Online Appendix for
"Churches as Social Insurance: Oil Risk and Religion in the U.S. South"

Andreas Ferrara* Patrick A. Testa†

January 5, 2022

*University of Pittsburgh, Department of Economics. Email: a.ferrara@pitt.edu.

†Tulane University, Department of Economics. Email: ptesta@tulane.edu.

List of Tables

List of Figures

A Additional Tables and Figures

Table A1: Oil Abundance and Religious Participation: Robustness to Spillover Effects

	Outcome: Membership in major Christian churches (% population)				
	(1)	(2)	(3)	(4)	(5)
Oil abundance	4.801***	6.627***	6.623***	6.966***	7.536***
	(0.892)	(0.921)	(0.922)	(0.956)	(1.039)
Observations	6808	4574	4565	3672	2778
Counties	774	520	519	418	317
Adj. R^2	0.750	0.761	0.761	0.757	0.751
Outcome mean	33.43	32.88	32.89	33.74	35.06
Excl. control counties	None	Adjacent	<50km	<100km	<150km

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. All regressions include county and sample year fixed effects. Regressions utilize different "donut" approaches to remove spillover effects. Column (1) excludes no counties. Column (2) excludes control counties that are adjacent to an oil county. The remaining columns exclude control counties within a certain distance threshold of 50, 100, and 150 kilometers (approximately 31, 62, and 93 miles) to the nearest oil county to test for sensitivity with respect to potential spillover effects. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A2: Oil Abundance and Religious Participation: Robustness to Pre-Oil Population Trends

	Outcome: Membership in major Christian churches (% population)					
	Panel a: Sample matched on pre-oil log population growth					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	5.718***	5.718***	2.776*	2.770*	6.069***	6.054***
	(1.374)	(1.374)	(1.506)	(1.508)	(1.542)	(1.546)
Observations	1811	1811	1811	1811	1811	1811
Counties	194	194	194	194	194	194
Adj. R^2	0.764	0.764	0.528	0.528	0.595	0.595
Outcome mean	36.54	36.54	36.54	36.54	36.54	36.54
County FE	yes	yes				
Pair FE			yes	yes	yes	yes
Propensity score control		yes		yes		yes
Geographic region control					yes	yes
	Panel b: Sample matched on pre-oil log population density growth					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	4.077***	4.077***	4.011***	3.728***	3.341**	3.142**
	(1.380)	(1.380)	(1.316)	(1.306)	(1.366)	(1.370)
Observations	1964	1964	1964	1964	1964	1964
Counties	201	201	201	201	201	201
Adj. R^2	0.761	0.761	0.565	0.566	0.595	0.599
Outcome mean	35.34	35.34	35.34	35.34	35.34	35.34
County FE	yes	yes				
Pair FE			yes	yes	yes	yes
Propensity score control		yes		yes		yes
Geographic region control					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, 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. Oil counties were matched to non-oil counties via propensity score matching using the log population growth (panel a) and log population density growth (panel b) in the pre-oil discovery years. All regressions include sample year fixed effects. We control for county fixed effects in columns (1) and (2), and for matched pair fixed effects in columns (3) and (4). Columns (2) and (4) also condition on the propensity score that was estimated by the matching algorithm. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A3: Oil Abundance and Religious Participation: Robustness to Population Trends

	Outcome: Membership in major Christian churches (% population)					
	Panel a: Sample matched on log population growth					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	5.307***	5.307***	2.598*	2.598*	3.870**	3.875**
	(1.226)	(1.226)	(1.514)	(1.513)	(1.525)	(1.536)
Observations	1695	1695	1695	1695	1695	1695
Counties	189	189	189	189	189	189
Adj. R^2	0.749	0.749	0.507	0.507	0.540	0.540
Outcome mean	36.11	36.11	36.11	36.11	36.11	36.11
County FE	yes	yes				
Pair FE			yes	yes	yes	yes
Propensity score control		yes		yes		yes
Geographic region control					yes	yes
	Panel b: Sample matched on log population density growth					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	4.474***	4.474***	4.935***	5.140***	4.985***	5.206***
	(1.206)	(1.206)	(1.410)	(1.391)	(1.456)	(1.438)
Observations	1650	1650	1650	1650	1650	1650
Counties	184	184	184	184	184	184
Adj. R^2	0.752	0.752	0.553	0.559	0.569	0.572
Outcome mean	34.92	34.92	34.92	34.92	34.92	34.92
County FE	yes	yes				
Pair FE			yes	yes	yes	yes
Propensity score control		yes		yes		yes
Geographic region control					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, 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. Oil counties were matched to non-oil counties via propensity score matching using the log population growth (panel a) and log population density growth (panel b) over all years in the sample. All regressions include sample year fixed effects. We control for county fixed effects in columns (1) and (2), and for matched pair fixed effects in columns (3) and (4). Columns (2) and (4) also condition on the propensity score that was estimated by the matching algorithm. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A4: Oil Abundance and Religious Participation: Robustness to Treatment Year

	Outcome: Membership in major Christian churches (% population)			
	(1)	(2)	(3)	(4)
Oil abundance	6.658***	6.202***	6.511***	7.147***
	(0.886)	(0.916)	(0.921)	(1.123)
Observations	4574	4574	4574	4574
Counties	520	520	520	520
Adj. R^2	0.760	0.759	0.760	0.759
Outcome mean	32.88	32.88	32.88	32.88
Treatment defined by	Earliest oil	Adjacent county	Oldest oilfield	Earliest, contig group

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. Regressions vary the year in which treatment turns on. Column (1) considers an oil-abundant county to be treated the year *any* oil was discovered there, even if not from a major oilfield. Column (2) considers an oil-abundant county to be treated the year it *or* an adjacent oil-abundant county discovered a major oilfield, whichever happened first. Column (3) considers an oil-abundant county to be treated the first year any of its major oilfields were discovered anywhere, even if not in that county. Column (4) 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 discovered a major oilfield. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A5: Oil Abundance and Religious Participation: Robustness to Spatial Autocorrelation

	Outcome: Membership in major Christian churches (% population)			
	(1)	(2)	(3)	(4)
Oil abundance	6.627***	6.627***	6.627***	6.627***
	(0.622)	(0.632)	(0.746)	(0.831)
Observations	4574	4574	4574	4574
Counties	520	520	520	520
Adj. R^2	0.139	0.139	0.139	0.139
Outcome mean	32.88	32.88	32.88	32.88
Distance cutoff	25km	50km	100km	150km

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. Conley standard errors that adjust inference for spatial autocorrelation in parentheses with the distance cutoff being reported in kilometers in the bottom table row. Distance cutoffs are 25, 50, 100, and 150 kilometers (approximately 15.5, 31, 62, and 93 miles). Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A6: Oil Abundance and Religious Participation: Robustness to Alternative Clustering

	Outcome: Membership in major Christian churches (% population)		
	(1)	(2)	(3)
Oil abundance	6.627***	6.627***	6.627***
	(0.921)	(1.057)	(1.583)
Observations	4574	4574	4574
Counties	520	453	382
Adj. R^2	0.761	0.761	0.761
Outcome mean	32.88	32.88	32.88
S.E. clustered by	County	Oilfield	Contiguous oil counties

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. Regressions consider alternative ways in which to cluster standard errors. Column (1) clusters standard errors at the county level, at which the treatment is defined. Column (2) clusters standard errors at the major oilfield level, where a county is assigned whichever of its major oilfields was discovered first in any county, as treatment is likely to have occurred in practice on that basis. Column (3) clusters standards errors by the set of contiguous oil-abundant counties in which it exists, even if they share no common major oilfield, as treatment in practice may span beyond the major oilfield level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A7: Oil Abundance and Religious Participation: Alternative Imputation Procedures

	Outcome: Membership in major Christian churches (% population)			
	(1)	(2)	(3)	(4)
Oil abundance	5.894***	6.495***	6.627***	3.930***
	(0.928)	(0.922)	(0.921)	(0.882)
Observations	4085	4497	4574	4574
Counties	520	520	520	520
Adj. R^2	0.761	0.761	0.761	0.784
Outcome mean	34.37	33.24	32.88	22.12
Procedure	Excludes 1906	Imputed from 1890	Imputed from 1906	Excludes Baptists

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. Regressions reflect different approaches to dealing with the aggregation of Northern, Southern, and National Baptists in the 1906 Religious Census. Column (1) simply excludes the 1906 Religious Census from the sample. Column (2) uses the ratio of Northern+Southern to National Baptists from the 1890 Religious Census to impute values for 1906. Column (3) uses the ratio from the 1916 Religious Census to impute values for 1906. Column (4) simply excludes both Baptists groups from the sample. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A8: Does Selective Migration Drive Effects? Movers-Only Sample

	Outcome: Household head moved to oil-abundant county, 1935-40					
	(1)	(2)	(3)	(4)	(5)	(6)
Biblical name	-0.004*	-0.004*	-0.003	-0.003	0.002	0.002
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
HH head works in oil		0.440***		0.440***		0.407***
		(0.013)		(0.013)		(0.004)
Observations	128152	128152	128152	128152	237921	237921
Families						
Adj. R^2	0.029	0.070	0.029	0.070	0.029	0.078
Outcome mean	0.152	0.152	0.152	0.152	0.543	0.543
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 and who across counties between 1935 and 1940. 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 A9: Does Selective Migration Drive Effects? NYSIIS-Soundex Agreement Sample

	Outcome: Household head moved to oil-abundant county, 1935-40			
	(1)	(2)	(3)	(4)
Biblical name	-0.001***	-0.001***	-0.001**	-0.001*
	(0.000)	(0.000)	(0.001)	(0.001)
HH head works in oil		0.200***		0.169***
		(0.007)		(0.003)
Observations	671454	671454	1218713	1218713
Families				
Adj. R^2	0.010	0.030	0.031	0.046
Outcome mean	0.022	0.022	0.083	0.083
Oil discovery period	1935-40	1935-40	Before 1935	Before 1935
Age in 1940	5-18	5-18	5-18	5-18

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 A10: Does Selective Migration Drive Effects? Children Aged 5-10 Sample

	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
	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)	(0.001)
HH head works in oil		0.239***		0.239***		0.192***
		(0.008)		(0.008)		(0.003)
Observations	374698	374698	374698	374698	678189	678189
Families						
Adj. R^2	0.012	0.039	0.012	0.039	0.039	0.058
Outcome mean	0.026	0.026	0.026	0.026	0.095	0.095
Oil discovery period	1935-40	1935-40	1935-40	1935-40	Before 1935	Before 1935
Age in 1940	5-10	5-10	5-10	5-10	5-10	5-10
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 10 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 A11: Negative Oil Shocks Hurt Oil-abundant Counties

	(1) % Mining workers	(2) Log mean mining wage	(3) Log mining per worker output	(4) % Agricultural workers
Oil abundance	2.968***	1.018***	1.267***	-3.993**
	(0.414)	(0.165)	(0.409)	(1.755)
Oil × Oil price increase	0.005	0.062***	0.062**	0.062***
	(0.004)	(0.010)	(0.025)	(0.008)
N	3116	928	926	3118
Counties	520	271	270	520
Adj. R^2	0.816	0.793	0.841	0.807
Outcome mean	3.005	10.66	12.59	22.13

	(1) % Mfg workers	(2) Log mean mfg wage	(3) % Unemployed	(4) Median family income
Oil abundance	-2.615***	0.058**	-0.079	2934.834***
	(0.780)	(0.028)	(0.273)	(846.442)
Oil × Oil price increase	-0.019***	0.001**	-0.028***	61.964***
	(0.005)	(0.000)	(0.002)	(18.098)
N	3519	4211	3640	2941
Counties	520	471	520	520
Adj. R^2	0.727	0.848	0.551	0.925
Outcome mean	11.63	10.14	4.869	36651.6

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. This indicator is interacted with a time-varying measure of world per barrel crude oil prices (in 2018 USD), which is normalized around the annual 1861 to 2000 mean. 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 covers various years from 1930 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 the shares of labor force in mining, agriculture, and manufacturing, log per worker annual mining output, log per worker annual mining wages, log per worker annual manufacturing wages, the unemployment rate, and median family annual income. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A12: Heterogeneous Effects: Oil Price Volatility Controlling for Oil Price

	Outcome: Membership in major Christian churches (% population)					
	(1)	(2)	(3)	(4)	(5)	(6)
Oil abundance	1.980*	4.393***	4.842***	6.512***	4.637***	4.408***
	(1.109)	(0.917)	(0.899)	(1.132)	(1.316)	(0.967)
Oil × 5 yr price s.d.	0.647***					
	(0.135)					
Oil × 10 yr price s.d.		0.183***				
		(0.045)				
Oil × 25 yr price s.d.			0.137***			
			(0.036)			
Oil × 5 yr log price s.d.				-0.129		
				(4.287)		
Oil × 10 yr log price s.d.					6.353*	
					(3.816)	
Oil × 25 yr log price s.d.						5.601***
						(1.687)
Observations	4574	4574	4574	4574	4574	4574
Counties	520	520	520	520	520	520
Adj. R^2	0.762	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. All specifications also interact the treatment variable with lagged real oil prices, normalized around the annual 1861 to 2000 mean. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A13: Heterogeneous Effects: Oil Dependence, Risk, and "Boomtowns"

	Outcome: Membership in major Christian churches (% pop.)					
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
Oil abundance	5.168***	7.782***	5.072***	10.487***	6.573***	9.502***
	(0.844)	(0.986)	(0.876)	(1.360)	(0.912)	(1.183)
Oil × Log pop density, 1900	-5.582***					
	(0.685)					
Oil × Urban in 1890		-8.104***				
		(1.755)				
Oil × Log mfg output pc, 1900			-1.886***			
			(0.301)			
Oil × Above-median mfg output, 1900				-8.747***		
				(1.583)		
Oil × % land in cotton, 1900					-0.447***	
					(0.115)	
Oil × Above-median cotton, 1900						-6.101***
						(1.613)
Observations	4566	4574	4226	4226	4526	4526
Counties	519	520	479	479	514	514
Adj. R^2	0.767	0.762	0.776	0.775	0.761	0.762
Outcome mean	32.87	32.88	32.27	32.27	32.91	32.91

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 county characteristics in 1890 and 1900, which for most counties is before major oil discoveries were made. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A14: Oil and Religion with Private Substitutes: Controlling for Urbanization

	Outcome: Membership in major Christian churches (% population)					
	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
Oil abundance	8.396***	8.585***	6.288***	6.464***	10.014***	9.991***
	(1.273)	(2.755)	(1.002)	(1.015)	(1.876)	(1.930)
Oil × Any savings & loans banks, 1950	-3.674**	-7.102*				
	(1.602)	(4.089)				
Oil × Any bank tellers, 1910			-4.096**	-4.141**		
			(1.836)	(1.947)		
Oil × Any insurance agents, 1910					-5.706***	-5.395***
					(1.927)	(1.979)
Observations	4574	3390	4529	4426	4529	4426
Counties	520	382	513	501	513	501
Adj. R^2	0.768	0.767	0.773	0.771	0.774	0.772
Outcome mean	32.88	31.83	32.73	32.71	32.73	32.71
Drops counties treated ≤ 1950?	No	Yes				
Drops counties treated ≤ 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. We also control for the log population density of each county in 1950 or 1910. As these are time-invariant, we interact them with sample year fixed effects. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A15: Oil and Demand for Private Substitutes for Consumption Smoothing

	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
	Log savings per capita		Log time deposits per capita		Log insurance agents per capita	
Oil abundance	0.501	0.674***	1.091***	0.623***	0.017***	0.019***
	(0.397)	(0.235)	(0.307)	(0.223)	(0.004)	(0.005)
Observations	4160	2405	4680	4394	2080	1721
Counties	520	435	520	515	520	483
Adj. R^2	0.743	0.903	0.849	0.912	0.673	0.674
Outcome mean	4.576	7.866	6.671	7.103	0.047	0.056
Outcome >0	No	Yes	No	Yes	No	Yes

Note: Estimates are from difference-in-differences regressions of county-level measures for insurance and banking 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. 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 savings per capita, log time deposits in banks per capita, and the log number of insurance agents per capita. Savings capital data are available for all counties for eight years between 1947 and 1982; time deposits data are available for all counties for nine years between 1940 and 1980; and decadal insurance agent data are available from the full count Censuses for all counties from 1910 to 1940. Columns (b) restrict the sample to those counties that had strictly positive outcome values to account for significant truncation at zero for the variables. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A16: Oil and Private Substitutes: Controlling for Urbanization

	(1a)	(1b)	(2a)	(2b)	(3a)	(3b)
	Log savings per capita		Log time deposits per capita		Log insurance agents per capita	
Oil abundance	0.454	0.680***	1.105***	0.648***	0.006	0.010**
	(0.423)	(0.232)	(0.307)	(0.216)	(0.004)	(0.004)
Observations	4160	2405	4680	4394	2080	1721
Counties	520	435	520	515	520	483
Adj. R^2	0.745	0.903	0.850	0.915	0.692	0.691
Outcome mean	4.576	7.866	6.671	7.103	0.047	0.056
Outcome >0	No	Yes	No	Yes	No	Yes

Note: Estimates are from difference-in-differences regressions of county-level measures for insurance and banking 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. 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 savings per capita, log time deposits in banks per capita, and the log number of insurance agents per capita. Savings capital data are available for all counties for eight years between 1947 and 1982; time deposits data are available for all counties for nine years between 1940 and 1980; and decadal insurance agent data are available from the full count Censuses for all counties from 1910 to 1940. Columns (b) restrict the sample to those counties that had strictly positive outcome values to account for significant truncation at zero for the variables. All regressions include each county's log population density. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A17: Religion, Oil Shocks, and Local Labor Composition: Robustness to Sample

	(1) % Unemployed	(2) Log pop density	(3) % Mining workers	(4) % Agricultural workers	(5) % Mfg workers
Oil × Oil price increase	-0.033***	0.001***	0.001	0.099***	-0.030***
	(0.003)	(0.000)	(0.005)	(0.012)	(0.007)
Oil × Oil price increase					
× Above-median Christian, 1936	0.013***	-0.001	0.007	-0.074***	0.025**
	(0.004)	(0.001)	(0.007)	(0.016)	(0.010)
Observations	3120	3120	3116	3118	3120
Counties	520	520	520	520	520
Adj. R^2	0.554	0.944	0.817	0.811	0.759
Outcome mean	5.310	2.307	3.005	22.13	12.47
Drops counties treated $\geq$ 1936?	No	No	No	No	No

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. Nevertheless, in this version we do not 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 * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Table A18: Oil and Religion: Denominational Effects

	Herfindahl index	% Southern Baptist	% Roman Catholic	% United Methodist
	(1)	(2)	(3)	(4)
Oil abundance	-0.038***	2.884***	4.199***	-0.423
	(0.009)	(0.671)	(0.821)	(0.292)
Observations	4574	4574	4598	4598
Counties	520	520	520	520
Adj. R^2	0.717	0.766	0.821	0.613
Outcome mean	0.558	10.40	9.930	7.546

Note: Estimates are from difference-in-differences regressions of measurements of membership in 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. In column (1), all denominations as well as the remainder are used to construct a county-level Herfindahl index of denominational concentration. Columns (2-4) measure county-level membership of Southern Baptists, Roman Catholic, and United Methodist denominations (% population), respectively. 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. Standard errors are clustered at the county level. Significance levels are denoted by * $p < 0.10$, ** $p < 0.05$, *** $p < 0.01$.

Appendix Figures

Figure A1: Balancing Test of the Oil Treatment on Pre-Discovery Observables

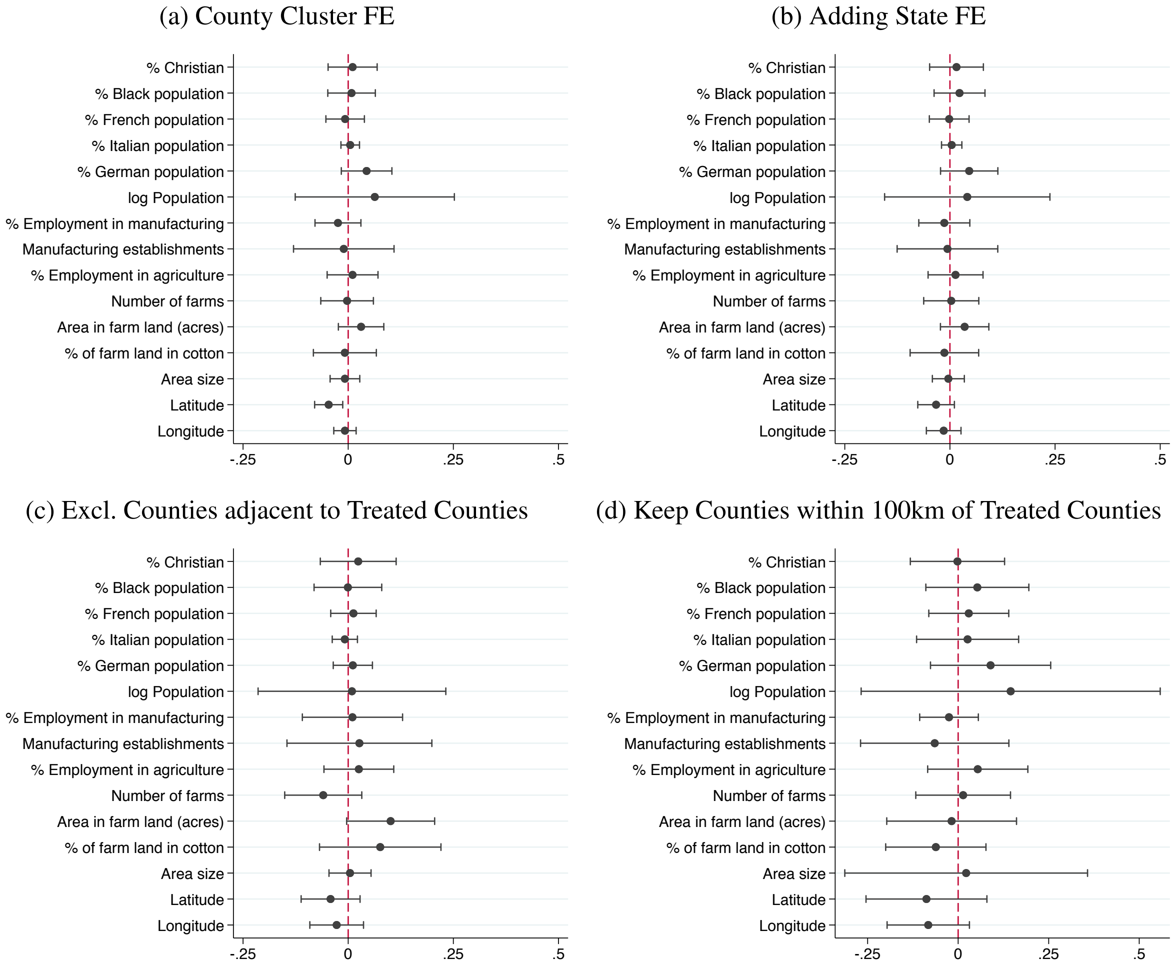

Note: Regressions of the oil treatment indicator on observable pre-oil discovery characteristics in the year prior to discovery. Due to the different timing of oil discoveries, we first define *clusters* of control counties and their nearest treated county. In each cluster, we consider all counties in the year before oil is discovered in the treated county of this cluster and regress the oil treatment indicator on observable county characteristics to test whether there are pre-determined variables that can predict discoveries. Since all counties are in the vicinity of an eventually treated county, they arguably had similar chances of discovering oil. All variables except latitude and longitude are standardized to have mean zero and variance one for comparability. The coefficients of these regressions with their 95% confidence intervals are plotted in each sub-graph which subsequently add different fixed effects, in (a and b), or exclude counties depending on proximity to the oil county, in (c and d). Major oilfields hold 100 million barrels of oil or more, which is also the definition of the oil treatment variable. Standard errors are clustered at the county level and the red dashed line marks zero.

Figure A2: Sensitivity of Oil Abundance and Religious Participation to Sample Changes

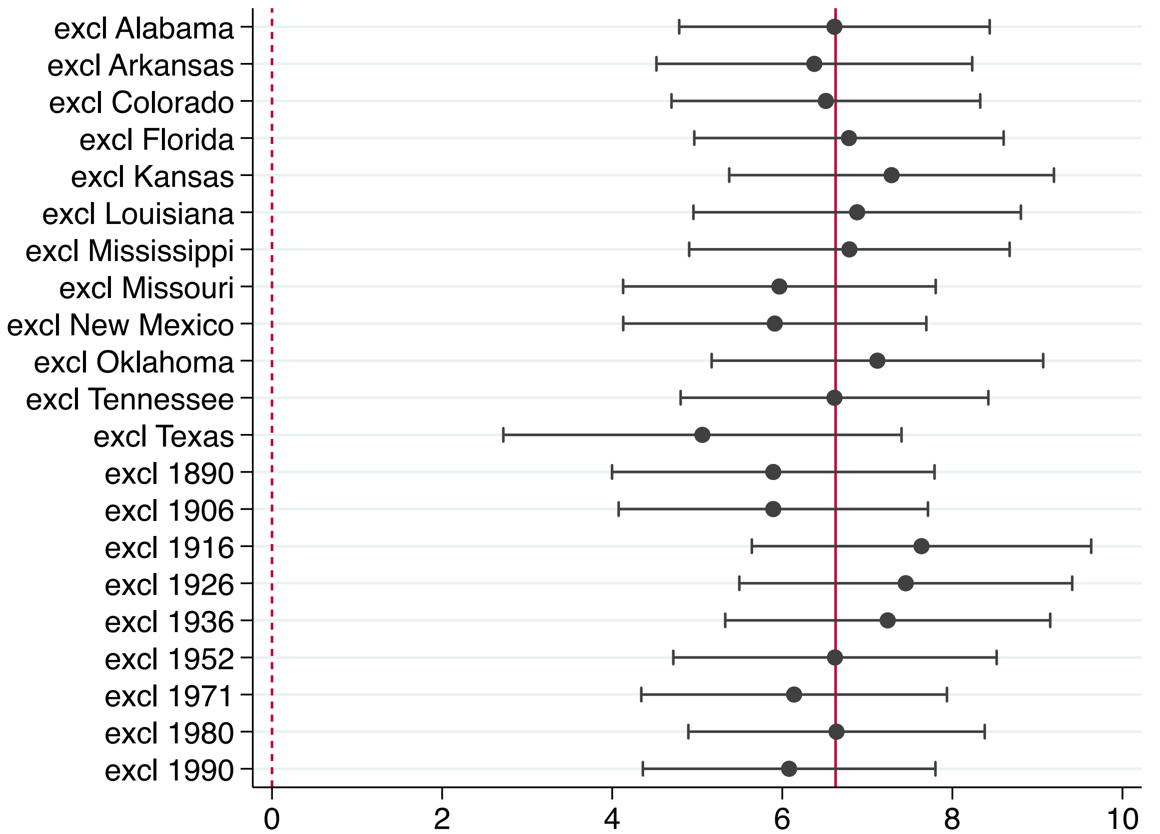

Note: Estimates are from difference-in-differences regressions of membership in 15 major, mainstream Christian denominations (% population) in county c in year t on an indicator called "oil abundance," which equals one for a county in years following a major oil discovery and is zero otherwise. Major 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. Each coefficient represents the oil abundance effect on membership among 15 major, mainstream Christian denominations, while excluding a certain state or sample year at a time. We exclude control counties that are adjacent to oil counties to avoid issues from spillover effects. The solid red line is the baseline effect, the dashed red line marks zero to show the distance of a specific coefficient from being a null effect. Standard errors are clustered at the county level and error bars show 95% confidence intervals.

Figure A3: Matched Sample on Pre-Oil Population Growth

(a) Unconditional Log Population Growth

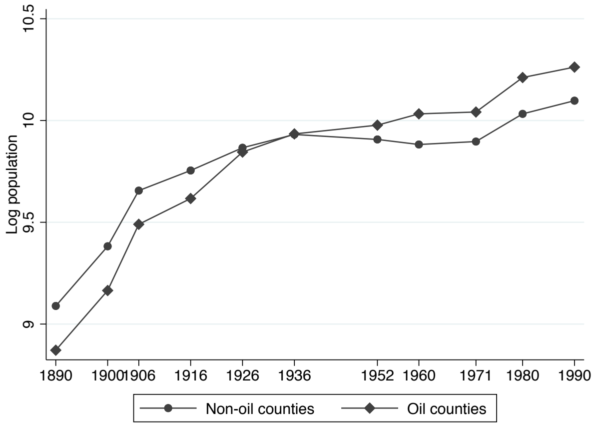

(b) Unconditional Log Population Density Growth

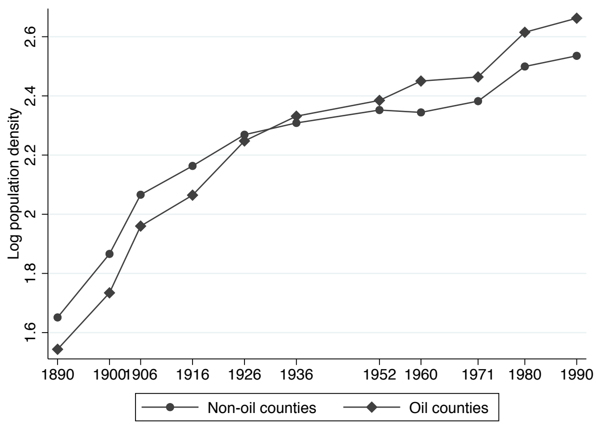

(c) Conditional Log Population Growth

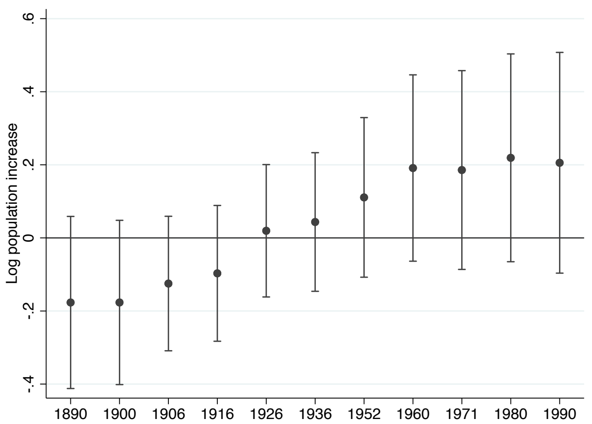

(d) Conditional Log Population Density Growth

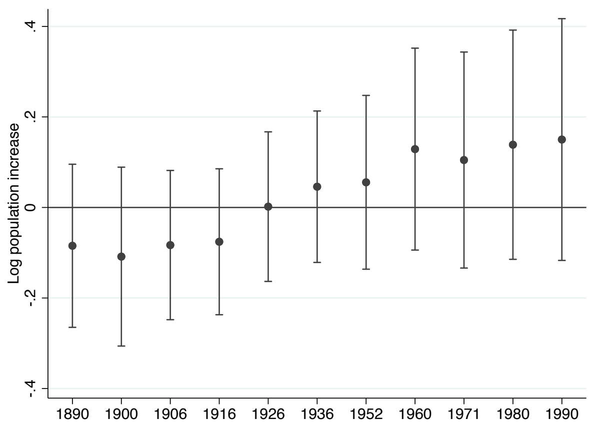

Note: Log population and log population density growth in oil and non-oil counties after matching on each variable for the years prior to the discovery of oil in the oil counties. Figures (a) and (b) show the unconditional evolution in the two groups for the matched sample over time. Figures (c) and (d) show the conditional evolution of these variables by regressing them on county and year fixed effects as well as the interaction of an indicator for whether a county ever had oil with the year fixed effects. The coefficients from this interaction are plotted in the two figures together with the 95% confidence interval which is represented by the error bars. Standard errors are clustered at the county level.

Figure A4: Matched Sample on Population Growth Over the Full Sample

(a) Unconditional Log Population Growth

(b) Unconditional Log Population Density Growth

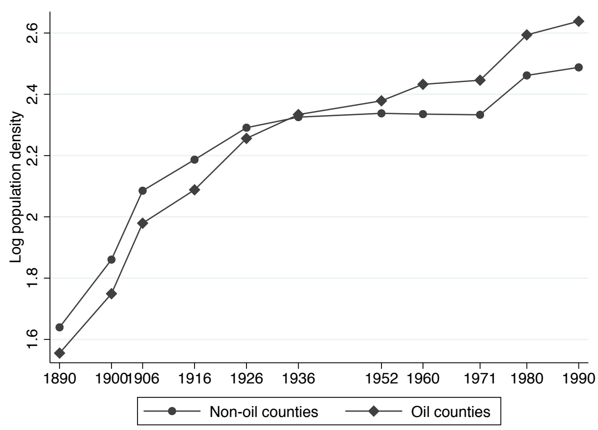

(c) Conditional Log Population Growth

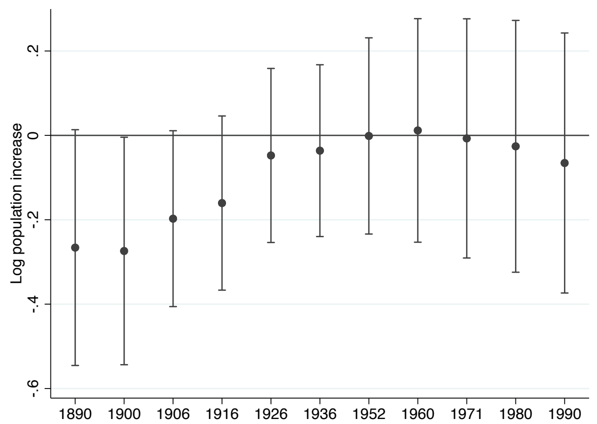

(d) Conditional Log Population Density Growth

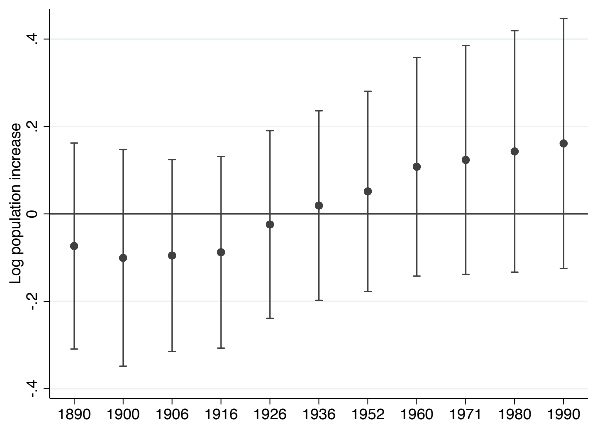

Note: Log population and log population density growth in oil and non-oil counties after matching on each variable over the entire sample period. Figures (a) and (b) show the unconditional evolution in the two groups for the matched sample over time. Figures (c) and (d) show the conditional evolution of these variables by regressing them on county and year fixed effects as well as the interaction of an indicator for whether a county ever had oil with the year fixed effects. The coefficients from this interaction are plotted in the two figures together with the 95% confidence interval which is represented by the error bars. Standard errors are clustered at the county level.

Figure A5: Spatial Distribution of Denominations in 1952

(a) All sample Christian churches (% population)

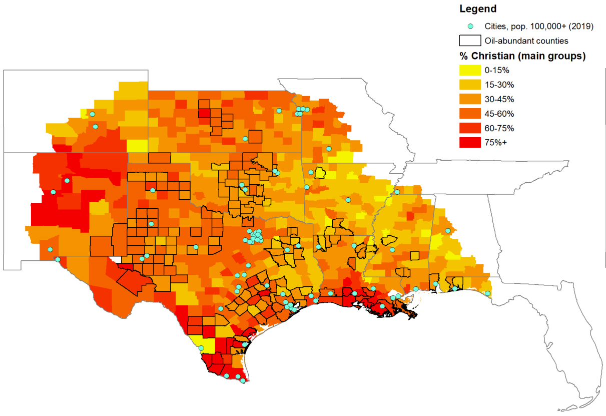

(b) United Methodist (% population)

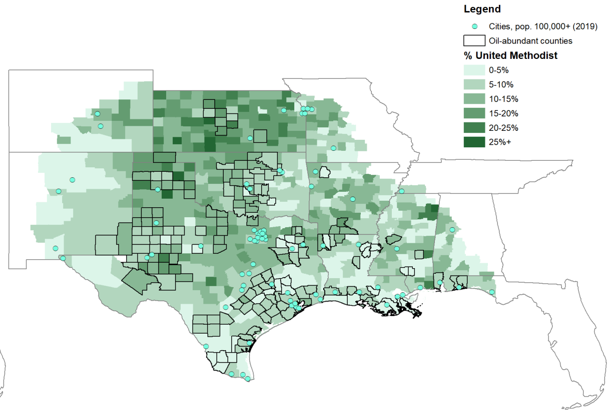

(c) Roman Catholic (% population)

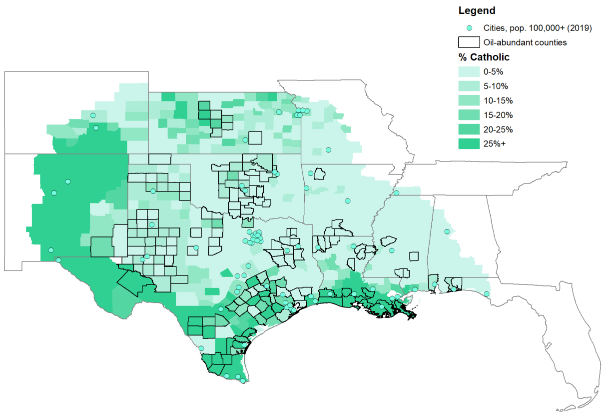

(d) Southern Baptist (% population)

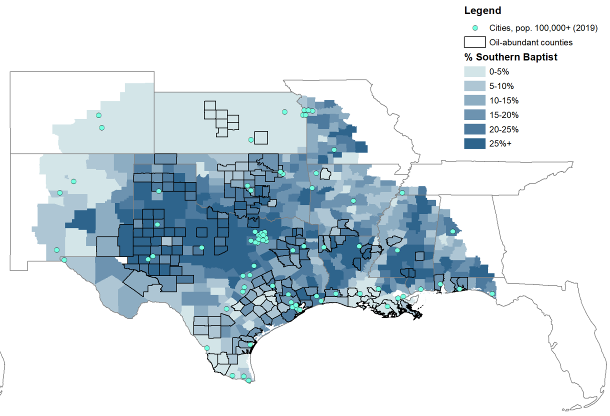

Note: Maps show the spatial distribution of different Christian denominations as a share of the total population in our sample counties, as reported in the 1952 religious census. Oil-abundant counties are outlined in black, while urban areas (cities with population >100,000 in 2019) are dotted in light 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).

Figure A6: Spatial Distribution of Denominations in 1990

(a) All sample Christian churches (% population)

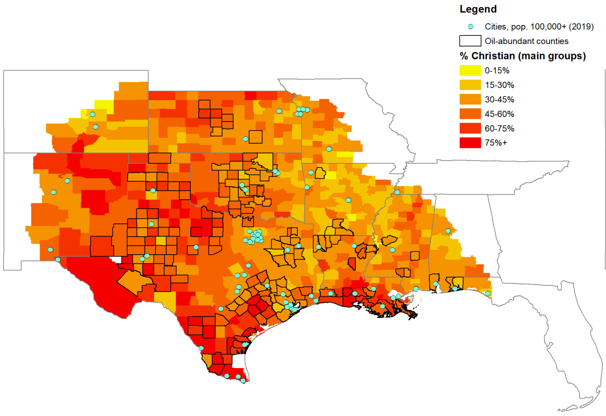

(b) United Methodist (% population)

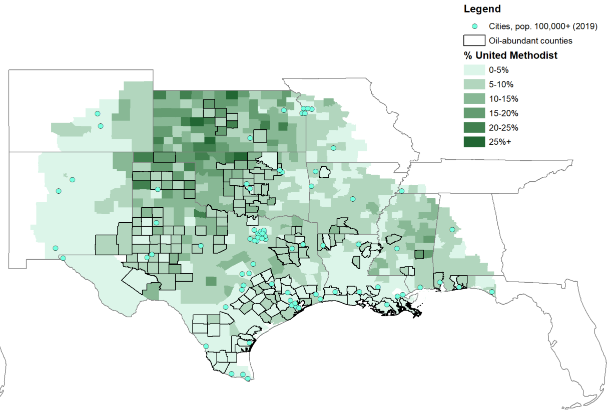

(c) Roman Catholic (% population)

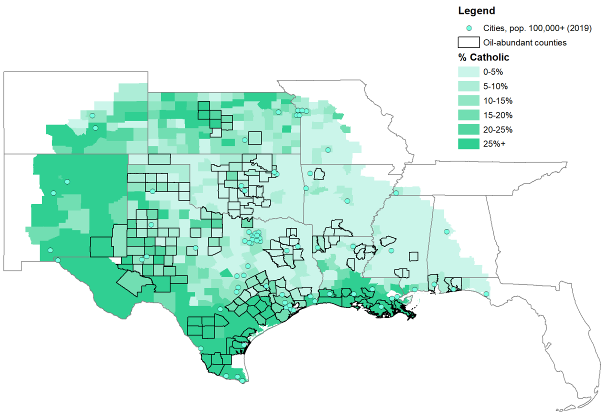

(d) Southern Baptist (% population)

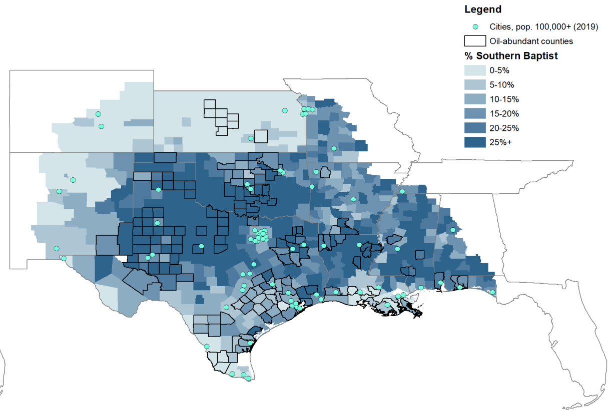

Note: Maps show the spatial distribution of different Christian denominations as a share of the total population in our sample counties, as reported in the 1990 religious census. Oil-abundant counties are outlined in black, while urban areas (cities with population >100,000 in 2019) are dotted in light 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).

B Theory Appendix

In this Appendix, we present a simple model to motivate our empirical analysis and illustrate how volatility in the return on one's endowment can lead to religious participation as a source of risk mitigation. In the model, as in the early oil South and in many undeveloped and developing settings today, there is a lack of strong formal insurance and lending institutions, for which religious communities may serve as a substitute. In the presence of uncertainty regarding the future return on one's endowment, this can generate incentives for workers to make religious investments, i.e. sacrifice some of their time and income to the church, which in turn may provide economic and other forms of support during hard times.

Endowment uncertainty in this setting is assumed to stem from a reliance on natural resources. Natural resource quantities are generally inelastic in price in the short-run, leading to large fluctuations in prices (van der Ploeg and Poelhekke, 2009; Ross, 2012). To the extent that labor demand in resource-abundant communities is derived from the market for such resources, relevant economic shocks can have a significant impact on real income in the short-run. If agents are risk averse, then such volatility may have important welfare implications for workers in those communities, upon which they will seek to improve intertemporally.

To show this, we adopt a two-period model with uncertainty in future endowment returns. In the first period, a risk-averse representative agent is endowed with some initial income y_0, which she can use for consumption c_0 or for religious investments r_1. In making such investments today, the agent provides the church with resources (e.g. money, manpower) needed for it to help her and others should hard economic times strike tomorrow.[1] Then, in the second period, the agent receives for consumption some additional income $y_1(s)$, the value of which depends on the state of the economy s, as well as support from the church *if and only if* that income is relatively low, the value of which depends on the agent's first period investment.

We define the agent's lifetime utility maximization problem as

$$\max_{c_0, c_1(s), r_1} \; u(c_0) + \beta \mathbb{E}[u(c_1(s))],$$

where $u(\cdot)$ is twice continuously differentiable with $u'(c_t) > 0$ and $u''(c_t) < 0$, and $\beta > 0$ is an intertemporal discount factor. Consumption in the initial period ($t = 0$) depends on initial income as well as the size of religious investments made: $c_0 \leq y_0 - r_1$. Consumption in period $t = 1$ depends on the state of the economy, with church support supplementing the agent's endowment in "bad" states, i.e. when $y_1(s)$ is relatively low:

$$c_1(s) \leq y_1(s) + \mathbb{1}\{s \in \mathit{bad}\} \times A r_1,$$

where $A > 1$ is a multiplier parameter, representing the supplemental effects of religious in-

[1] In an extended model with more time periods, such sacrifices would occur in an ongoing manner, not just in the period prior, while shocks would be modeled as idiosyncratic across several communities or markets.

vestments from across different markets with the *same* religious institutions but *idiosyncratic* state realizations.[2]

Maximization yields:

$$u'(c_0^*) = \beta APr[s \in bad]\mathbb{E}[u'(c_1^*(s)|s \in bad], \tag{1}$$

with budget constraints binding in equilibrium due to the strictly increasing nature of utility. In order to see how consumption and in turn religious investments respond to endowment volatility, one must place some restrictions on the distribution from which $y_1(s)$ is drawn. We now evaluate the problem using a normal probability distribution with quadratic utility functions.[3]

Let $y_1 = \theta + \varepsilon$, where θ is some measure of long-run real income and $\varepsilon \sim N(0, \sigma^2)$ is a normally-distributed random variable, scaled by $\sigma > 0$. Similar to before, the church provides support if and only if real income is below average in $t = 1$, i.e. $\varepsilon \leq 0$. We then adopt a quadratic utility function, which is a common choice in settings such as this, as it satisfies risk aversion while enabling one to pass the expectations operator through $u'(\cdot)$ to evaluate $\mathbb{E}[\varepsilon|\varepsilon \leq 0]$. Letting $u(c_t) = c_t - \frac{\alpha}{2}c_t^2$, equation (3) becomes

$$1 - \alpha(y_0 - r_1^*) = \frac{\beta}{2}A(1 - \alpha(\theta + \mathbb{E}[\varepsilon|\varepsilon \leq 0] + Ar_1^*)). \tag{2}$$

Evaluating the expectation yields equilibrium religious investments of r_1^*, which are positive if and only if shocks are sufficiently large relative to real income, θ. That is, there exists some threshold value of $\sigma \equiv \tilde{\sigma}$, only above which religious investments are made in equilibrium. This is because the agent only receives church support in return for her investments during below-average income periods, such that there must be a lot at stake, relatively speaking, for her to invest. Then, conditional upon $\sigma > \tilde{\sigma}$, the size of religious investments is always increasing in σ. Altogether, these results can be summarized with the following proposition:

Proposition 1. *There exists an equilibrium in which agents forgo some consumption in favor of religious investments if and only if:*

(i) The relative impact of negative economic shocks on real income is sufficiently large: $r_1^* > 0$ *if and only if* $\sigma > \tilde{\sigma}$, *with* $\frac{\partial \tilde{\sigma}}{\partial \theta} > 0$*; where*

(ii) The size of religious investments is further increasing in endowment volatility: $\frac{\partial r_1^*}{\partial \sigma} > 0$ *when* $r_1^* > 0$.

In other words, if typical shocks to the return on one's natural resource endowment are economically salient, then an agent will engage in religious participation, with greater volatility

[2]For example, if shocks were independently and identically distributed across a large number of identical regions, with half in bad states, and all religious investments went to providing church support, then $A = 2$.

[3]The same insights can be yielded using general utility functions by assuming a binary probability distribution. For that analysis, please see the alternative model below.

in the return on one's endowment increasing religious investments thereafter. Given this, we expect to find that counties with greater dependence on oil will experience increased religious participation (at least for these major denominations that span beyond just oil regions), with such demand further increasing with greater oil price volatility.

B.1 Proof of Proposition 1

Proof. The Lagrangian is:

$$L = u(c_0)+\lambda_1[y_0-r_1-c_0]+\beta Pr[\varepsilon \leq 0](\mathbb{E}[u(c_1)|\varepsilon \leq 0]+\lambda_2[\mathbb{E}[y_1|\varepsilon \leq 0]+Ar_1-c_1(\varepsilon \leq 0)]) + \beta Pr[\varepsilon > 0](\mathbb{E}[u(c_1)|\varepsilon > 0] + \lambda_3[\mathbb{E}[y_1|\varepsilon > 0] - c_1(\varepsilon > 0)]). \quad (3)$$

Maximizing this yields four first order conditions:

$$r_1^* : -\lambda_1^* + \lambda_2^*\beta Pr[\varepsilon \leq 0]A = 0, \quad (4)$$

$$c_0^* : u'(c_0^*) - \lambda_1^* = 0, \quad (5)$$

$$c_1^*(\varepsilon \leq 0)) : \beta Pr[\varepsilon \leq 0]\mathbb{E}[u'(c_1^*)|\varepsilon \leq 0] - \lambda_2^*\beta Pr[\varepsilon \leq 0] = 0, \quad (6)$$

$$c_1^*(\varepsilon < 0)) : \beta Pr[\varepsilon > 0]\mathbb{E}[u'(c_1^*)|\varepsilon > 0] - \lambda_3^*\beta Pr[\varepsilon > 0] = 0, \quad (7)$$

Combining (7) and (8) using (6) and using the fact that $u'(c_t) = 1 - \alpha c_t$ under quadratic utility yields:

$$1 - \alpha c_0^* = \beta A Pr[\varepsilon \leq 0](1 - \alpha\mathbb{E}[c_1^*|\varepsilon \leq 0]).$$

Then, using the fact that $Pr[\varepsilon \leq 0] = \frac{1}{2}$ if $\varepsilon \sim N(0, \sigma^2)$, $c_0^* = y_0 - r_1^*$, and $\mathbb{E}[c_1^*|\varepsilon \leq 0] = \theta + \mathbb{E}[\varepsilon|\varepsilon \leq 0] + Ar_1^*$, it is easy to see that this is equivalent to equation (4).

From here, solving for r_1^* requires that one evaluate $\mathbb{E}[\varepsilon|\varepsilon \leq 0]$. $\mathbb{E}[\varepsilon|\varepsilon \leq 0] = \sigma\mathbb{E}[\frac{\varepsilon}{\sigma}|\frac{\varepsilon}{\sigma} \leq \frac{0}{\sigma}]$, where $\frac{\varepsilon}{\sigma} \sim N(0, 1)$ is standard normal with a density function of $\phi(\epsilon)$ and a cumulative distribution function of $\Phi(\epsilon)$. Hence, this conditional expectation can be evaluated as follows:

$$\mathbb{E}[\varepsilon|\varepsilon \leq 0] = \sigma\mathbb{E}[\frac{\varepsilon}{\sigma}|\frac{\varepsilon}{\sigma} \leq 0] = \sigma\frac{\int_{-\infty}^{0} \epsilon\phi(\epsilon)d\epsilon}{Pr[\frac{\varepsilon}{\sigma} \leq 0]}.$$

For the standard normal distribution, $\phi'(\epsilon) = -\phi(\epsilon)\epsilon$, so we can rewrite this as

$$\mathbb{E}[\varepsilon|\varepsilon \leq 0] = -2\sigma\int_{-\infty}^{0} \phi'(\epsilon)d\epsilon = -2\sigma\phi(0) = -2\sigma(2\pi)^{-\frac{1}{2}},$$

which upon plugging into (4) simplifies to

$$r_1^* = \frac{\alpha[2y_0 - \beta A(\theta - 2\sigma(2\pi)^{-\frac{1}{2}})] - 2 + \beta A}{\alpha(2 + \beta A^2)}.$$

It is straightforward to show from here that is strictly increasingly in σ but strictly positive for only some values of σ, namely:

$$r_1^* > 0 \Leftrightarrow \sigma > \frac{2(1-\alpha y_0) - \beta A(1-\alpha\theta)}{2\alpha\beta A(2\pi)^{-\frac{1}{2}}} \equiv \tilde{\sigma},$$

which by inspection is increasing in θ. □

B.2 Alternative Model: General Utility With Binary Income States

Let $y_1(s) \in \{y_1(b), y_1(g)\}$ be bad and good state endowment incomes, respectively, where $y_1(g) = y_1(b) + \eta > y_1(b) > 0$ and where η represents the income differential between good and bad states. As before, the church provides support if and only if $s = b$.

In this setting, the Euler equation becomes

$$u'(y_0 - r_1^*) = \beta APr[s = b]u'(y_1(b) + Ar_1^*).$$

Setting $y_1(b) = y_1(g) - \eta$ and defining $Pr[s = b] = p(b)$, we can derive comparative statics for r_1^* by implicitly differentiating this equation with respect to η, which yields:

$$\frac{\partial r_1^*}{\partial \eta} = \frac{\beta Ap(b)u''(y_1(g) - \eta + Ar_1^*)}{u''(y_0 - r_1^*) + \beta A^2 p(b)u''(y_1(g) - \eta + Ar_1^*)} > 0.$$

As η indicates both the relative impact of negative shocks on real income as well as overall income dispersion, this comparative static is more or less synonymous with Proposition 1 without restricting the utility function, albeit at the cost of using a much simpler distribution of income states.

C Data Appendix

C.1 Religious Data

The data on religious bodies and church memberships were obtained primarily from the Association of Religion Data Archives (ARDA).[4] This includes data from Statistics of Churches in the United States for 1890, the U.S. Census of Religious Bodies for 1906, 1916, 1926 and 1936, Churches and Church Memberships in the United States for 1952, 1971, 1980, and 1990. The final dataset uses the Censuses of Religious Bodies included in Haines (2010), for reasons of data cleanliness and formatting. As far as possible,[5] we harmonized church memberships and denominations across years to generate a stable measure of church membership. For the final sample, this includes membership information by county for the Roman Catholics Church, Latter Day Saints, several mainline Protestant groups (United Methodist Church, Evangelical Lutheran Church in America, American (Northern) Baptist Church, Episcopal Church, Presbyterian Church, United Church of Christ, Disciples of Christ, Reformed Church of America), and several evangelical Protestant groups (Southern Baptist Church, Wisconsin Evangelical Lutheran Synod, Lutheran Church–Missouri Synod, Christ Reformed Church, and Seventh Day Adventists), which we aggregate to construct a measure of (mainstream, predominantly-white) Christian participation in sample counties. Other groups, including various Pentecostal churches, Black Methodist churches, and Black Baptist churches, are not included due to missing data for several religious censuses. These were matched to county boundaries for religious census years. Between-census county boundaries (e.g. for 1906) were determined using the Atlas of Historical County Boundaries to modify the Tiger/Line boundaries from the U.S. Census Bureau.[6] Controls for nearest-census year variables were then combined by harmonizing boundaries from the nearest census year from the U.S. Census Bureau, following the procedure in Hornbeck (2010) and Ferrara, Testa and Zhou (2021). To construct the panel dataset, all boundaries were harmonized again to 2000 boundaries, using the same approach, to create a unified panel.

Some additional measures are taken to construct the full sample:

- For 1906, major Baptist groups are combined and counts must be imputed from their 1890 or 1916 relative membership counts. This is the subject of Table A7.
- Through 1916, Wisconsin and Missouri Lutheran members are counted as part of the Evangelical Lutheran Synodical Conference of North America, which at times included some other minor groups. We treat these groups as synonymous.
- Some groups were notably undercounted in 1936, namely Southern Baptists and Southern Methodists (Ager, Hansen and Lonstrup, 2016). Results are not sensitive to dropping

[4]These can be accessed online at http://www.thearda.com/Archive/ChCounty.asp.

[5]A frequent issue with the church census and religious data is that several churches and denominations merge, split, or change name over time.

[6]See https://publications.newberry.org/ahcbp/ (last accessed on Dec. 1, 2019).

1936, as shown in Figure A2.

- From 1971 forward, Catholics are only reported by number of adherents, not members, the former being a superset of the latter. We thus adopt this measure in place of membership for this denomination for these years.

- Measurement error causes a small number of county-years' Christian membership to exceed 100% of the population upon aggregation. These counties are censored at 100%. To account for this, we create a dummy variable, given a value of 1 for these counties, for which we control in all specifications. Results are not sensitive to instead dropping these county-years.

C.2 Oil Data

Data for major U.S. oilfields come from the Oil and Gas Journal Data Book (2000), which lists the universe of oilfields in the United States with 100 million barrels (bbl) of oil or more, both on land and offshore, their locations by state, and their *overall* discovery years. We link on-land major oilfields with data for all county-oilfields from the Oil and Gas Field Code Master List (U.S. Department of Energy, 2004). This lists all oil and gas fields in the United States, their county(s), each county-field's discovery year, and its composition (for example, oil, natural gas, both, etc.). Only county-fields with oil are kept. Then, as treatment is at the county level, we compile a list of all major oilfields for treated counties in the sample, along with their county-specific discovery year, and then assign to that county the earliest of those years as its treatment year.

Data for crude oil prices are based on U.S. average spot prices for 1861 to 1944, Arabian Light prices as posted at Ras Tanura in Saudi Arabia for 1945 to 1983, and Brent dated prices after 1983. Oil price data was compiled by BP and collected from Quandl.[7]

C.3 County Level Data

The two main data sources we use to construct economic information at the county level for our estimation sample are the U.S. Census of Population and Housing digitized by Haines (2010) and the Census of Agriculture which was collected by Haines, Fishback and Rhode (2018). From the population Census we harmonized variables over time. The variables we harmonized from the Census include the total population, number of urban population in cities of at least 25,000 people, county area, the percent of Black, native-born, and foreign-born from various countries of origin, and the number of firms, employment, wages, and output in manufacturing. From the Census of Agriculture we harmonized the number of farms, farm values, acres of land in farms, output, and the value of machinery and implements. Nominal values were deflated

[7]These data can be accessed from https://www.quandl.com/data/BP/CRUDE_OIL_PRICES, from where we retrieved our data on July 27, 2020.

to 2018 U.S. dollars using the CPI provided by the Minneapolis Fed.[8] A few county-level variables were taken from secondary sources. Additional decadal data on sector employment for 1960-90 from Michaels (2011), as well as an indicator he constructed for whether a county had a least some of its population in an urban area of 25,000+ in 1890, were also merged into our dataset, as were manufacturing data for 1910 from Matheis (2016). To create consistent county boundaries over time, we follow the approach by Hornbeck (2010). The year to which these boundaries are harmonized is 2000.

Information on the number of bank tellers (305), insurance agents (450), actuaries (83), clergy (9), religious workers (78), and social workers (79) per county was obtained from the 1910, 1920, 1930, and 1940 full count Census files provided by Ruggles, Flood, Goeken, Grover, Meyer, Pacas and Sobek (2018). Numbers in parentheses indicate the occupation code contained in the `occ1950` variable, which uses the 1950 U.S. Census Bureau occupational definitions to classify workers into these occupational groups. In each Census year, we kept individuals who were between 15 and 75 years of age, active in the labor force, and who did not report to currently attend school.

Stata-level data for unemployment insurance and workers compensation benefits from 1940 through 1990 are also merged with county-level data. These measures are derived in Fishback (2020) and downloaded from that paper's data archives.

C.4 Generating Matched Samples Based on Population Growth and Migration

This section describes the matching procedure based on the full population growth trajectory of oil and non-oil counties, as well as for the pre-oil discovery population trajectories that are used in Appendix Tables A2 and A3, and Appendix Figures A3 and A3. One potential concern we sought to address by controlling for the log population size in columns (3) and (4) of table 1 is that oil counties may have been substantially different from the comparison counties in terms of their local development. For instance, new oil discoveries were typically made in sparsely populated areas. This might have had several implications not only for local development but also for the incentives of religious organizations to establish footholds in such areas.

To further address this potential issue, we attempt to make oil and non-oil counties more comparable by matching them on their population growth over time. We employ two different strategies. First, we match counties based on their log population growth and log population density growth in the years before the oil discovery using propensity score matching. For each oil county we determine the year in which oil is found and then match potential comparison counties on our population measures for those pre-oil discovery years. The idea is to limit the possibility that post-oil outcomes are driven by differences in pre-oil population growth which would imply a violation of the parallel trends assumption in our difference-in-differences setting. Second, we also match oil and non-oil counties on the three population measures over

[8] The CPI series is available online and can be accessed at https://www.minneapolisfed.org/about-us/monetary-policy/inflation-calculator/consumer-price-index-1800-.

the entire sample period. While population growth might be a direct result of oil discoveries, this strategy allows us to verify that post-oil discovery increases in religion are not merely a result of population increases in the respective counties. Since we are then comparing counties that have the same population growth over all time periods, any changes in religion cannot be driven by differential population growth across oil and non-oil counties.

We show the unconditional and conditional evolution of the population and log population size across oil and non-oil counties from both matching exercises in Figures A3 and A4, respectively. The conditional plots were generated by regressing the corresponding population measure on year fixed effects as well as their interaction with an indicator for oil counties. The latter takes the value of one if oil is discovered at any point in the county over the sample period and is zero otherwise. As it turns out, the difference between matching on population variables in pre-oil discovery years compared to matching along the whole population growth trajectory over the entire sample makes little difference. This implies that population growth is relatively stable and is not strongly affected by oil discoveries themselves. While oil counties tend to start out with both lower log population and log population density in the earlier years and growing over time, the matched sample from our propensity score matching exercise is good enough such that those differences are relatively small and never significant.

We then use the matched samples to re-estimate equation (1). The results are reported in Table A2 using the sample matched on pre-oil discovery population outcomes, and in Table A3 using the matched sample on population outcomes along the entire sample period. When we condition on matched pair fixed effects, we also estimate effects controlling for geographic region dummies,[9] as not all matched oil and non-oil counties are in comparable geographic areas. This is ordinarily taken into account by county fixed effects. Because we are interested in the effect of population growth on our estimates, we partial out such geographic differences. Both tables show a positive and significant effect of oil abundance on Christian membership (% population) with effect sizes that are generally in the range of our main results in Table 1. We can therefore conclude that the treatment effect is not driven merely by differential effects of population growth between oil and non-oil counties.

[9]We define five geographic regions in our sample at the state-level: (i) Great Plains (Oklahoma and Kansas), (ii) Ozarks (Arkansas and Missouri), (iii) Deep South (Louisiana, Mississippi, Southwest Alabama and the Florida panhandle), (iv) the Southwest (New Mexico and Colorado), and (v) Texas.

References

Ager, Philipp, 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.

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.

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.

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

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

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.

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.

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

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