6) over 5, 10, and 25 years as measures of income risk associated with oil. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 6: Heterogeneous Effects: Access to Public Social Insurance

	Outcome: Membership in major Christian churches (% population)					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	10.551***	10.412***	11.614***	12.166***	10.298***	11.548***
	(2.591)	(2.115)	(2.408)	(2.465)	(2.132)	(2.327)
Oil × State max weekly UI	-7.055**	-21.319***	-13.130***			
	(2.971)	(4.229)	(3.413)			
State max weekly UI	0.521	-6.325**	-0.722			
	(2.055)	(2.718)	(2.561)			
Oil × State max weekly WC				-8.261***	-17.486***	-11.212***
				(1.877)	(3.100)	(2.333)
State max weekly WC				-8.109***	-7.091***	-6.212***
				(1.948)	(2.582)	(2.371)
Observations	2600	2600	2600	2544	2544	2544
Counties	520	520	520	520	520	520
Adj. R^2	0.791	0.795	0.792	0.791	0.791	0.790
Outcome mean	39.58	39.58	39.58	40.09	40.09	40.09
Adj for poverty line	Yes			Yes		
Adj for per capita income		Yes			Yes	
Adj for mfg income			Yes			Yes

Note: Estimates are from difference-in-differences regressions of membership in 15 major, mainstream Christian denominations (% population) in county c in year t on an indicator called "oil abundance," which equals one for a county in years following a major oil discovery and is zero otherwise. Major oil discoveries are defined as oilfields holding 100 million barrels of oil or more. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, during the post-New Deal period when states had developed social insurance systems. We exclude counties that are adjacent to oil counties to limit spillover effects that might dilute the treatment. All regressions include county and sample year fixed effects. The additional regressors are interactions of the oil abundance indicator with measures of state-level maximum weekly unemployment insurance and workers compensation benefits from Fishback (2020), beginning in 1940 through 1990, each of which is adjusted to reflect relative value, by (i) national poverty line weekly income equivalent for a 4 person family, (ii) average weekly earnings plus benefits for manufacturing workers in the state, and (iii) state weekly per capita income, respectively. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 7: Heterogeneous Effects: Access to Credit and Private Insurance

	Outcome: Membership in major Christian churches (% population)					
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
Oil abundance	9.342***	10.067***	7.614***	7.781***	13.354***	13.442***
	(1.259)	(2.952)	(1.003)	(1.018)	(1.697)	(1.744)
Oil × Any savings & loans banks, 1950	-5.625***	-9.814**				
	(1.581)	(4.337)				
Oil × Any bank tellers, 1910			-6.330***	-6.508***		
			(2.037)	(2.161)		
Oil × Any insurance agents, 1910					-8.997***	-8.830***
					(1.859)	(1.905)
Observations	4574	3390	4529	4426	4529	4426
Counties	520	382	513	501	513	501
Adj. R^2	0.762	0.761	0.763	0.761	0.765	0.763
Outcome mean	32.88	31.83	32.73	32.71	32.73	32.71
Drops counties treated $\leq$ 1950?	No	Yes				
Drops counties treated $\leq$ 1910?			No	Yes	No	Yes

Note: Estimates are from difference-in-differences regressions of membership in 15 major, mainstream Christian denominations (% population) in county c in year t on an indicator called "oil abundance," which equals one for a county in years following a major oil discovery and is zero otherwise. Major oil discoveries are defined as oilfields holding 100 million barrels of oil or more. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, covering the nine church and religious censuses held between 1890 and 1990. We exclude counties that are adjacent to oil counties to limit spillover effects that might dilute the treatment. All regressions include county and sample year fixed effects. We interact the oil abundance indicator with indicators for alternative insurance possibilities such as banks and private insurance companies. Those include dummies for whether a county had any savings and loan associations in 1950, or whether there were any bank teller or insurance agents in the county in 1910. The latter two variables come from the full count Census of 1910, while data on savings and loan associations were not available in the U.S. Census County Data Books until the mid-19th century. To minimize bad control concerns, secondary specifications in all columns (b) exclude counties treated prior to the year of the interaction term. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table 8: Oil Price Shocks, Religious Communities, and Labor Outcomes, 1940-90

	(1) % Unemployed	(2) Log pop density	(3) % Mining workers	(4) % Agricultural workers	(5) % Mfg workers
Oil × Oil price increase	-0.035***	0.002***	-0.008	0.123***	-0.030***
	(0.003)	(0.001)	(0.006)	(0.014)	(0.008)
Oil × Oil price increase × Above-median Christian, 1936	0.011***	-0.002**	0.007	-0.073***	0.020*
	(0.004)	(0.001)	(0.009)	(0.019)	(0.012)
Observations	2718	2718	2714	2716	2718
Counties	453	453	453	453	453
Adj. R^2	0.556	0.946	0.821	0.810	0.759
Outcome mean	5.351	2.317	2.698	21.93	12.97
Drops counties treated $\geq$ 1936?	Yes	Yes	Yes	Yes	Yes

Note: Estimates are from regressions of county-level economic outcomes in county c in year t on an "oil" indicator which equals one if and only if a county lies above an oilfield holding 100 million barrels of oil or more. This oil dummy is interacted with a time-varying measure of world per barrel crude oil prices (in 2018 USD). This is in term interacted with a time-invariant indicator of whether a county was above the sample median in Christian membership in 1936. This year is chosen because most outcomes are reported in the U.S. Census County Data Books beginning in 1940. We therefore exclude counties treated after 1936, which as such may see increases in Christian participation later in the panel thus biasing estimates. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, and is decadal from 1940 to 1990. We exclude counties that are adjacent to oil counties to limit spillover effects that might dilute the treatment. All regressions include county and sample year fixed effects. Outcomes include log population density, shares of labor force in mining, agriculture, and manufacturing, and the unemployment rate. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Figures

Figure 1: Real Crude Oil Prices, 1861 to 2000

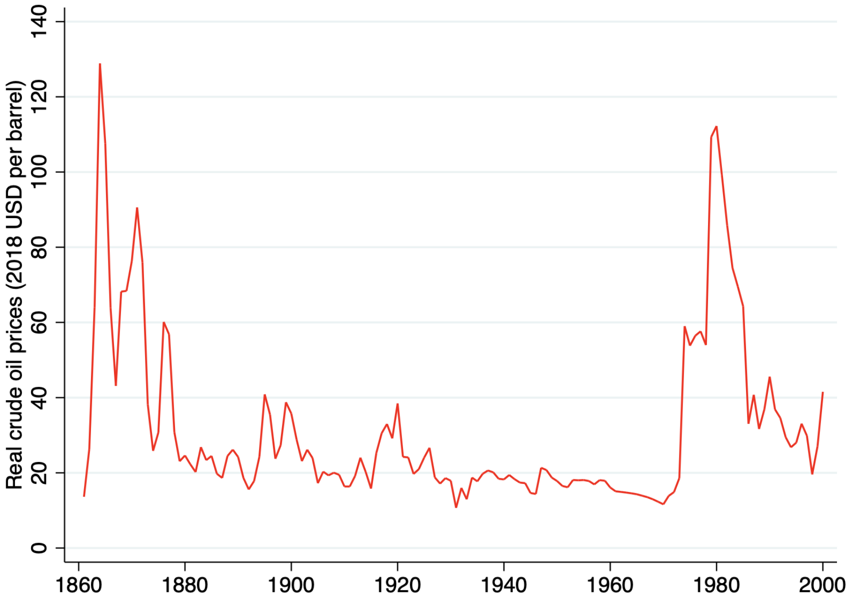

Note: Prices are expressed in 2018 USD per barrel. Prices from 1861 to 1944 are U.S. average spot prices, 1945 to 1983 are Arabian Light prices, and 1984 to 2010 are Brent dated prices. Oil price data was compiled by BP and collected from Quandl at https://www.quandl.com/data/BP/CRUDE_OIL_PRICES (date retrieved: July 27, 2020).

Figure 2: Map of All Oil-Abundant Counties in the Sample

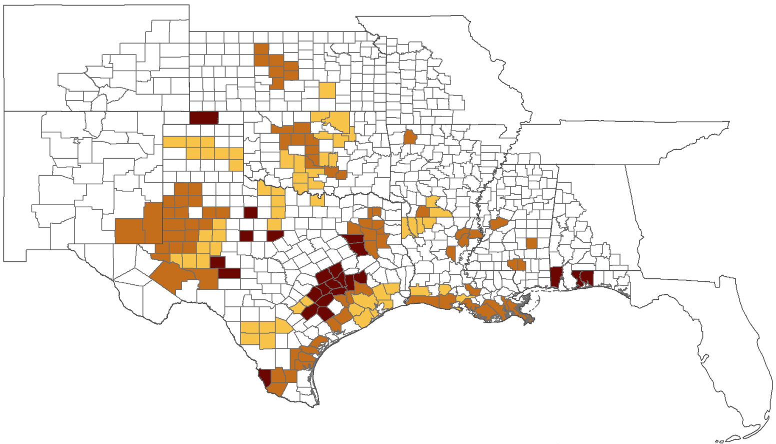

Note: Major oilfield with ≥ 100 million barrels discovered in each county (i) 1893-1925 (light orange), (ii) 1926-50 (medium orange), (iii) 1951-82 (dark orange). White indicates no major oilfields. Counties included in the sample are outlined. These are limited to counties within 200 km of oil-abundant counties in Louisiana, Oklahoma and Texas as in Michaels (2011) to limit the geographic heterogeneity of the sample.

Figure 3: Oil Discoveries and Membership in the Major Christian Churches, 1890 to 1990

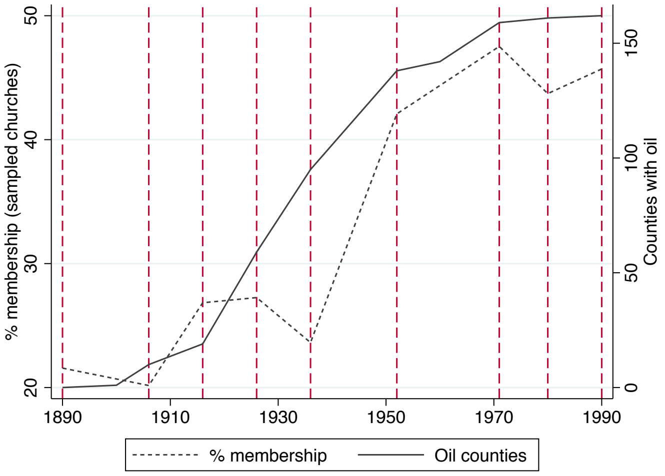

Note: The solid line plots the number of treated counties, based on year of a county's first major (≥ 100 million barrels) oilfield discovery (right axis), spanning 1893 to 1982. The dashed lined plots the evolution of membership (% population) in sampled Christian churches (left axis). Membership generally entails baptism or confirmation and is a strictest definition of religious participation. Note that the 1936 religious census underreported some Baptist and Methodist groups (Ager et al., 2016). Results are not sensitive to dropping 1936 (see Figure A2). Also note the rise in membership after 1936, which reflects the growth of evangelicalism in the South and dramatic rise in church attendance after WWII (Pew Research Center, 2018). Red dashed lines indicate the religious censuses in the sample, the first in 1890 and the last in 1990.

Figure 4: Oil and Religion: Event Study Plots

(a) Baseline result

(b) Partialling out pre-trends

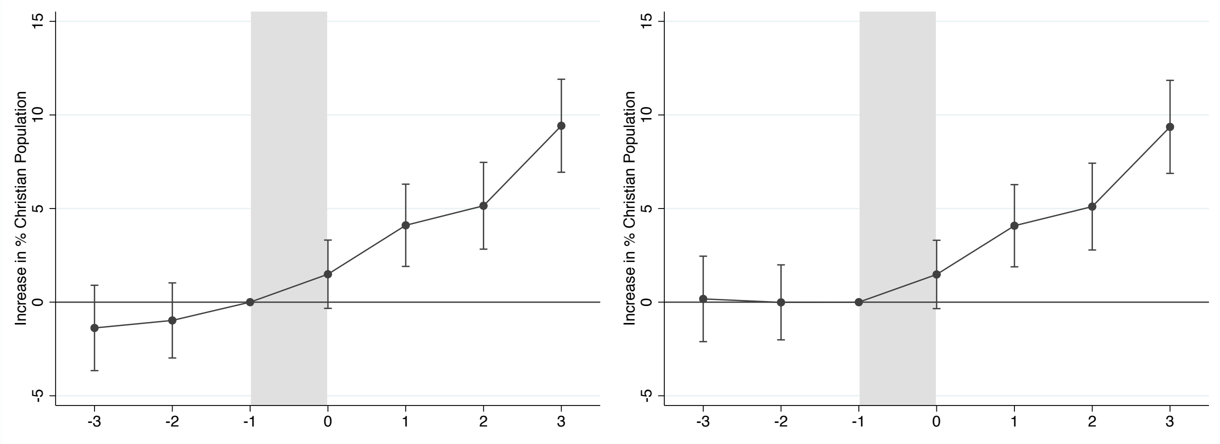

(c) Matching on pre-oil log pop. density growth

(d) Using Sun and Abraham (2021)'s estimator

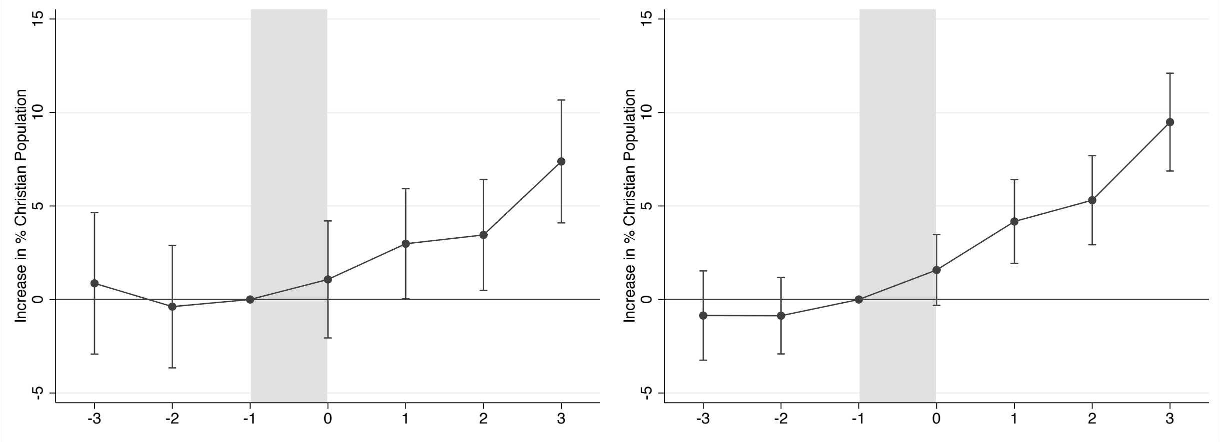

Note: Coefficient plots from event-study difference-in-differences analyses that regress membership in 15 major, mainstream Christian denominations (% population) in a county on both year and county fixed effects as well as an indicator for a major oil discovery in the county interacted with event time fixed effects. Panel (d) adopts the estimator proposed by Sun and Abraham (2021) to remove contamination from other treatment timing cohorts in the presence of heterogeneous treatment timing. Major oil discoveries are defined as oilfields holding 100 million barrels of oil or more. Event time is defined as the three periods before and after the occurrence of the first major oil discovery. The omitted baseline period is $t = -1$, which is the last pre-treatment period. The gray shaded area indicates the time frame within which oil is discovered between $t = -1$ and $t = 0$. The sample consists of counties in Louisiana, Oklahoma, and Texas as well as surrounding counties in Alabama, Arkansas, Colorado, Florida, Kansas, Mississippi, Missouri, New Mexico, and Tennessee, covering the nine church and religious censuses held between 1890 and 1990. We exclude counties that are adjacent to oil counties to limit spillover effects that might dilute the treatment. Standard errors are clustered at the county level and error bars represent 95% confidence intervals.

Figure 5: Spatial Distribution of Denominations in 1916

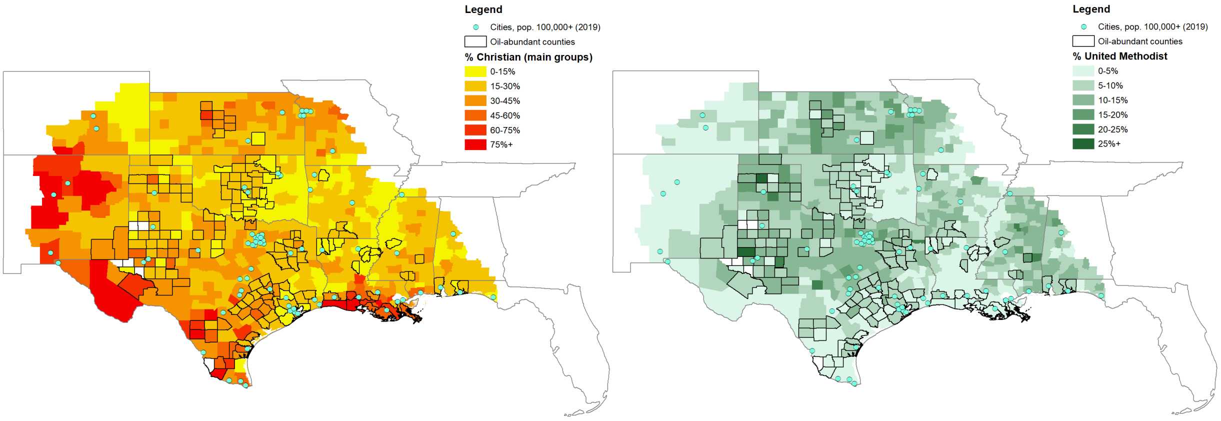

(a) All sample Christian churches (% population)

(b) United Methodist (% population)

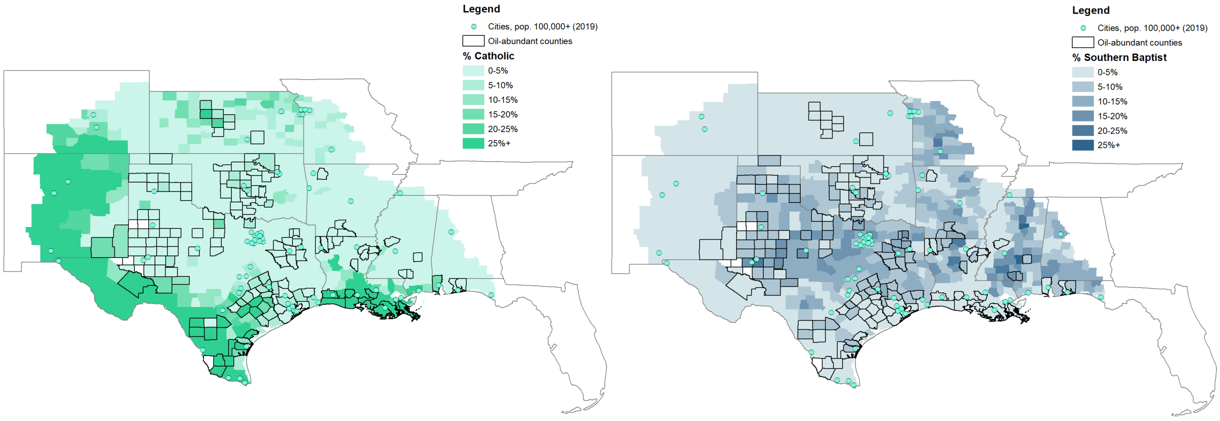

(c) Roman Catholic (% population)

(d) Southern Baptist (% population)

Note: Maps show the spatial distribution of different Christian denominations as a share of the total population in our sample counties, as reported in the 1916 United States Census of Religious Bodies. Oil-abundant counties are outlined in black, while urban areas (cities with population >100,000 in 2019) are dotted in light blue. Note the sudden decline in Southern Baptists at the Kansas border, which generally marks the edge of the Bible Belt. City population and longitude-latitude data from SimpleMaps.com at https://simplemaps.com/data/us-cities (date retrieved: August 20, 2020).