NBER WORKING PAPER SERIES

RELIGIOUS WORSHIP ATTENDANCE IN AMERICA:
EVIDENCE FROM CELLPHONE DATA

Devin G. Pope

Working Paper 32334
http://www.nber.org/papers/w32334

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
April 2024

I am grateful to Carla Colina, Carolina Cannon, and Rory Lawson for incredible research assistance on this project. I also thank Josh Davis, Josh Dean, Stefano DellaVigna, Kareem Haggag, Alex Imas, Zack Nelson, David Ridge, Scott Shurtliff, Avner Strulov-Shlain, Justin Sydnor, and seminar participants at various universities for helpful comments and suggestions. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

Religious Worship Attendance in America: Evidence from Cellphone Data
Devin G. Pope
NBER Working Paper No. 32334
April 2024
JEL No. Z12

ABSTRACT

Religious worship is integral to the lives of millions of Americans. In this paper, I provide a descriptive analysis of religious worship attendance using geodata from smartphones for over 2 million Americans in 2019. I establish several key findings. First, 73% of people step into a religious place of worship at least once during the year on the primary day of worship (e.g. Sundays for most Christian churches). However, only 5% of Americans attend services "weekly", far fewer than the ~22% who report to do so in surveys. The number of occasional vs. frequent attenders varies substantially by religion. I estimate that approximately 45M Americans attend worship services in a typical week of the year, but with large changes around Holidays (e.g. Easter). I document how start times, duration of attendance, and average household income all differ meaningfully across religious traditions. The intensity of religious observance correlates with a host of other activities. For example, relative to non-attenders and infrequent attenders, frequent religious attenders are less likely to go to strip clubs, liquor stores, and casinos. While cell phone data has limitations, this paper provides a unique way of understanding worship attendance and its correlates.

Devin G. Pope
Booth School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637
and NBER
devin.pope@chicagobooth.edu

Religion is an important aspect of life for a large number of people in the US and around the world. For example, in a recent poll, 70% of Americans said that religion is fairly or very important in their lives (GALLUP 2022). Despite the centrality of religion in the lives of so many people, religion is often mentioned as an understudied topic in economics (Iannaccone 1998; Iyer 2016; Hungerman 2020). In part, this is due to limited data on religious observance relative to many other aspects of life (consumption, labor force activities, etc.). With a few exceptions, what we know about religious observance is the result of surveys.[1]

Survey data on religious worship, however, might be biased in several ways. For example, survey responses have been shown to be unreliable in certain contexts like accurate reports of voting behavior (Anderson and Silver 1986; Bernstein, Chadha, and Montjoy 2001; Ansolabehere and Hersh 2012). Religious attendance may be similarly overstated due to misremembering or social desirability bias (and activities like strip club attendance may be understated for similar reasons).[2] In addition, survey data is unable to capture other aspects of religious attendance. For example, while surveys can ask participants about their income, it is difficult to gather information about the income of fellow congregants (because survey takers don't typically know the income of the people they worship with). Thus, survey data is unable to identify the level of economic diversity that exists within congregations across America – a potentially important factor for social mobility (Chetty et al. 2022). Lastly, the sample sizes of surveys are often quite small. This leads to statistical power issues, especially when discussing religious groups with a small number of affiliates (Jews, Buddhists, Hindus, etc. in an American context).

In this paper, I avoid many of the pitfalls associated with survey data and provide a descriptive analysis of religious worship attendance in America. The results make use of location data from a company that aggregates de-identified, geospatial data points for millions of U.S. smartphones. I restrict the sample to ~2.1 million cellphones that generate consistent location data over approximately a one-year period leading up to the COVID pandemic (April 2019 – February 2020). Is this 2.1M sample representative of the U.S. as a whole? Surely it is an imperfect representation. For example, it does not represent small children and other adults who do not use smartphones or whose phones do not generate consistent geodata due to a) turning off their phones for long periods of time or failing to pay their cellphone bill, b) not allowing location-tracking services on their

[1] There are a few notable exceptions such as Dube, Blumenstock, Callen (2022).

[2] Some studies have directly provided evidence of over reporting in a religious context (e.g. Hadaway, Marler, and Chaves 1998; Brenner 2011; Brenner and DeLamater 2017)

phone, c) regularly switching cellphones, or d) leaving the country for extended periods of time. However, I provide a host of checks that suggest that the sample is reasonably representative of the US. For example, it is representative by location (e.g. there are more cellphones in California than in Delaware), by population density, and by income (although the sample skews a bit towards being wealthier as one might expect given the requirement of consistent smartphone ownership). Further, I show that my cellphone sample predicts visit counts for locations/events that have a known number of visitors with reasonable accuracy. For example, I am able to accurately predict the number of people at a baseball game or the number of visitors each year to places like Home Depot, McDonalds, or Six Flags.[3] Thus, while the sample is clearly imperfect, it is arguably as representative as surveys that also struggle to reach small children, people who are not willing to take surveys, etc.

Treating the sample as if it is representative, I push forward and identify all visits in the data to a religious location (church, temple, mosque, synagogue, etc.). I define a religious worship visit to be a visit that occurred on the typical day of worship for each religion. For example, a visit to a Baptist church on a Sunday is defined as a religious worship visit whereas a visit to a Seventh-day Adventist church on a Saturday is defined as a religious worship visit. This restriction helps eliminate visits to a religious location for other reasons (choir concerts, preschools, etc.). I then assign each individual (i.e. cellphone) in the data as either unaffiliated (if they do not have a religious worship visit in the dataset) or I assign them to a particular religion (the religious location they frequented the most).

I produce eight key findings based on these religious worship visits. The first finding relates to overall religious worship attendance and frequency of attendance. 73% of my sample attended a religious worship service at least once during the ~1 year sample. This is higher than the 46% of Americans who report in surveys to go to a religious service at least once per year. This number may be higher due to individuals attending funerals or other events that the survey's explicitly ask responders to exclude. However, I find that only 21% of people in the sample attend worship services at least monthly (they attended at least 11 times throughout the 47 weeks of the data sample). This is lower than the 30% of people who claim in surveys to attend at least monthly. Finally, my analysis finds that only 5% of people are weekly-attending worshippers (I generously define weekly attendance as attending 36 weeks throughout the 47 weeks of data). This is much lower than the 22% of people who claim to be weekly attenders in the survey data.

[3] Not all data checks are perfect. For example, I undercount the number of people who go to an AMC theater or attend NBA basketball games and provide a discussion of these mispredictions.

The second key finding is that the frequency of religious attendance varies significantly by religion. Members of some religions, Latter-day Saints (Mormons) and Jehovah's Witnesses, have a relatively high fraction of members who are weekly attenders while members of other religions, Catholics and Jews, have a relatively low fraction of members who are weekly attenders. For example, there are approximately 14 times more Americans identifying as Catholics than Latter-day Saints, but slightly more Latter-day Saints who are weekly attenders than Catholics.[4] Overall, the vast majority (85%) of weekly-attending worshippers are Protestants.

Cellphone data allow for observing worship attendance on any given week. The third finding that I report is that approximately 45M Americans attend a worship service during a typical week of the year. There is limited week-to-week variation/seasonality in attendance with holidays being the major exceptions. For example, Easter Sunday and Christmas have nearly a 50% higher religious attendance than a typical week of the year. Similar spikes in attendance around holidays are found in the data for non-Christian religions (e.g. high Muslim worship during Ramadan and Buddhist attendance during Mahayana New Year).

The fourth finding uses the detailed nature of the cellphone data to observe the exact arrival and departure time for each congregant. I produce several figures that illustrate the variance in religious worship services that exists across religions. For example, the data show extreme consistency/uniformity in some religions both in terms of start times and durations (Muslims, Latter-day Saints, and Jehovah's Witnesses) while other religions are much less uniform (Buddhists and Hindus). Overall, the duration of religious worship visits ranges from an average of approximately 51 minutes (Muslims) to 115 minutes (Orthodox Christians, Latter-day Saints, and Jehovah's Witnesses).

The census tract of residence for each cellphone in the database can be determined with location patterns and I use this to assign median household income to each cellphone. The fifth finding in this paper shows the income distribution for attenders of each religion. Overall, religious individuals have very similar income to non-religious individuals ($79k vs. $80k). However, individuals that attend weekly are slightly less wealthy ($74k) than less-frequent attenders ($78k) and never attenders ($80k). Average congregant income varies by religion from Jews ($104k) and Orthodox Christians ($91k) to Protestants ($76k) and Jehovah's Witnesses ($71k). Perhaps an

[4] Some, but not all, of this difference in religious attendance frequency can be picked up by the survey data. For example, the survey data suggests that there are 14 times more Americans identifying as Catholics than Latter-day Saints. The survey data further suggests that there are approximately 5 times more weekly-attending Catholics than Latter-day Saints (where the data suggests more weekly-attending Latter-day Saints than Catholics)

equally interesting statistic is the variance of income by religion. The religions with the most economic diversity nationwide are Jews and Buddhists while those with the least are Jehovah's Witnesses and Latter-day Saints (Mormons).

It is possible that a religion with a large amount of income variation nationwide is still very segregated in terms of worship experience due to having some very wealthy congregations and other very poor congregations. Using the arrival time of each worshipper, I calculate the income difference between two random congregants who are at the exact same religious location at the same time. Across all religions, two congregants worshipping together are likely to differ in income by $21k. The group with the most income variability at the same worship service is Muslims ($31k) while the most economically uniform congregations are Latter-day Saints and Jehovah's Witnesses ($13k and $16k). I can compare the level of economic diversity of religious congregations to the economic diversity of other gathering spots. For example, the average income difference between two people shopping at the same Walmart at the same time is $21k (similar to the average religious congregation), Trader Joe's is $35k, and Nordstrom is $45k. Overall, it appears that religious congregations tend to be a bit more homogeneous than the majority of other locations that I consider, perhaps reflecting the fact that many congregations gather from immediate surrounding neighborhoods.

It is common in economics to use religiosity across geographic regions as a correlate with other measures. For example, religiosity has been shown to correlate with social mobility, productivity, and human capital (Chetty et al. 2022; McCleary and Barro 2003; Balan and Knack, 2012). There are several reasonable measures of religiosity by geographic area that have been used in the literature. For the sixth finding of this paper, I produce a new measure of religiosity based on cellphone data for counties, commuting zones, and states that can be used by other researchers. The correlation between my measure and others commonly used in economics is high, but far from a perfect correlation.

I merge in weather data to produce the seventh finding of this paper. I show that temperature and precipitation on the day of service are strong and significant predictors of religious worship attendance. The relationship shows that cold temperatures and precipitation lead to less worship attendance.

For the eighth and final finding, I exploit the fact that cellphone data can observe visits to locations other than religious establishments. One could easily produce a whole host of correlations (e.g. are frequent religious attenders more likely to go to Walmart or Target?). I focus, however, on correlations between religious attendance and going to locations that are often linked with religiosity in some way (e.g. many religions consider certain locations to be morally dubious). Specifically, I look at visits to strip clubs, liquor stores, tobacco stores, casinos, and gym and fitness centers. Importantly, these results are not meant to be causally interpreted. In fact, many of them could easily be explained by factors such as the age and gender composition of religious adherents. However, the simple correlations in visit patterns can still be of interest from a descriptive perspective. I find that frequent religious attenders are less likely to go to strip clubs, liquor stores, casinos, or go to the gym. For example, approximately 2% of individuals who are unaffiliated or attend religious worship services yearly go to a strip club during my sample period whereas 1.4% of monthly attenders and 0.9% of weekly attenders go to a strip club. Strip club attendance is highest among Muslims (5%) and lowest among Latter-day Saints (1.7%).

I show that the eight key findings summarized above are resilient to several key robustness checks. For example, the definition I use for a worship visit (a visit that takes place on the primary day of worship service) may be considered too narrow. Perhaps there would be far more weekly-attending worshippers if I allowed, for example, a Baptist to "worship" in church during the week and not just on Sunday. When I relax my definition of religious worship to allow for a visit at any point during the week to count as a religious worship visit, I find that 9% of Americans are weekly attenders as opposed to 5%. Almost surely 9% is too high given that it captures people who show up regularly at a religious location for preschools or other events. But it provides a nice upper bound. I also show how the key results vary when changing the cellphones that are included in the sample (for example, choosing phones that are even more likely to be providing regular location data). I find that the results are incredibly robust to different sample restrictions.

While cellphone data has many advantages relative to survey data, there are also some important limitations. I will briefly describe three important limitations here. First, the data do not represent individuals who attend worship services, but leave their phones at home. This could lead to estimates of religious worship attendance that are biased downwards. Perhaps more importantly though, the bias may differ by religion. For example, orthodox Jews attending synagogue services on the Sabbath will be drastically undercounted due to prohibition of phone use on the Sabbath. To help address this limitation, I run a large survey that asks whether people who attend religious services take their cell phone with them. I find that 87% of individuals say they always or almost always take their cellphone with them. So, while this is certainly a limitation of the study, the results are likely only biased downward by a small amount. Second, it is difficult for me to measure worship that takes place in a location that is not designated as a place of worship. For example, congregations that meet in a park, a school, or rent out space during the weekend at a place of business will be undercounted. Third, the cellphone data do not allow me to identify the race, gender, or age of the cellphone user – all variables that would be interesting to analyze if they were available.

This paper adds to a growing literature on the economics of religion. However, this paper differs from a typical paper in this space. Most papers in this field attempt to address one of two causal questions: a) what leads to increased or decreased religiosity (religiosity as the dependent variable)? And b) what are the impacts of being religious on various outcomes/behaviors (religiosity as an independent variable)? Regarding the first question, evidence has shown that the many factors can have a causal impact on people becoming more religious including earthquakes (Levy and Razin 2012; Bentzen 2019; Sibley and Bulbulia 2012), drought (Dube, Blumenstock, and Callen 2022), the need for social insurance (Ager, Hansen, Lonstrup 2014; Ager, Hansen, and Lonstrup 2017; Hungerman 2015), income (Buser 2015), charity (Gruber 2004), and education (Hungerman 2013; Bazzi, Hilmy, and Marx 2020). In regards to the second question, research has provided causal evidence that increased religiosity can lead to lower levels of depression (Lucchetti, Koenig, and Lamas 2021), fewer deaths of despair (Giles, Hungerman, and Oostrom 2023), more prosocial behavior (Benjamin, Choi, and Fisher 2016; Bottan and Perez-Truglia 2015), less crime (Moreno-Medina 2023), and less drinking and drug use (Gruber and Hungerman 2004). In addition to estimating causal parameters, the economics of religion has produced theory that attempts to understand the marketplace for religion and religious observance (see reviews by Iannaccone (1998) and Iyer (2016)). My paper differs markedly from previous work that addresses the causal questions listed above and instead provides a more descriptive analysis of religious observance. The hope is that this deeper understanding of what day-to-day religious worship attendance looks like and the introduction of a new, high-frequency measure of religious attendance can lead to further research that better understands the reasons for and impacts of religious observance.

The paper proceeds as follows: Section I reviews survey data results on religious participation and attendance in America. Section II discusses the cellphone data, its representativeness, and how I measure religious worship attendance using the data. In Section III, I present the 8 key findings. Section IV provides robustness to the main results by exploring various sample restrictions and definitions of what is a worship visit and also provides the results of the survey about phone behavior in a religious context. Section V concludes and provides a brief discussion of the broader implications and direction of this work.

I. Measures of Religiosity in America

Statistics on religious worship attendance in America come almost exclusively from surveys. Several religions report baptisms or other affiliation numbers, but it is very hard to find data on attendance/observance. Given the difficulty that many religions have with retention (Pew Research 2014), conversion statistics (e.g. baptisms) may not be a very good proxy for attendance. Some religions may gather attendance data, but use the information internally without sharing with the public.[5]

In Table 1, I review 8 primary survey sources for religious attendance data. For a few of these data sources (e.g. Baylor Religion Survey), the primary objective of the survey was to learn about religious practice. Others (e.g. Cooperation Election Study) ask about religiosity without that being the primary goal of the survey. For each of these 8 survey sources, I identify the year of data that allows for the most direct comparison with the cellphone data. I am careful to avoid data after March 2020 given the disruption that the COVID pandemic had on religious observance. The sample size for most of these surveys ranges from 1,000 to 30,000 people with the Nationscape data being an exception with over 300,000 respondents. Most of the surveys ask people if they are affiliated with a religion and go on to ask if they have attended a religious or worship service in the past year. Many of the surveys further ask if the person attends religious services monthly and weekly. The exact questions vary by survey, but are similar. For example, the affiliation question that is asked by Pew Research (the study that I will use as a comparison in the paper is) is, "What is your present religion, if any?" followed up with a list of options. The attendance question that is asked by Pew Research is, "Aside from weddings and funerals, how often do you attend religious services..." followed up with a list of options.

Overall, the results in Table 1 show reasonable consistency across surveys. 63-79% of Americans affiliate with a religion.[6] 45-65% claim to go to a religious worship service at least once a year, 30-42% claim to go at least once a month, and 22-30% claim to go at least once a week. The survey data also provide these statistics by religious tradition. I will show the Pew Research results by religion as a comparison to the cellphone results in Section 3 of the paper.

One additional important survey result is the American Time Use Survey (ATUS). This government sponsored and conducted survey asks about the time spent doing various activities the

[5] There are a few exceptions. For example the Episcopalian faith releases attendance records based on self-reported statistics from local leaders.

[6] The one exception is the 2020 US Religion Census that indicates 49% of Americans are religious affiliated. This is a survey that asks religious bodies how many affiliates they have (as opposed to asking individual people if they are affiliated).

previous day for approximately 2,500 people each year. The style of the survey makes the results less susceptible to social desirability bias than the other surveys in Table 1. The nature of the data does not allow for an estimate of weekly or monthly attendance, but it can be used to estimate the total number of people who attend worship services on a given day. Because I do not know the religious affiliation of survey participants, I focus on Sundays (the primary day of worship for the majority of Americans). The 2019 sample provides an estimate that 19% of Americans attended worship services on a typical Sunday.

As a comparison to these surveys, I present the main results using my cell phone sample in Table 1 as well. The rest of the paper will provide details for how these results were produced.

II. Cellphone data and Religious Attendance

Selecting cellphone sample. I use geolocation data from Veraset, a company that provides de-identified geospatial data for millions of U.S. smartphones. During our sample period, the company has a unique cellphone ID for tens of millions of phones (~225M) that transmit latitude-longitude coordinates ("pings"). The inclusion of a cellphone in the dataset is a bit of a black box, but is primarily a result of a user allowing location-tracking services on apps (e.g. the weather app). Many of the cellphones show up in the data only occasionally, whereas other cellphones are consistently pinging. This difference is likely the result of users only allowing location-tracking services while the app is in use vs users who allow tracking even when the app is not in use. A typical phone that pings consistently records a lat-long coordinate every 5 minutes or more frequently.

I use a Veraset dataset called the "Visits" data that uses a machine learning algorithm and identifies visits that a cellphone makes to various locations based on clusters of GPS signals. For example, if a cellphone pings within the geofence or polygon of a Walmart, Veraset will list Walmart as a visit and record the duration of the visit based on the first and last ping that occurred within the polygon. Veraset records visits to an individual's home, work, and all other business establishments and locations that are deemed "places of interest". The dataset that I use covers all visits made by cellphones between April 2019 and February 2020.

The selection of which cellphones to use in the analysis is important. The goal is to focus on phones where I am able to track their movement nearly continuously. The vast majority of cellphones in the Veraset sample do not ping regularly enough to allow for an accurate assessment of the visits that someone makes throughout the year. I calculate the number of days during the year (335 possible days between April 2019 and February 2020) that each cellphone in the Veraset data records at least one visit. I plot the distribution of days with at least one visit in Appendix Figure 1a. As can be seen, more than 35% of the phones in the dataset only have a visit on 1 day of the year and a majority of phones have visits on less than 20 days of the year. These phone are clearly not providing regular geolocation tracking. Appendix Figure 1b zooms in on the right tail of the distribution and shows that there are hundreds of thousands of cellphones that have at least one visit on the vast majority of days in the sample.

For the main analysis, I choose a cutoff of 325 days that a cellphone must have at least one visit in the dataset in order to be included in the sample. This leads to the inclusion of 2,131,844 cellphones.[7] Even though these cellphones provide pings enough for at least one visit for almost every day of the year in the sample, one might worry why they might be missing a visit on as many as 10 days during the year. A missing day could be caused by traveling out of the country, turning off a cellphone or not paying a bill that leads to it being turned off, etc. I provide robustness checks in Section IV to determine what happens if I limit the sample to cellphones with fewer missing days during the sample.

The 2.1M cellphone sample is pinging enough to register at least one visit nearly every day. But, are the phones pinging enough throughout each day to accurately reflect visits taken each day? On average, these 2.1M phones record ~4 visits each day. For example, a phone might show that someone spent time at home, at work, went to a restaurant, went to the grocery store, and then went back home. Because I know the duration of each of these visits, I can calculate on average how much of each day is accounted for by known location visits. For example, if a person was at home for 13 hours, work for 6 hours, and at the grocery store for 1 hour, then I was able to track them for 20 hours of the day. Appendix Figure 2 plots the distribution of average hours accounted for my main cellphone sample. I know the precise location of the typical cellphone in my main sample for approximately 18-20 hours of each day. Why is this less than 24 hours? One explanation could be that some of the cellphones are not pinging consistently enough throughout the day to fully account for their location. But there are other very good explanations as well. The visit data do not record information when the cellphone is not at home, work, or a place of interest. So, I would not expect 24 hours of coverage. For example, any time spent commuting, going on a walk, visiting a neighbor's house, etc. will not be part of a visit and therefore will be unaccounted for. So, it is not surprising that the total time that people show up

[7] Of these ~2.1M phones, I know the census tract of residency for ~1.4M of them. So, the sample is ~2.1M for most of the analysis, but ~1.4M when doing analyses that require census tract information (e.g. analysis of median household income).

in the visits data is less than 24 hours on average since most people commute and do other things during the day.

Overall, I think the 2.1M sample is a fairly conservative cut of the data and represents situations where phones are being tracked nearly continuously. However, for robustness, I run and report in Section IV several robustness checks using different and much tighter sample restrictions (both in terms of days with a visit and average hours during the day) and find nearly identical results.

Representativeness of sample. Even after assuming that I am able to continuously track the movement of 2.1M cellphones, a potential threat to the interpretation of my results is the representativeness of the sample. Just like a survey company that is worried about who is willing to answer the survey, I am worried about which smartphones are part of the sample and which are not. I do not expect the sample to be a perfect representation of the US. After all, the data do not represent the movement of small children who are unlikely to own a smartphone. The data may also under sample older or low-income Americans who may be less likely to own a smartphone.

The first approach that I take to think about representativeness is to compare the cellphone sample along certain dimensions to the US as a whole. Appendix Figure 3 plots the population of each state as a percentage of the US alongside the number of cellphones in each state as a percentage of the cellphone sample. Overall, the geographic distribution of cellphones mirrors the US fairly well. States with a higher population are states where the sample contains more cellphones. There are fewer cellphones in New York and California than would be commensurate with their populations and more cellphones in Texas, Missouri, and Illinois. Appendix Figure 4 plots the distribution of median household income across the US and the median household income across the cellphones in the sample. As expected, the cellphone sample skews a bit wealthier than the US as a whole. I also compare the population density of the cellphone sample with the US as a whole in Appendix Figure 5. The cellphone data undersamples super rural and super urban areas a bit, but matches the population density of the US nicely.

Another approach that helps to understand the representativeness of the cellphone sample is to identify businesses or locations that have a known number of visitors and then compare that to how many visits I would have predicted based on the cellphone sample. For example, sports teams provide attendance numbers on game days. I can check to see how many visits to the game occurred from the 2.1M sample of cellphones, multiply up to the US population level, and compare my predictions to the truth. If my sample is doing a good job representing the US as a whole, then the predictions should consistently match the truth.

The first test along these lines is with a random sample of home games for every major league baseball team in 2019. Appendix Figure 6 plots the actual attendance for each baseball game against the attendance that the cellphone data predicts. Dots above the 45-degree line are situations where the cellphone data estimate a higher attendance than the actual attendance while dots under the 45-degree are the opposite. The average prediction from the cellphone data (30,364) is almost exactly the same as the average actual attendance (30,313). Further, the data reasonably fit the variation in attendance across stadiums. I do the same analysis for all NBA teams in Appendix Figure 7. I undercount on average in this case by about 30% (average attendance prediction from cellphone data: 12,532; True attendance: 17,874). Undercounts are especially pronounced in high-urban areas in California and New York (Clippers, Lakers, Warriors, Nets, Knicks, etc.).[8]

Other than sporting events, it is hard to get actual attendance numbers for locations. Table 2 attempts to look at some additional examples. I compare how many visits are predicted by the cellphone sample to the number of actual visits to AMC theaters, Lowe's, McDonald's, Olive Garden, Planet Fitness, Six Flags, Target, The Home Depot, and Walmart. For many of these examples, I don't actually know the number of visitors like I do at sports stadiums, but I know the number of transactions.[9] Transactions isn't a perfect proxy for visits since people can enter a store/restaurant and not buy anything or multiple people can go to the store/restaurant together and leave with just one transaction, but it could serve as a reasonable proxy for total visitors.

I find that that overall the cellphone data does a reasonable job of predicting actual attendance measures (often being proxied for by transactions). If anything, the cellphone data appears to over count the number of people walking into a store (e.g. Target). This could be due to the fact that multiple people might walk in and have only one transaction. The one location where I undercount total visits is AMC theaters (undercount by ~30%). A partial explanation for this is that the number of AMC theaters listed as places of interest in the cell phone data (539) is less than the actual number of AMC theaters (636) for some reason. The 15% fewer theaters can explain about half of the undercounting that I find for AMC theaters. As a whole, the cellphone data appear to do a reasonable

[8] I have explored why undercounting occurs for NBA games. Part of this could be driven by basketball arenas being located in highly dense urban areas. Exploring the data also reveals cellphone visits to stores/shops/bars that are located within the NBA arena. Someone visiting a bar inside the arena for 3 hours is likely to be someone who bought a ticket and attended the game itself, but leads us to undercount game attendance.

[9] In some cases I don't know the number of transactions, but I know the total revenue and the average transaction amount and can back out a reasonable approximation for transactions.

job predicting actual attendance, although one may reasonably concerned by over or under predictions in some cases.

Identifying religious visits. Each visit in the dataset is categorized using the North American Industry Classification System (NAICS). For the sample of 2.1M cellphone users, there are ~32M total visits to a location that is classified as a religious location by NAICS. The data contain the name for each of the religious locations (e.g. "Saint Mary's Church of God in Christ"), but the locations are not categorized by religious tradition. Research assistants categorized each religious location with a religion. In many cases, this was easy given the name of the church (e.g. if the church's name had the word Presbyterian in it, it was classified as Presbyterian and Protestant). In other cases, the church had to be Googled and sometimes even then it was unclear and had to be categorized as non-denominational Christian. In categorizing, I followed the religious classification set forth by Pew Research and established the following primary religious affiliations: Protestant, Catholic, Church of Jesus Christ of Latter-day Saints (Mormon), Jehovah's Witness, Orthodox Christian, Jewish, Muslim, Buddhist, Metaphysical+, Other Non-Christian Religions, Unitarian+, and Hindu.[10] Following Pew Research, I categorize non-denominational Christian churches as Protestant.[11] [12]

Churches and other religious locations can be used for activities in addition to religious worship services. For example, many church buildings are used for preschools, choir practices, weddings, funerals, etc. In order to avoid counting preschool visits as religious worship visits, I define a religious worship visit as one that takes place on the primary worship day for each religion.[13] I relax

[10] Metaphysical+ is called "Other Christian" in the Pew Research surveys. It includes Metaphysical and other Christian traditions (but not including non-denominational Christian). Other Non-Christian includes Baha'i, Jainism, Sikh, and Taoism. Unitarian+ includes Unitarian, Scientology, and Freemasonry.

[11] In some reports, Pew Research subdivides Protestants into Evangelical, Mainline, or Historically Black Protestants. I am unable to do this subcategorization by just looking at the names and locations of churches. For example, Pew categorizes Historically Black Protestant churches as those where the majority of attendees and leaders are reported to be black (something I can't do with the data). Pew further subdivides, for example, Evangelical Protestants into denominations (Baptist, Lutheran, etc.). I tried to do this with the data, but worry that the categorization is far from perfect. For example, if "Baptist" wasn't in the name of the church, I am likely to have categorized it as non-denominational Christian under the Protestant umbrella. Because of the categorization error, I treat Protestant as a unified category in all analyses.

[12] In the results section, I will provide overall results including all religions. When breaking down the results by individual religion, I will focus on the 8 most visited religions in the sample and therefore exclude Metaphysical+, Other Non-Christian, Unitarian+, and Hindu. These four categories are very small in the US context and are difficult to analyze due to small sample sizes in many of the analyses.

[13] I define worship days for Protestants, Catholics, Latter-day Saints (Mormons), Jehovah's Witnesses, Orthodox Christians, Metaphysical+, and Unitarian+ to be on Sundays with the following exceptions: Seventh-Day Adventists worship days are on Saturdays and Freemasonry can occur on any day of the week. For Christian religions, I also include Christmas and Christmas Eve as worship days since special worship services often occur on those days (Easter is also a special day, but always occurs on a Sunday). I define worship days for Muslims to be on Fridays. I also include four additional days as worship service days for Muslims; two days each around Eid al-Fitr and Eid al-Adha. I define worship days for Jews to be from Friday at 5pm to end of day on Saturday and I include 8+ additional days for four holidays

this restriction as a robustness check in Section IV to see what religious attendance looks like if I include visits to religious locations even on days that aren't typical worship days.

For each individual in the sample, I tabulate the total number of religious worship visits that they made during the 47 weeks of the sample. The max number of visits allowed for an individual is one per week. Thus, total religious worship visits range from 0 to 47. Lastly, I also categorize each person into a particular religion. If an individual had no religious worship visits in the sample, they are categorized as a non-attender. If they had one or more visits, I categorize them into the religion that they visited the most (e.g. if they went to a Jewish synagogue six times and a Catholic church once, they are categorized as Jewish). In the event of a tie (uncommon), they are categorized at random.

I make one more important adjustment to all the results that are reported in the paper (see Appendix 1 for full details). The total number of religious places of interest that the Safegraph/Veraset data contain during my sample period is ~228k. This number is lower than it should be. Estimates suggest that there are ~380,000 total churches or other religious places of worship in the US (National Congregational Study Survey). Safegraph/Veraset has been continually updating its places of interest (POIs) database, and religious locations clearly were not fully represented in 2019. One might be tempted to simply inflate all of the results by some weight in order to account for the missing POIs. However, this would not be wise since the missing data might not be random. For example, the missing POIs might be places of worship that are smaller or more likely to be from a particular religious tradition. Fortunately, I am able to identify what POIs are missing in the data. To do this, I obtain Safegraph POI data from Fall 2023 (downloaded from a 3rd party data provider, Dewey). By 2023 Safegraph had updated its POI file and the new dataset contains ~392k legitimate religious POIs (much closer to the estimated 380k religious locations in the US). Research assistants once again went through and coded up the religious tradition of the new POIs. Using this dataset, I can up weight the estimates in the paper by the appropriate number of missing religious POIs in each religious tradition. As described in Appendix 1, I also use the square footage of each religious POI in the weighting process to account for potential differences in the size of the missing religious POIs relative to the non-missing data. I believe this process allows me to carefully account for the missing 2019 data, but certainly also adds a degree of noise to the estimates in the paper.

(Rosh Hashanah, Sukkot, Simchat Tora, and Yom Kippur). For Buddhists, Hindus, and Other Non-Christians, I define worship days to be any day of the week except for Baha'i which is restricted to Fridays.

III. Results

Result 1: Overall attendance. The first key finding shows overall religious worship attendance using the cellphone sample. In the figures that follow, I show both the cellphone results along with the Pew Research survey results as a comparison. The top bar in Figure 1 shows that 66% of Americans surveyed claim affiliation with a religion. I do not have a comparable number from the cellphone data because I am not able to see stated affiliation with movement data. The second set of bars in Figure 1 show that 46% of survey participants claim to have gone to a religious worship service at least once in the past year. Based on cellphone data, I find a larger number (73%) of people who visit a religious place of worship on the primary day of worship (e.g. Sunday for most Christian churches). There are several reasons why I might find larger effects for visiting at least once a year than are reported in surveys. For example, in the survey question asked by Pew Research, it asks about religious worship attendance NOT including weddings, funerals, etc. In my data, I am not able to exclude weddings, funerals, choir concerts, or other social gatherings that might occur in a place of worship on the primary day of worship. The 3rd set of bars in Figure 1 show that 30% of surveyed individual claim to go to church at least once a month. By comparison, I find that only 21% of individuals in the cellphone data attended a worship service at least 11 out of 47 times during the sample frame. Lastly, the final set of bars show that 22% of people claim to attend worship services at least weekly. However, I find that only 5% of the cellphone sample attended at least 36 out of the 47 weeks of the year (I generously allow a "weekly" attender to miss once per month due to sickness, travel, etc.).

Result 2: Overall attendance by religion. The results presented in the previous paragraph mask interesting patterns that are taking place across different religions. In Figures 2 Panels A-H, I show the same result separated out by religion.[14] Not surprisingly, Protestants are the dominant sect in America. I assign 55% of Americans as Protestants in the cellphone sample that attend church at least once each year. Thus Protestants make up 75% of Americans who attend at least one worship service. Protestantism becomes even more dominate when focusing on weekly attenders: 3.90% of Americans are weekly-attending Protestants (85% of the total weekly attenders in the sample).

Catholics are the next largest group when it comes to attending at least once a year (13.45% of Americans). Yet, there are relatively few weekly-attending Catholics (.26% of the US population).

[14] I show the results for just the 8 largest religious groups in these 8 panels.

Similarly, 1.87% of Americans are Jews that attend at least once per year, yet the number of weekly-attending Jews is very small (.01% of the US population).

In contrast, Latter-day Saints (Mormons), and to a lesser extent Jehovah's Witnesses and Orthodox Christians, have a much smaller base of members but a large number of weekly attenders. For example, 1.99% of Americans are Latter-day Saints that attend at least once per year and .29% of Americans are weekly-attending Latter-day Saints. Thus, approximately 1 out of 7 people who I classify as a Latter-day Saint attends weekly.

Result 3: Week-to-week variation in attendance. Unlike survey data that asks about overall attendance patterns, the cellphone data allow me to analyze religious worship attendance week-by-week over nearly a one-year period. Figure 3 estimates the total number of Americans who attended a religious worship service each week of the year from April 2019 to Feb 2020. In a typical week, I estimate that ~45M Americans attend a worship service (~14% of the US Population). One can think of this as consisting of the 5% of Americans who are weekly attenders plus a portion of the 21% who are monthly attenders plus some of the yearly attenders who are scattered in throughout the year. The two key spikes in the data that can be seen in Figure 3 are Easter and Christmas. Easter Sunday and the week of Christmas experience attendance that is nearly 50% higher than a typical week. I also see small drops in attendance on holiday weekends (Memorial Day, Labor Day, and Thanksgiving).

Figure 4, Panels A-H, show these same week-by-week attendance figures for each religion separately. These figures yield some very interesting religion-specific effects. For example, I see large increases in worship attendance for Muslims during Ramadan and ending with Eid al-Fitr and also an attendance spike for Eid al-Adha. I also see large increases in attendance for Jews on Yom Kippur and Rosh Hashanah. I see a large increase in attendance for Buddhists towards the end of the calendar year that coincides with New Year celebrations. A few other quirks of religious attendance also show up in these figures. For example, unlike nearly all other Christian religions, Latter-day Saints (Mormons) do not experience an increase in attendance around Easter or Christmas, but there are two huge drops in attendance for Latter-day Saints when religious services are canceled and replace with semiannual online conferences.

In addition to the interesting spikes and troughs in the data series, these figures provide perhaps the best statistics available for how many people of different religions attend worship services every week. For example, the fact that ~38M Protestants, ~6M Catholics, ~2M Latter-day Saints, ~400k Jehovah's Witnesses, ~220k Orthodox Christians, ~400k Jews, and ~220k Muslims go to a worship service in a given week may even be interesting and novel information for these religions themselves where detailed data are not kept.

Result 4: Start time and duration of attendance. Because of the granular geodata, I can determine with a high degree of accuracy when (using local times) someone enters a religious place of worship and when they leave. This allows me to identify differing patterns of religious worship behavior. Overall, most religious worship in America is occurring on Sunday mornings. But there are interesting patterns by religion shown in Figure 5 Panels A-H. Protestants primarily worship on Sundays with main arrival times between 9 and 11am. For Catholics, arrival is between 8:30 and 11am in the morning, but then there is also a group of worshippers who arrive for evening services (e.g. 5pm). Latter-day Saints (Mormons) have three clear peaks of arrivals that take place during Sunday mornings. Jehovah's Witnesses also have 3 main peaks, but spread more evenly throughout the day. Orthodox Christians arrive mostly in the morning from 9-11am. For Jews, I allow a worship visit to occur on Friday night or Saturday. I see arrival times occurring both in the morning and the evening. The primary worship day for Muslims is Friday and I see a peak at 1pm which coincides with the time of prayer. Buddhists show very irregular start time patterns, which is consistent with their worship experience being more individual.

The data also can be used to indicate the typical duration of a religious worship visit. Across the entire sample of religious worship visits, the mean duration is 97 minutes. Once again, this masks interesting differences that exist across religious traditions. I plot the distribution of duration by religion in Figure 6 Panels A-H. The religions with the longest average visit duration are Orthodox Christians (116 minutes), Latter-day Saints (115 minutes), and Jehovah Witnesses (115 minutes). The religions with the shortest average visit duration include Muslims (51 minutes), Catholics (66 minutes), and Buddhists (71 minutes). Jews (92 minutes) and Protestants (102 minutes) have average durations in the middle of the distribution.

Result 5: Income. As discussed in Section 3, I obtain median household income for cellphones in the sample based on Census tract information for the location of their home. With this in hand, I can explore the level of household income associated with members of different religions and how income differs by frequency of attendance. I can also explore diversity of socio-economic status across religions and across congregations within a religion.

Panel A of Figure 7 shows the overall median household income distribution of the Americans who are never attenders and the Americans who attend at least once per year. As can be seen, these distributions are very much on top of each other with only a slight difference in mean income (Never attenders: $80k; Attenders: $79k). Panel B of Figure 7 splits the attenders into narrower groups based on frequency of attendance (those who attend yearly or more, monthly or more, and weekly). I find that the average income of weekly attenders ($74k) is significantly less than non-weekly attenders and never attenders.

In Appendix Figure 8 Panels A-H, I also show the income distributions separately for each religion by frequency of attendance. Some interesting patterns emerge. As noted in the previous paragraph, weekly attenders are on average lower income than less frequent attenders. However, there are three exceptions to this rule. For Latter-day Saints (Mormons), Muslims, and Buddhists, the weekly-attending members are actually more wealthy than the less frequent attenders.

In Figure 8 Panels A-H, I graph the distribution of median household income separately for each religion. In each panel, the mean and standard deviation of attendance is noted in the upper right-hand corner. The religious worshippers with the highest income are Jews ($104k), Buddhists ($99k), and Orthodox Christians ($91k) while the religious worshippers with the lowest income are Jehovah's Witnesses ($71k) and Protestants ($76k). Based on the standard deviations, I can also identify religions with the most nationwide socio-economic diversity. These include Jews and Buddhists whereas the religions with the least nationwide socio-economic diversity are Jehovah's Witnesses and Latter-day Saints (Mormons).

The statistics discussed in the previous paragraph describe the level of socio-economic diversity for each religion when using data across the entire US. However, those statistics do not reveal the level of economic diversity that exists within a particular congregation. For example, it is possible to have a large amount of socio-economic diversity in a religion nationwide, but congregations can be very homogenous (e.g. some very rich and some very poor). The cellphone data uniquely allow for the measurement of diversity within a congregation by focusing on cellphones that are in the same location at the same time.

In every instance where two people are worshipping in the same religious location and at the same time, I compare the income of a congregant to a randomly assigned co-congregant. For example, if a religious location has 10 people all worshipping at the same time, each of the 10 congregants will be randomly paired with one of the other 9 congregants. The income difference between these random pairs is calculated. For each religion, I can then plot the distribution of income differences between randomly assigned pairs who are worshipping together. In Figure 9, I show the overall difference in income across all religions for people worshipping together. Figure 9 reveals that the average income difference between two random worshippers in America is $21k. One thing to note in Figure 9 is that there is a large spike in the distribution at exactly $0 of income difference. This occurs when two people who live in the same Census tract are randomly paired due to showing up at the same congregation at the same time. Because many places of religious worship draw from a very local area, it is not surprising that there is a spike at zero as neighbors who live in the same census tract might be worshipping together.

In Figure 10 Panels A-H, I show the distribution of income differences for co-worshippers separately by religion. Some religions have congregations that are much more homogenous that others. The most homogenous religions in the data are Latter-day Saints (mean absolute income difference: $13k) and Jehovah's Witnesses ($16k). The most diverse congregations are Muslims ($31k). This difference may in part stem from institutional differences across religions. For example, Latter-day Saints (Mormons) are assigned to a congregation based on the location of residence. This can lead to many people with similar median household income (as measured with Census tract data) attending the same church at the same time. Muslims on the other hand tend to congregate from a more disperse geographic area leading to more socio-economic diversity.

A natural question might be how the economic diversity of religious congregations compares to the economic diversity of other gathering places (restaurants, amusement parks, retail stores, etc.). I analyze the data for other locations in the same fashion as was done for co-worshippers. For example, I compare the income of two randomly assigned pairs of individuals who are shopping at the same Walmart at the same time and calculate the mean absolute income difference among co-shoppers at Walmart around the country just as I did with religious congregations. In Figure 11 Panels A-F, I plot the results for six different locations. The most economically diverse location in Figure 11 is Nordstrom where the average income difference between two shoppers is $45k. This is significantly larger than the average difference between two co-worshippers in America ($21k). The other examples include Disney World ($35k), Trader Joe's ($35k), Walgreens ($21k), Walmart ($21k), and Applebee's ($18k). Overall, it appears that religious congregations are less diverse than many public gathering spots and somewhat similar in terms of economic diversity with stores and restaurants that serve a mostly local cliental.

Result 6: Geographic measures of religiosity. It is not uncommon in economics to correlate geographic measures of religiosity with various economic outcomes. To do these correlations, one needs to have a reasonable measure of religiosity at the state, county, or sometimes commuting zone (CZ) level. Various measures have been used in the literature, most commonly the Religion Census.

The cellphone data allow for a measure of religiosity that is tied to actual attendance as opposed to surveys of individuals or religious bodies. I compute state, county, and CZ measures of religiosity where religiosity is measured as the fraction of people in the area who attend at least once a year or more, once a month or more, or weekly.[15]

Figure 12 Panel A provides a state-level map for the percentage of people in each state who had at least one religious visit during the sample period. Panel B maps the same information at the county level. These maps show that religious attendance is highest in the South, Midwest, and Utah. Similar maps can be created for religiosity as measured by more frequent attendance (monthly+ and weekly).

How does the cellphone measure of religious attendance across geographies compare to the other commonly used measure of religious observance? In Figure 13, I provide a scatter plot at the state level that graphs the number of yearly+ attenders by state from the data against the fraction of the population that are considered religious adherents by the US Religion Census (based on a survey of religious bodies). Overall the correlation between my measure and the Religion Census measure is high (R-squared: .40), but leaves plenty of room for differences that could exist between the measures.

Result 7: Weather and religious attendance. Using data from the National Oceanic and Atmospheric Administration (NOAA), I analyze precipitation and maximum temperature for each census tract and each day in the data. I merge these data at the census tract level with the home locations for each of the cellphones in the sample. With this data in hand, I can explore how church attendance varies with respect to weather changes.

Figure 14 Panel A shows the relationship between worship attendance and precipitation.[16] Specifically, I run a regression at the individual level with fixed effects for date*religion and individual.[17] Thus I am absorbing a person-specific mean and a date*religion specific mean and identifying off of variation in weather that occurs across time and geography. The date*religion fixed effects help capture any jumps or drops in attendance that may occur for specific religions and at certain dates. For example, it helps to capture the fact that Latter-day Saints (Mormons) have two online conferences each year that lower attendance to almost zero. Without these fixed effects, a particularly wet day on one of those days could lead to very biased estimates. I restrict the sample to

[15] I assign a missing value to geographic areas (e.g. counties) that have 5 or fewer cellphones in the sample and label those areas as "insufficient data" when showing maps.
[16] Precipitation is measured in millimeters of rain plus the liquid equivalent of any frozen precipitation (e.g. snow).
[17] Because of the individual-level fixed effects, the sample is restricted to individuals who attended worship services at least once during the year.

locations that worship on Sunday.[18] Relative to a base group of no precipitation, living in a location that experienced precipitation on a particular Sunday leads to approximately 0.15 percentage point lower probability of attendance. The base rate level of attendance for this sample is 18.8%. Thus, precipitation leads to approximately .5% lower attendance. Perhaps surprisingly, I do not see differences between a small amount of precipitation (less than 5 millimeters) compared to location-days with a larger amount of precipitation (greater than 5 millimeters). While significant and robust, the effect is economically small. For every 200 people going to church, only 1 is deterred when it rains on the day of the service. One potential reason for the small effect is that I am using rain that occurs at any time on the day of worship. Presumably, the effect would be larger if I restricted to rainfall that occurs in the morning hours (the time of most Sunday worship services).[19]

Panel B of Figure 14 shows the impact of temperature on worship attendance. Once again, I run a regression at the individual level with fixed effects for date*religion and individual. Relative to a base group where the daily high temperature is less than 25 degrees Fahrenheit, warmer weather (e.g. temperature more than 35 degrees leads to an increase in attendance by approximately 1.5 percentage points. The base rate level of attendance for this sample is 18.8%. Thus, a daily high temperature that is greater than 35 degrees leads to approximately 5% higher attendance relative to cold days (temperature < 25 degrees).

Result 8: Religious attendance and other religious-adjacent activities. Given the nature of the cellphone data, it is possible to not only see worship attendance, but also to analyze visits that people make to other establishments. I could produce a nearly infinite number of correlations (Are people who attend religious services frequently more likely to go to the grocery store?, Do they prefer Walmart or Target?, etc.). These correlations, however, are not necessarily meaningful in an article trying to better understand religion in America. However, there are locations that I can identify in the cellphone data that at least tangentially relate to religion/morality. For example, many religious traditions discourage their congregants from going to strip clubs. Thus, it is somewhat relevant to look at the correlation between strip club attendance and religious attendance to see if – at least descriptively – frequent worshippers are less likely to attend strip clubs. I identify five different visit locations that

[18] This restriction makes for a very simple analysis even though it leaves out a few smaller US religions (e.g. Buddhists). Because, for example, I allow for Buddhists to worship on any day of the week, a regression would need to be done with 335 days as opposed to 47 days. This does not fit nicely into the same regression framework, so for simplicity here I restrict ourselves to religions that worship on Sundays.

[19] Moreno-Medina (2023) finds a 17% drop (p = .04) in church attendance for any amount of precipitation that occurs between 9am and 1pm.

might be relevant in a moral/religious sense: strip clubs, casinos, liquor stores, tobacco stores, and fitness/gyms.[20] [21] Many religions counsel against alcohol and tobacco use and encourage worshippers to take care of their physical bodies (hence, fitness/gyms). Gambling is also frowned upon by many religions. Clearly these categories might be too broad and not capture behavior well. For example, walking into a liquor store isn't always for the purpose of purchasing alcohol. Furthermore, many religions do not consider moderate alcohol consumption to be immoral. Similarly, going to a gym does not necessarily signal someone who takes care of their body more than someone who doesn't go to the gym (especially if they work out at home, run outside, etc.). So all of these correlational results should be thoughtfully considered.

Panels A-E of Figure 15 show the percentage of people in the main sample who visited each of the five location types at least once during the sample period. The results are shown separately for the intensity of religious attendance. Panel A illustrates that approximately 2% of individuals who never attended a worship service in the data went to a strip club at some point. Worship service attenders overall (those who attend at least yearly or more) show a very similar strip club attendance rate (~2%). However, more frequent religious attenders go to strip clubs at a lower rate (1.4% for monthly+ attenders and 0.9% for weekly attenders).

Panel B shows the results for casino attendance by frequency of worshipping. ~37% of never attenders go to a casino during the sample. Overall, worshippers go to casinos slightly more often (~38%). But once again, frequent religious worshippers go less often (34% for monthly+ attenders and 28% for weekly attenders).

Panel C provides the results for visiting liquor stores. 46% of never attenders go to a liquor store at least once during the sample. Yearly+ worshippers are slightly more likely to visit a liquor store (47%). However, frequent attenders are less likely to visit (42% of monthly+ worshippers and 39% of weekly attenders).

[20] A couple of other locations were considered, but then discarded due to data issues. For example, a visit to a correctional facility or jail is likely to signal criminal behavior. However, the data made it very hard to identify people who go to a correctional facility due to an arrest. Phones are typically taken away upon entry and thus do a poor job of signaling a visit. I was also not able to easily parse out arrested individuals from police officers, lawyers, staff workers, and jail visitors.

[21] Visits to these 5 location types relied primarily on NAICS codes, but with a few deviations. Strip clubs are a narrower category than the NAICS codes to which they belong (motion picture theaters and drinking places). I identified strip clubs based on the name of the location with the help of Google searches. The NAICS code for liquor stores was slightly too broad (e.g. it included non-alcoholic soda locations), therefore, I restricted the set of locations to those that had "liquor", "wine", or "spirits" in the title. Similarly, I restricted tobacco stores to those that had "tobacco", "cigar", "vape", "smoke", or "vapor" in the title. No restrictions were needed for casinos and gyms.

Panel D shows the results for tobacco stores. Approximately 60% of never attenders go to a tobacco store at least once during the sample. Religious worshippers are slightly more likely to attend, even those that are frequent religious attenders (63% for Yearly+, 62% for monthly+, and 61% for weekly attenders).

Lastly Panel E shows the results for gym attendance. For this measure, I plot the average number of visits for individuals in the sample as opposed to the percent of individuals who go at least once. The reason for this change is that the results are truncated due to nearly 100% of the sample going to a gym or fitness center at least once during the sample period. Never attenders go to a fitness center 34 times on average. Religious worshippers attend more often (38 times on average), but frequent attenders less (35 times for monthly+ and 29 times for weekly attenders).

I can also look at visits to the five location types separately by assigned religion of worshippers in the data. These results can be found in Panels A-E in Figure 16. Panel A shows that strip club attendance is highest among Muslims (~5%) and lowest among Latter-day Saints (1.7%). For casino attendance in Panel B, visiting rates are highest for Buddhists, Jews, and Catholics and lowest for Jehovah's Witnesses. Panel C presents the results for liquor store attendance. The religions with the highest attendance rate is Jewish (60%) and Orthodox Christian (57%) while the lowest is Latter-day Saints (32%). Panel D shows that tobacco store attendance is highest among Jehovah's Witnesses and Buddhist and lowest among Latter-day Saints (Mormons) and unaffiliated individuals. Lastly, Panel E finds that fitness center attendance is highest among Buddhists, Catholics, and Jews and lowest among Jehovah's Witnesses.

IV. Robustness

In this section, I provide robustness results when thinking about two important changes to the analysis. First, I consider what happens when a visit to a church on any day of the week (not just the day of worship) is considered a worship visit. Second, I consider what happens when the selection criteria for which phones are included in the sample is adjusted.

A broader set of worship days. For the main results in the previous section, I define worship visits to be those that occur on the primary day of worship for each religion (Sunday for most Christians, Fridays for Muslims, etc.). Because churches are often used for preschools, arts/dance centers, funerals, weddings, etc., it seems unnatural to consider all visits to a religious location as a worship visit. However, perhaps my definition of worship is too restrictive and some people go to a place of worship on a day different than the primary day of worship and consider that to be worship.

In fact, some religions (e.g. Catholics) may have specific worship services during the middle of the week.

I relax that restriction in this section and consider someone who steps foot in a religious establishment at any point during the week to have participated in a worship service for the week. I do not allow for someone to worship more than once a week (if they go to the religious location 4 times in the same week, they still just get counted as having worshipped once that week). Thus to be counted as a "weekly attender", they would just need to go to a religious location at least once each week for 36 out of the 47 weeks in the data. But now, the visit each week can occur on any day of the week.

Figure 17 reproduces the main key finding #1 of the paper using this less restrictive measure of worship attendance. The top bar in Figure 17 continues to show that 66% of Americans surveyed claim to be affiliated with a religion and 46% claim to have attend a worship service at least once a year. When restricting to the primary day of worship only, I was finding that 73% of people attended at least one worship service in the sample – a very good match with the survey data. However, now that I use the much looser measure of worship attendance, I find that nearly all of the sample attended a worship service at least once. Almost surely this is including people who attended weddings, funerals, preschools, voting (which often occurs in a place of worship), etc. Before I was finding that 21% of people attended monthly or more. Now with the looser measure of worship attendance, I find that 39% of the sample are "monthly+ attenders". Lastly, I was finding that just 5% of the cellphone sample attended worship services weekly. When relaxing the worship day restriction, it is now 9%. Thus, even with the much less restrictive measure of worship attendance (that likely includes parents dropping of kids at preschool weekly, etc.), I am finding far fewer weekly attenders than the number of claim to be weekly attenders in surveys (22%).

I also reproduce the week-to-week variation figures and include what the results would look like if a visit to a religious location on any day of the week was considered a worship visit. These results are interesting because not only does it give an upper bound on religious worship, it also shows for each religion the amount of foot traffic at places of worship that occur on days other than the main day of worship. Figure 18 shows the total number of Americans who attended a religious worship service each week of the year from April 2019 to Feb 2020 using both the old definition of religious worship (a visit that occurred on the primary day of worship) and the new definition (a visit that occurred on any day of the week). Based on the old definition, I estimate that ~45M Americans attend a worship service (~14% of the US Population) each week of the year. This is primarily Christians attending church on Sundays. By contrast, I find that 75M Americans go to a place of worship each week of the year when attendance is allowed to occur on any day of the week.

Figure 19, Panels A-H, show these same week-by-week attendance figures for each religion separately using the two different measures of religious worship. These figures allow one to gauge how much visitation is occurring to places of worship for each religion on the primary day of worship and on other days during the week. Religions vary in how much larger the worship visits look when relaxing the primary worship day restriction. Approximately ~38M Protestants worship each Sunday, but this effect is ~50% higher when relaxing the primary day of worship restriction. ~6M Catholics worship each Sunday, but it is 50% higher with looser assumption. ~2M Latter-day Saints, but it is 35% higher. 400k Jehovah's Witnesses, but 50% higher. 220k Orthodox Christians, but 40%, 400k Jews, but nearly 100% more, 220k Muslims, but 80% more. The Buddhist figure does not change, because I was already counting any day of the week as a worship day.

Main sample restrictions. As discussed in Section II of this paper, an important yet somewhat arbitrary decision was made to limit the main sample to cellphones that recorded at least one visit on 325 of the 335 days in the sample. This is to ensure that I am capturing phones that are pinging regularly. I did not restrict the sample based on the average number of hours in each day that the phone is in a known location, but that is also a somewhat arbitrary choice.

In Table 2, I document what happens to the main results of the paper (% of people who are yearly+, monthly+, and weekly attenders) if further restrictions are placed on the data. The top line of the table shows the results from the main sample (using the restriction of 325+ days with a visit and no average hour restriction). The main sample has ~2.1M cellphones and I find 73% yearly+ attenders, 21% monthly+ attenders, and 5% weekly attenders.

The next five rows in the data place tighter restrictions on the number of days in the sample with at least one visit. I increase the required number of days from 325+ to 327+, 229+, 331+, and 333+ out of the 335 days in the sample. When imposing the most binding restriction (333+) the sample size drops to below 1M. However, this smaller sample continues to produce nearly identical results for the attendance for each frequency.

The middle section of Table 2 returns to the 325+ days restriction (just like the main sample), but imposes increasingly tighter restrictions on the average number of hours that I track people. Since the vast majority of the main sample was already being tracked for most of the day on average, restricting to 12+ hours does not drop the observation count by much at all. However, requiring an average hour of 20+ (the most restrictive robustness check) shrinks the dataset by about 75%. As a reminder, restricting to seeing someone with 20+ hours means that I do not allow that cellphone to spend much time outside of home, work, and specific visit locations. This restriction will remove cellphones for anyone who commutes a large amount, spends significant time outdoors (walking, biking, etc.), or even spends a lot of time at a friend's house. The results show that when requiring a higher number of average hours of coverage each day, the number of yearly+ attenders decreases by a significant amount (as low as 59% in the 20+ restriction) and monthly+ attenders decreases by a small amount (to as low as 17% in the 20+ restriction). The percent of the sample who are weekly attenders remains at 5% even when the sample is cut to ¼ the size.

The lower section of Table 2 imposes a more restrictive number of days in sample requirement (330+) and also restrictions for average hours of coverage each day. The findings are very similar to those discussed in the previous paragraph.

Overall, the main results of the paper are reasonably consistent and robust when choosing the main sample with very different restrictions. This leads me to believe that the results are not sensitive to the arbitrary decisions that led to choosing the 2.1M main sample.

Cellphones at places of worship. A concern that one might have regarding using cellphone data to measure religious worship is if worshipers do not take their cellphone with them to the service. This is certainly a problem with some very small religious groups (e.g. Orthodox Jews) where cellphones are unlikely to be used on the primary day of worship. But is this a problem more broadly? To shed light on this data limitation, I run a large survey (N = 4,993) on Prolific. I ask each survey participant about their religious affiliation and frequency of attendance and present those results in Appendix Figure 9. The sample is slightly less religious than the self-reported data from the Pew Research data (e.g. ~16% of the Prolific sample claim to attend services once a week compared to 22% in the Pew data), but is broadly consistent with national patterns.

I then ask those survey participants who attend religious services at least once a year about whether they bring their cellphone with them to services and report these results in Figure 20 of the paper. Overall 87% of respondents claim to "always" or "almost always" take their cell phone with them to worship. The percentage of people taking their phones to worship with them is similar across group with differing levels of attendance frequency. There is some variation by religion in the response rates about taking cellphones to worship services. For example, Jews are the least likely to report taking cellphones to services (e.g. 70% claim to always or almost always do so). More than 80% survey participants from all other religious groups report taking their cellphone to worship always or almost always.

The bottom three bars in Figure 20 compare the percent of people who take their cellphone with them to religious worship services compared to claims about taking cellphones to movie theaters and grocery stores. The percent of people who take their cellphones to religious services always or almost always (87%) is less (but not dramatically less) than the percent who take their cellphones to movie theaters (95%) and grocery stores (96%).

Clearly questions remain about the representativeness of the Prolific sample and the accuracy of self-reported cellphone carrying, but these results at least provide some evidence that the bias due to leaving cellphones at home during worship services is not too large.

V. Discussion and Conclusions

Religion is an important part of life for millions of Americans. Even though this paper provides evidence that the frequency of religious worship visits is lower than claimed in surveys, I still find that approximately 45 million Americans are spending more than an hour each week attending religious worship. And this doesn't include the large amount of time spent on religious devotion that occurs outside of places of worship (prayer, scripture study, etc.). In this paper, I attempt to better understand what religious worship attendance looks like by providing a descriptive analysis using a new data source, cellphone geodata.

As discussed throughout the paper, using cellphone data has limitations. One has to worry about the representativeness of the sample, whether phones are pinging consistently enough, and whether the cellphones are always tethered to their owner. It is also not possible to know if a cellphone's owner identifies with a religion. Instead, this has to be uncovered based on movement. Despite these limitations, the cellphone data are able to provide a new angle at understanding worship attendance. Hopefully this new data source allows researchers to not just better understand religious worship, but to better understand the connection of actual behavior and survey data. For example, why does there appear to be so much social desirability bias when it comes to religious worship frequency? When survey data suggest a drop in religious attendance over the last two decades, how much of this drop is actual attendance drop and how much is a change in social desirability (e.g. a greater willingness to admit to not being a regular worshipper)?

The detailed nature of these data can also help address causal issues related to religious observance. What happens to attendance when certain political or social news events happen? What happens when events occur at the local level (e.g. sex abuse scandals) that may impact religious worship.

Future research can hopefully build on the descriptive data presented in this paper to answer some of the most important causal questions related to religion such as what leads to increased/decreased religiosity and what are the impacts of religiosity on the attitudes and behaviors.

References

Ager, P., Hansen, C. W. and Lonstrop, L. (2014). Church Membership and Social Insurance: Evidence from the Great Mississippi Flood of 1927. *SSRN Electronic Journal,* Working Paper 2530868. https://ssrn.com/abstract=2530868

Ager, P. and Hansen, C. W. and Lonstrop, L. (2016). Church Membership and Social Insurance: Evidence from the Great Mississippi Flood of 1927. *SSRN Electronic Journal,* Working Paper 2862290. https://ssrn.com/abstract=2862290

American Time Use Survey, *US Bureau of Labor Statistics*

Anderson, B. A., Silver, B. D. and Abramson, P. R. (1986). Who Overreports Voting? *American Political Science Review,* 80(2), 613-624. DOI: 10.2307/1958277

Ansolabehere, S. and Hersh, E. (2012). Validation: What Big Data Reveal About Survey Misreporting and the Real Electorate. *Political Analysis,* 20(4), 437-459. DOI: 10.1093/pan/mps023

Balan, D. and Knack, S. (2012). The correlation between human capital and morality and its effect on economic performance: Theory and evidence. *Journal of Comparative Economics*, 40(3), 457-475. https://doi.org/10.1016/j.jce.2011.12.005

Bazzi, S., Hilmy, M. and Marx, B. (2020). Religion, Education, and the State. *NBER Working Papers*. DOI: 10.3386/w27073

Benjamin, D. J., Choi, J. J. and Fisher, G. W. (2016). Religious Identity and Economic Behavior. *The Review of Economics and Statistics,* 98(4), 617-637. https://doi.org/10.1162/REST_a_00586

Bentzen, J. S. (2019). Religiosity and Natural Disasters Across Subnational World Districts. *The Economic Journal*, 129(622). https://doi.org/10.1093/ej/uez008

Bernstein, R., Chadha, A. and Montjoy R. (2001). Overreporting Voting: Why It Happens and Why It Matters. *Public Opinion Quarterly*, 65(1), 22–44. https://doi.org/10.1086/320036

Bottan, N. L. and Perez-Truglia, R. (2015). Losing my religion: The effects of religious scandals on religious participation and charitable. *Journal of Public Economics*, 129, 106-119. https://doi.org/10.1016/j.jpubeco.2015.07.008

Brenner, P. S. (2011). Exceptional Behavior of Exceptional Identity? Overreporting of Church Attendance in the US. *Public Opinion Quarterly*, 75(1), 19-41. https://doi.org/10.1093/poq/nfq068

Brenner, P. S. and DeLamater, J. (2016). Lies Damned Lies, and Survey Self-Reports? Identity as a Cause of Measurement Bias. *Social Psychology Quarterly,* 79(4), 333-354. https://doi.org/10.1177/0190272516628298

Buser, T. (2015). The Effect of Income on Religiousness. *American Economic Journal: Applied Economics*, 7(3), 178-195. DOI: 10.1257/app.20140162

Chetty, R. et al. (2022). Social capital I: measurement and associations with economic mobility. *Nature*, 608, 108-121. https://doi.org/10.1038/s41586-022-04996-4

Dube, O., Blumenstock J. and Callen, M. (2022). Measuring Religion from Behavior: Climate Shocks and Religious Adherence in Afghanistan. *NBER Working Papers*. DOI: 10.3386/w30694

Gallup (2022). In Depth: Topics A to Z, Religion. https://news.gallup.com/poll/1690/religion.aspx

Giles T., Hungerman, D. M. and Oostrom, T. (2023). Opiates of the Masses? Deaths of Despair and the Decline of American Religion. *NBER Working Papers*. DOI: 10.3386/w30840

Gruber, J. (2004). Pay or pray? The impact of charitable subsidies on religious attendance. *Journal of Public Economics*, 88(12), 2635-2655.

Gruber, J. and Hungerman, D. M. (2008). The Church Versus the Mall: What Happens When Religion Faces Increased Secular Competition? *The Quarterly Journal of Economics*, 123(2), 831–862. https://doi.org/10.1162/qjec.2008.123.2.831

Hadaway, C. K., Marler, P. L., and Chaves, M. (1998). Overreporting church attendance in America: Evidence that demands the same verdict. *American Sociological Review*, 63(1), 122–130. https://doi.org/10.2307/2657484

Hungerman, D. M. (2005). Are church and state substitutes? Evidence from the 1996 welfare reform. *Journal of Public Economics,* 89(11-12), 2245-2267. https://doi.org/10.1016/j.jpubeco.2004.12.009.

Hungerman, D. M. (2014). The effect of education on religion: Evidence from compulsory schooling laws. *Journal of Economic Behavior and Organization*, 104, 52-63. https://doi.org/10.1016/j.jebo.2013.09.004

Hungerman, D. M. (2020). Religious Institutions and Economic Wellbeing. *The Future of Children.* 30(1), 9-28.

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

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

Levy, G. and Razin, R. (2012). Religious Beliefs, Religious Participation, and Cooperation. *American Economic Journal: Microeconomics*, 4(3), 121-151.

Lucchetti, G., Koenig, H. G., & Lucchetti, A. L. G. (2021). Spirituality, religiousness, and mental health: A review of the current scientific evidence. *World journal of clinical cases*, *9*(26), 7620–7631. https://doi.org/10.12998/wjcc.v9.i26.7620

McCleary, R. M. and Barro, R. J. (2003). Religion and Economic Growth across Countries. *American Sociological Review*, 68(5), 760-781. https://doi.org/10.2307/1519761

Moreno-Medina, J. (2023). Sinning in the Rain: Weather Shocks, Church Attendance, and Crime. The *Review of Economics and Statistics,* 105(1), 54-69. DOI: 10.1162/rest_a_01024

National Congregational Study Survey (year).

Climate Data Online. *National Oceanic and Atmospheric Administration.* https://www.ncei.noaa.gov/cdo-web/

Pew Research (2014). “Pew Research 2014 U.S. Religious Landscape Study.” https://www.pewresearch.org/religion/dataset/pew-research-center-2014-u-s-religious-landscape-study/

Sibley, C. G., & Bulbulia, J. (2012). Faith after an earthquake: a longitudinal study of religion and perceived health before and after the 2011 Christchurch New Zealand Earthquake. *PloS one*, 7(12), e49648. https://doi.org/10.1371/journal.pone.0049648

Appendix 1

There are 228,721 religious places of interest (POIs) in the 2019 Safegraph/Veraset data. This is certainly an undercount of total religious places of interest (estimates suggest closer to ~380k). In some ways, this is a standard missing data problem and one could impute data under the assumption that the missing religious POIs are missing at random from the full sample. However, there is no reason to think that the missing POIs were a random set of religious locations. It is possible that Safegraph was less likely to include smaller POIs or POIs from certain religious traditions.

Fortunately, I am able to identify the missing POIs in this situation. This is done by obtaining Safegraph's POI data from Fall 2023 (downloaded from a 3rd party data provider, Dewey). In 2023, there are 427, 275 POIs. Research assistants once again went through these POIs to categorize each location with a religious tradition (e.g. Protestant). After doing several random deep dives into the new POIs (e.g. using google maps street view), we concluded that a large group of the 2023 POIs did not appear to actually be places of worship, but rather were the homes of the director/owner of a non-profit association that had some relation to religion. Our process led us to eliminate approximately 35,000 POIs from the 2023 dataset leaving 392,220 religious locations in the 2023 data. This number is now quite close to the estimated 380k religious POIs in the US.

This allows us to identify what POIs were missing in the 2019 dataset. For example, we can calculate what percent of Catholic churches were missing in the 2019 dataset and this number will allow us to adjust the results from the 2019 analysis. However, we do not know the number of visits to each of the 2023 POIs that were missing in the 2019 POI dataset. And it is possible that these missing Catholic churches differ from the non-missing Catholic churches in a systematic way. For example, perhaps the Catholic churches that were missing in 2019 are smaller than those that were non-missing. However I can estimate how many visits each missing Catholic church was likely to have using the 2019 sample. Specifically, I regress the number of visits to a 2019 religious location on the size of the location (square meterage) and the population density of the county it is located in. These regressions are run separately by each religion, and then separately by each state[22]. Then for each location, I calculate fitted values separately by religion and then by state, which produces religion-level predicted visits and state-level predicted visits.

[22] These regressions are run separately by religion and then by state because the ultimate goal is to estimate total visits by religion, and then by state. To understand how many visits there might be in Vermont, there is no need to incorporate visits data from other states, which are less likely to be relevant. Similarly, when determining visits for a protestant church, it is unlikely that the number of visits to a Buddhist temple of a given size in a given county would be relevant.

These predicted visits are used to calculate religion and state-level weights, which are then used to update figures 1, 2, 3, 4, 12, 13, 14, 17, 18, and 19. Religion-level weights are calculated by summing up the religion-level predicted visits and calculating the percent change from the total religion-level visits in the 2019 sample to the total religion level predicted visits for the full sample (for each religion at a time). For Protestants, this ends up being .66, meaning that one would expect 66% more protestant visits in 2019 than are found in the 2019 visits data. The state-level weights are calculated in an analogous fashion but use state-level predicted visits and state-level totals for visits in 2019 rather than religion-level ones.

The religion-level weights are used to update figures 2, 4, and 19. The state level weights are used to update figures 12 and 13. Finally, the average of the religion level weights (which is very close to the protestant level weight) is used to update figures 1, 3, 14, 17, and 18 as well as tables 1 and 3. The religion level-weights vary between .26 and 1.54 mean (1.66) while the state level weights vary between .34 and 1.08 (mean 1.66). For figures that present attendance as a total (Figures 3, 4, 18, and 19) or that present the marginal effects of weather on attendance (figure 14), I simply multiply the visits obtained from the 2019 religious location sample by the corresponding weight. For example, if there were 20 million observed visits in the 2.1 million cellphone sample for 2019, and the weight is 1.2 (a 20% increase in visits), the weighted estimate for this figure would be 24 million visits. However, for figures that present religious attendance as a percent of Americans (figures 1, 2, and 17, and tables 1 and 3), the approach described above is problematic as it can produce an attendance share that is greater than 100%[23], which is unrealistic and not very informative.

This is particularly of concern for relatively higher initial attendance shares. To illustrate this, imagine that with the original sample of religious locations, it is found that 65% of Americans visit a religious location yearly using the visits data. With an average religion weight of .66 (a 66% increase), the upweighted share of Americans would be $.65 * (1 + .66) = 1.02$, or 102%.

On the other hand, for lower unweighted attendance shares, it seems reasonable for the share to increase roughly in equal proportion to the weight. For example, if the weight implies a doubling of visits ($w = 1$), then an increase from 5% to 10% very well could be justified. Therefore, it is critical that the application of weights treats lower percentages differently from higher percentages. Another way to think about this is that if there already were 60% of people attending

[23] In practice the weight would be applied as such: let the weight = w, and the initial share = p. Then the final share = $p * (1+w)$. Since $w > 0$, $1 + w > 1$ for all w.

church in the sample of 228k churches, adding more religious locations will almost surely have less of an effect on visits than if one began with 5% attendance in the visits data. Therefore, if weighted attendance shares are a function of the attendance share found in the 2019 visits data (holding the weight constant) then one would expect this function to be concave. Rather than impose a continuous functional form, the weights are applied according to the following stepwise function:

$$f(p,w) = \begin{cases} p \geq .25 \mid p * (1+w) \\ .25 < p \leq .50 \mid (.25 * (1+w)) + \left((p - .25) * \left(1+\frac{w}{2}\right)\right) \\ .50 < p \leq 1 \mid (.25 * (1+w)) + \left(.25 * \left(1+\frac{w}{2}\right)\right) + \left((p - .5) * \left(1+\frac{w}{4}\right)\right) \\ f(p,w) > 100 \mid f(p,w) = 100 \end{cases}$$

where *p* is the unweighted attendance share, *w* is the weight and *f(p,w)* is the weighted attendance share. For example, if *p* = .05 and *w* = 1 (indicating a 100% increase), then the weighted attendance share would be 10%. For *p* = .80, and *w* = 1, the final weighted share is 125%, which is actually then converted to 100%. Here, the first 25% is applied the full weight, the second 25% is applied half of the weight (a 50% increase instead of 100%), and the final 30% (80%-50% = 30%) is applied one quarter of the weight (a 25% increase). As can be seen here, the first 25 percent is given greater upweighting than the subsequent 25 percent, which has a greater upweighting than the final 50 percent. With this parameterization, *f(w,p)* is bounded in the interval [0,1][24].

[24] Note that if *p* = 0, then *f(p,w)* = 0 for all values of w. However, if p = 0 in the data, that would mean that there were no religious visits at all within a particular state or religion. Although this scenario is theoretically possible, it is practically impossible, and is not observed in the data.

Figure 1: Religious Attendance Based on Self-reported vs. Cellphone Sample

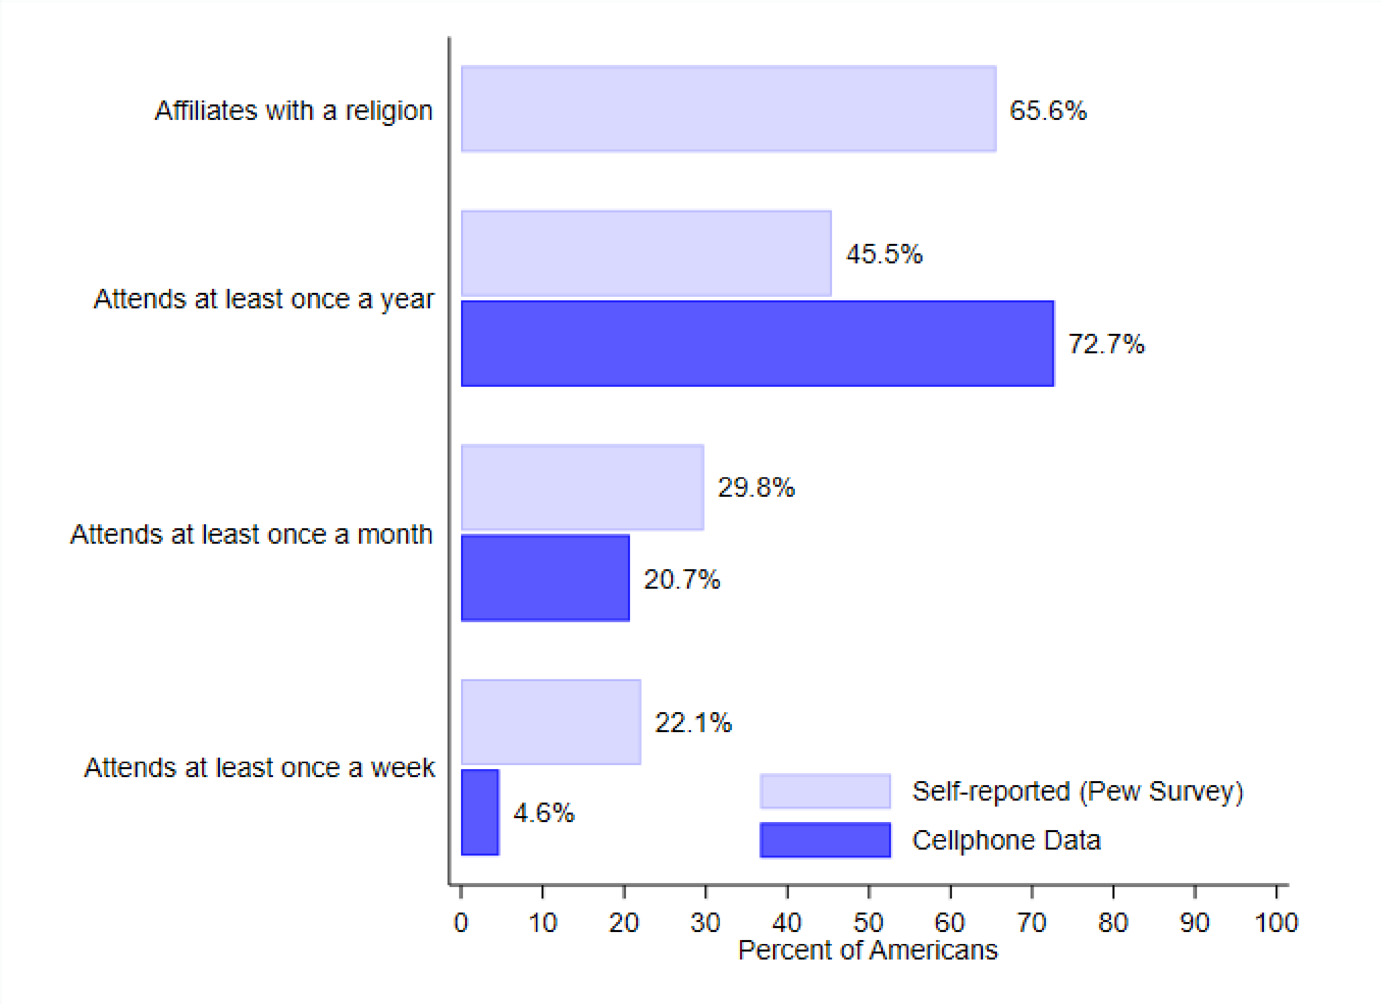

Notes: The light bars show self-reported religious attendance and religious affiliation among Americans based on data from Pew Research. The dark bars show religious attendance using the main sample of cellphone data, which only considers visits to places of worship during a given religion's primary day(s) of worship.

Figure 2: Religious Attendance by Religion

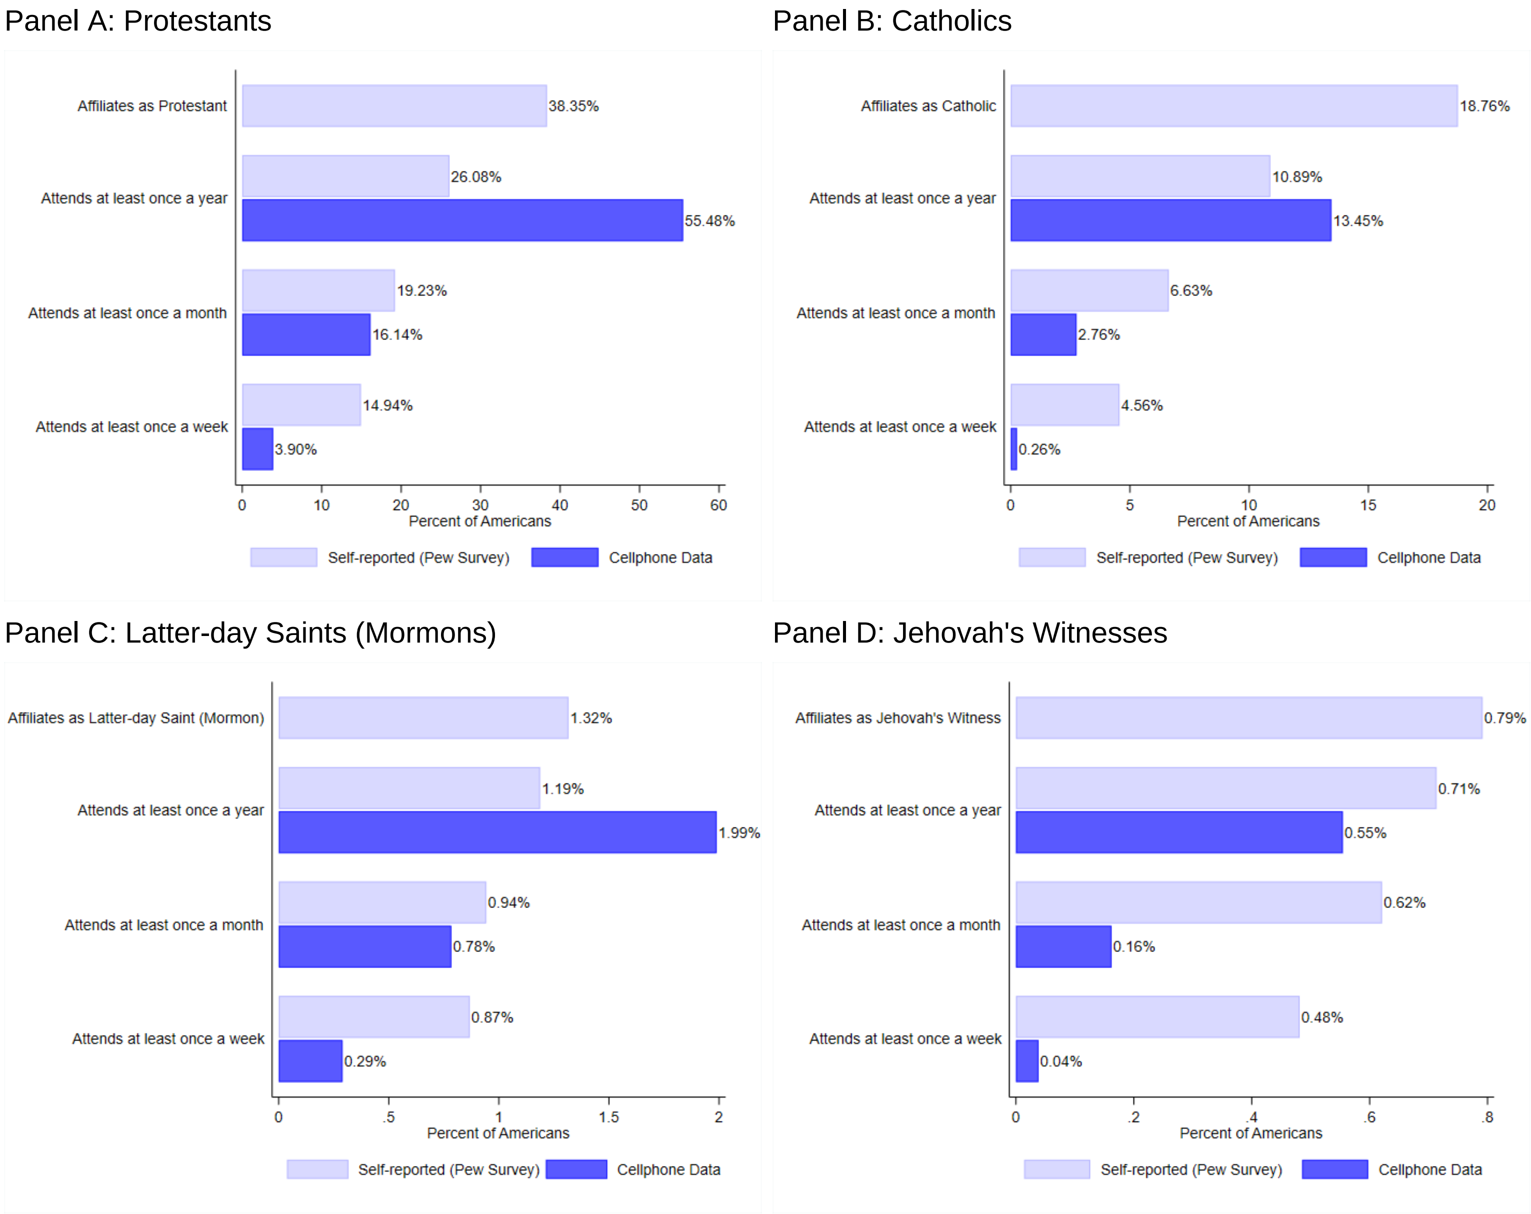

Notes: The light bars show self-reported religious attendance and religious affiliation among Americans based on data from Pew Research. The dark bars show religious attendance using the main sample of cellphone data, which only considers visits to places of worship during a given religion's primary day(s) of worship. Each panel presents religious attendance and affiliation for a given religion as a percent of the American population in 2019.

Figure 2 (Continued): Religious Attendance by Religion

Panel E: Orthodox Christians

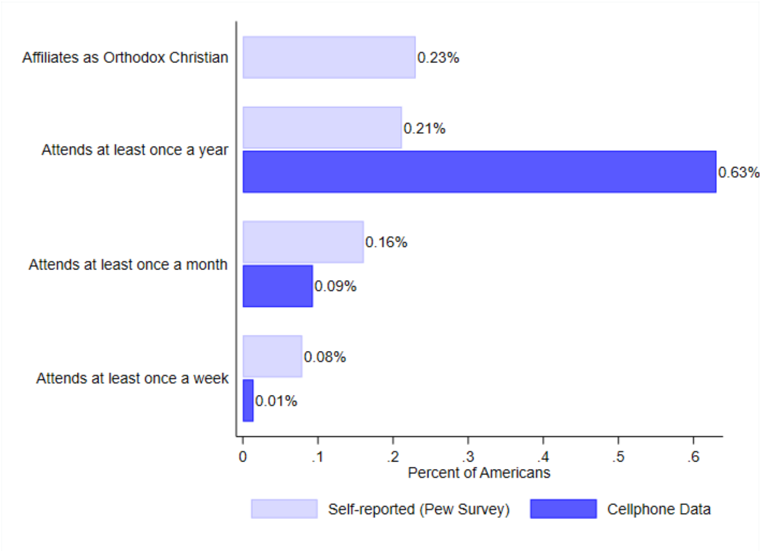

Panel F: Jews

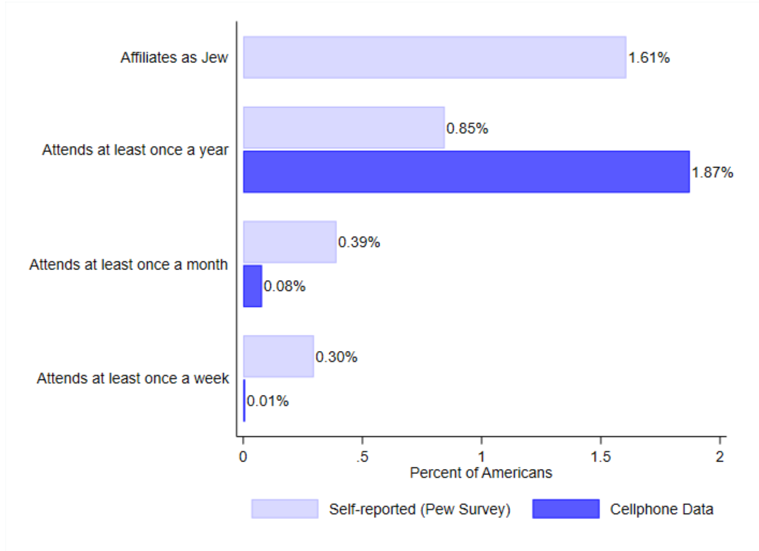

Panel G: Muslims

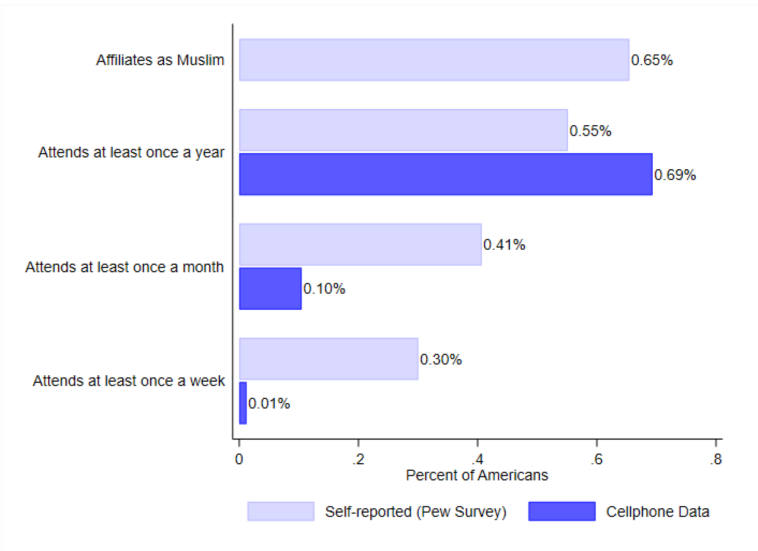

Panel H: Buddhists

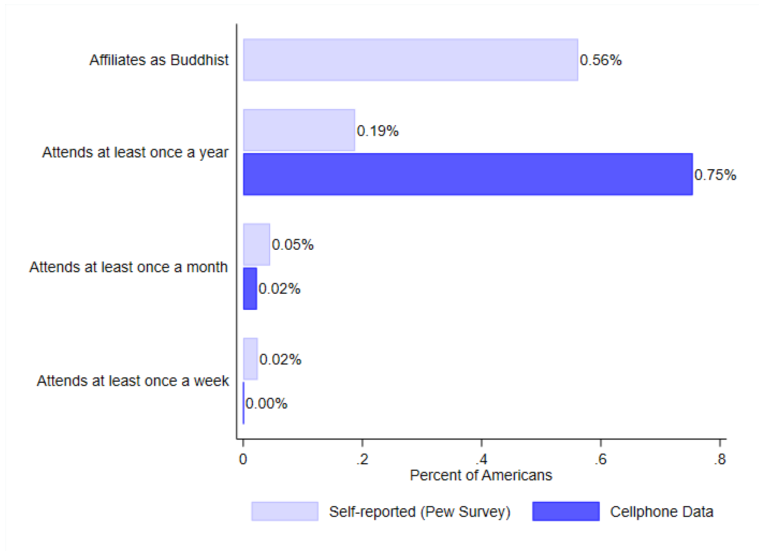

Notes: The light bars show self-reported religious attendance and religious affiliation among Americans based on data from Pew Research. The dark bars show religious attendance using the main sample of cellphone data, which only considers visits to places of worship during a given religion's primary day(s) of worship. Each panel presents religious attendance and affiliation for a given religion as a percent of the American population in 2019.

Figure 3: Week-to Week Variation in Religious Attendance

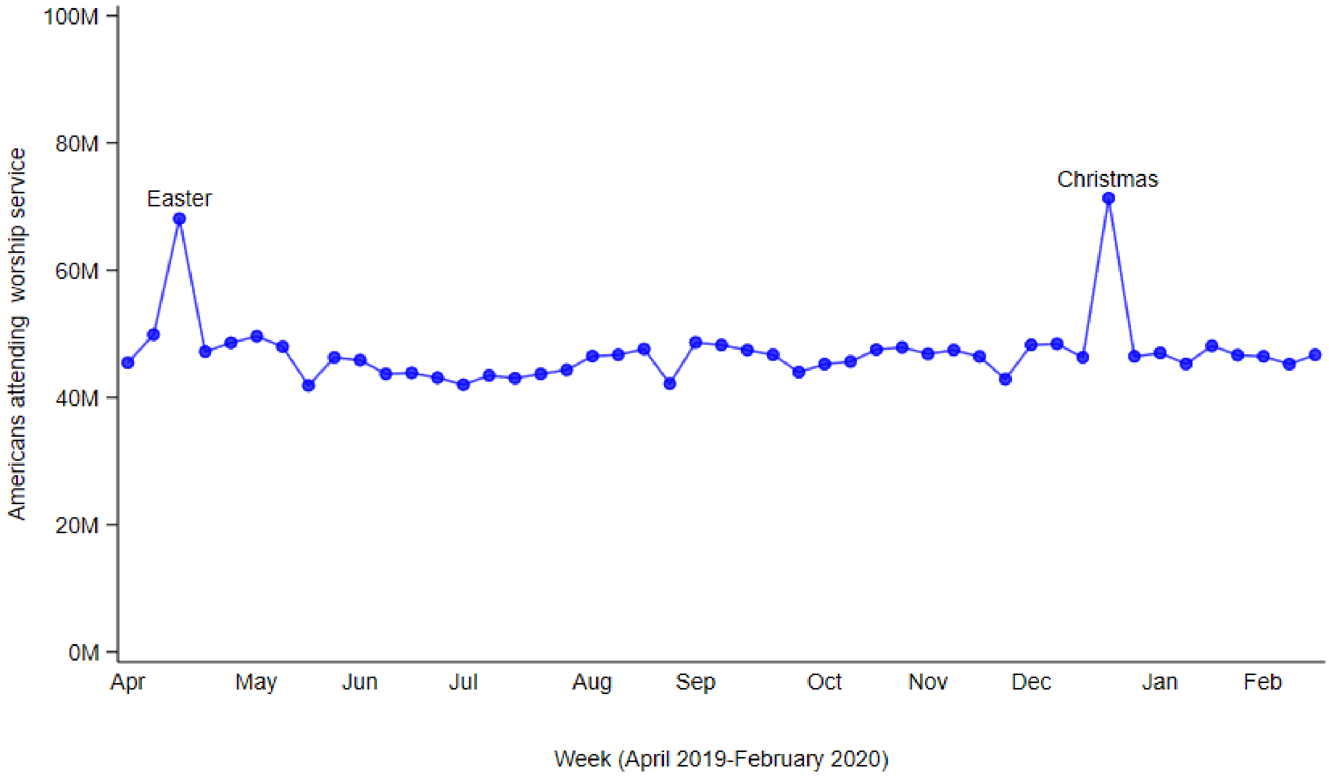

Notes: Each dot represents the weekly total of Americans attending a religious worship service during their religion's primary day(s) of worship each week during the cellphone sample period.

Figure 4: Week-to Week Variation by Religion

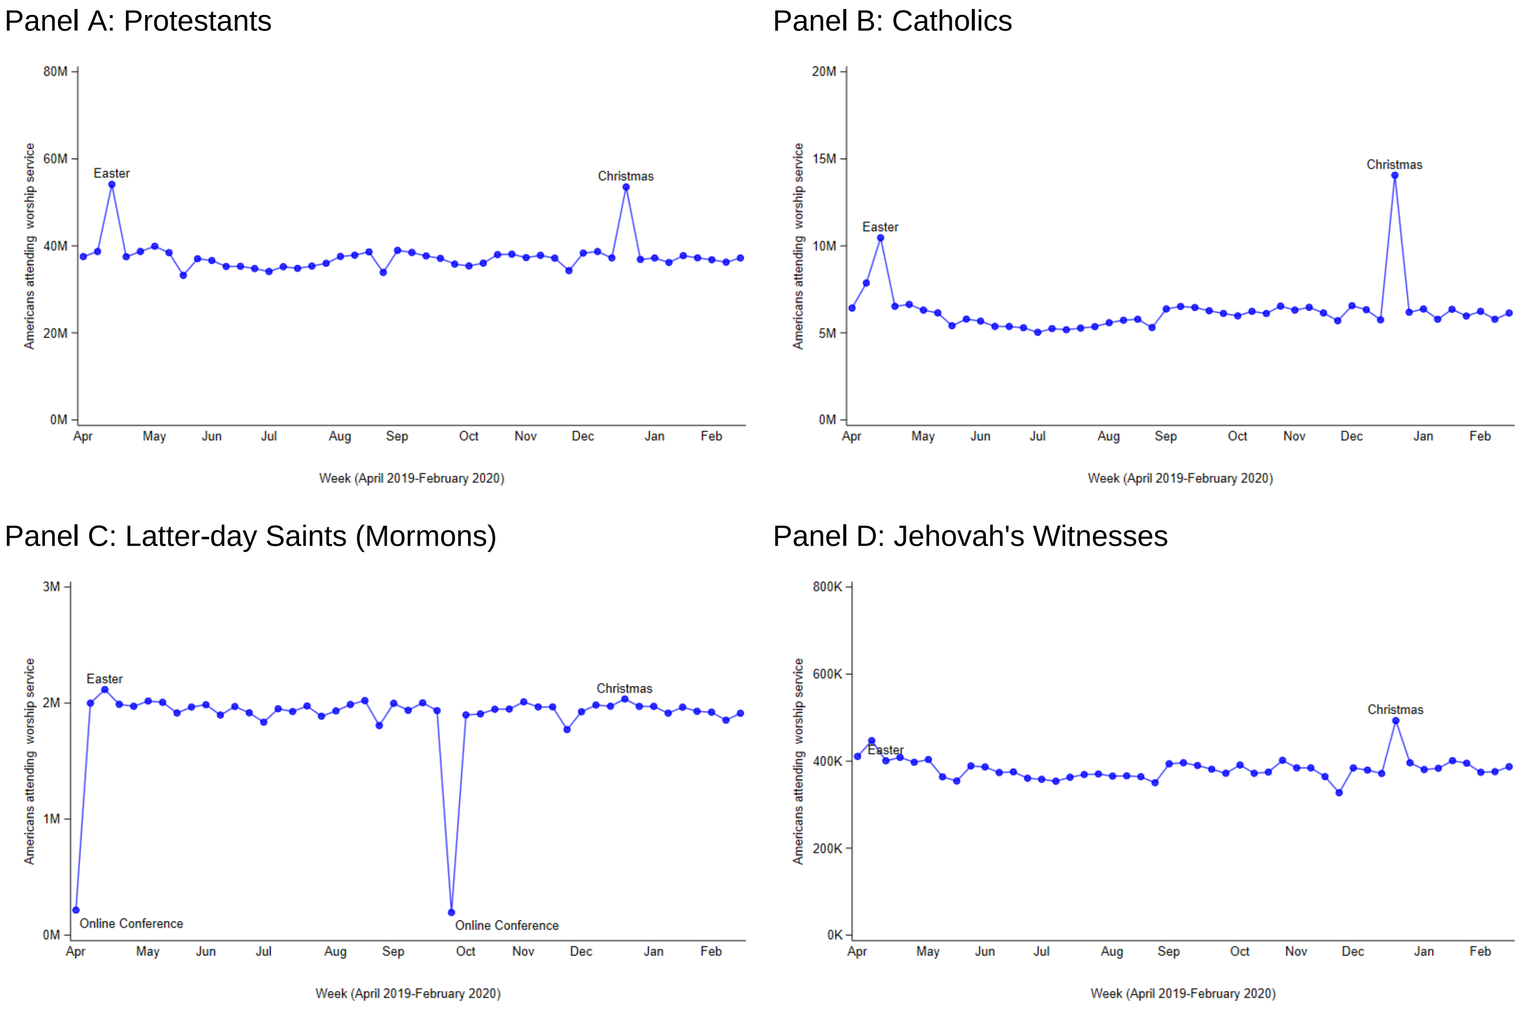

Notes: Each dot represents the total number of Americans attending a religious worship service during their religion's primary day(s) of worship from April 2019 to February 2020 using the main cellphone sample. Each panel presents this weekly total for a given religion. Important religious holidays and other events that one might expect to significantly impact religious worship attendance in a given week are indicated by text (and shaded area if the holiday occurs over multiple days e.g. Ramadan). The y-axis is scaled by either millions (M) or thousands (K) depending on the size of that religious constituency in the United States.

Figure 4 (Continued): Week-to Week Variation by Religion

Panel E: Orthodox Christians

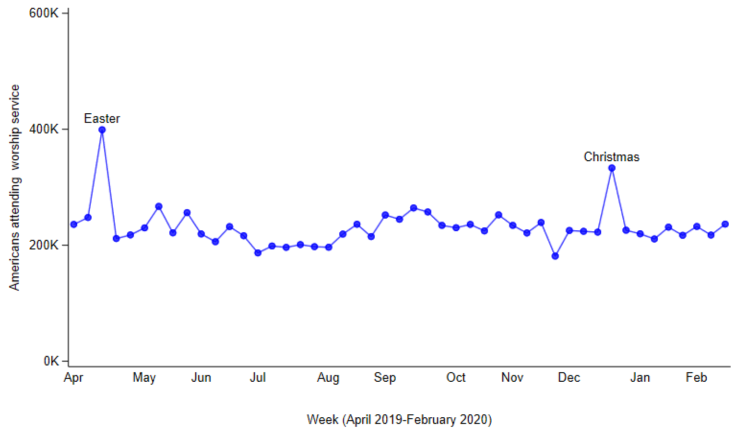

Panel F: Jews

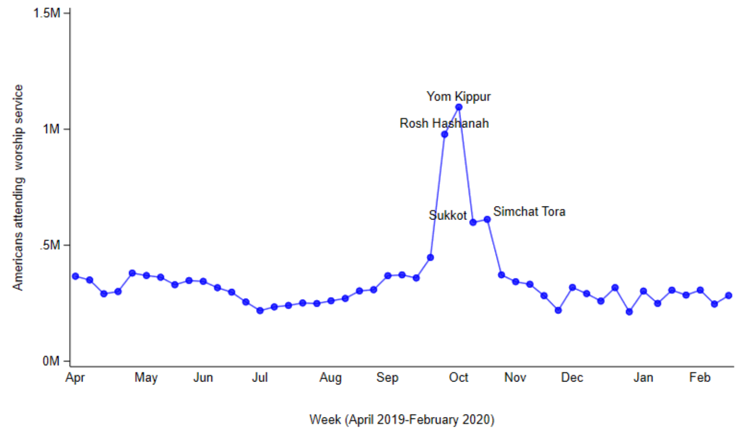

Panel G: Muslims

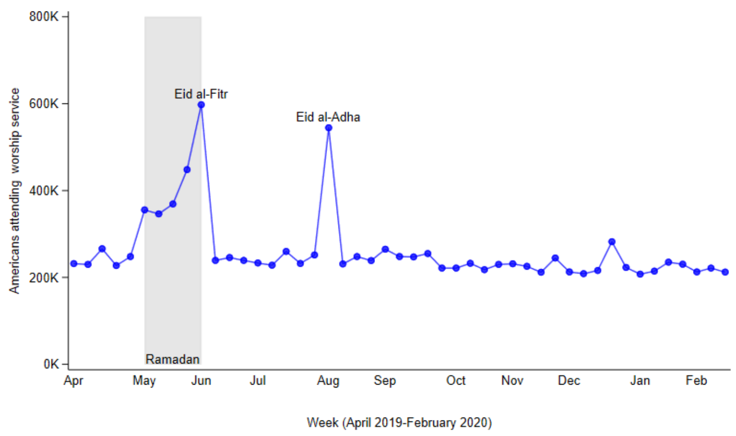

Panel H: Buddhists

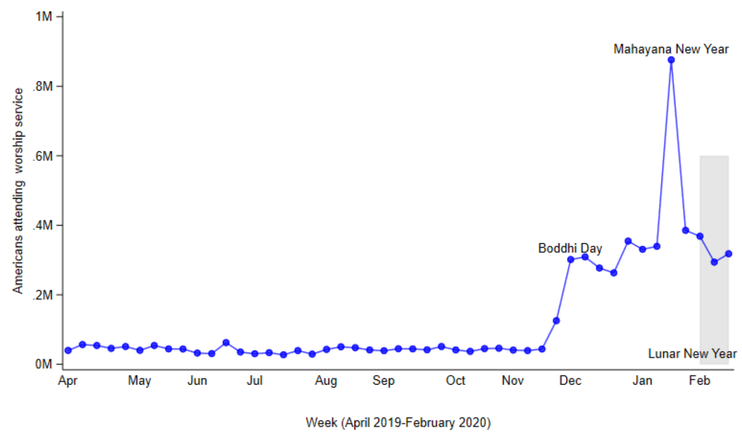

Notes: Each dot represents the total number of Americans attending a religious worship service during their religion's primary day(s) of worship from April 2019 to February 2020 using the main cellphone sample. Each panel presents this weekly total for a given religion. Important religious holidays and other events that one might expect to significantly impact religious worship attendance in a given week are indicated by text (and shaded area if the holiday occurs over multiple days e.g. Ramadan). The y-axis is scaled by either millions (M) or thousands (K) depending on the size of that religious constituency in the United States.

Figure 5: Worship Arrival Time by Religion

Panel A: Protestants

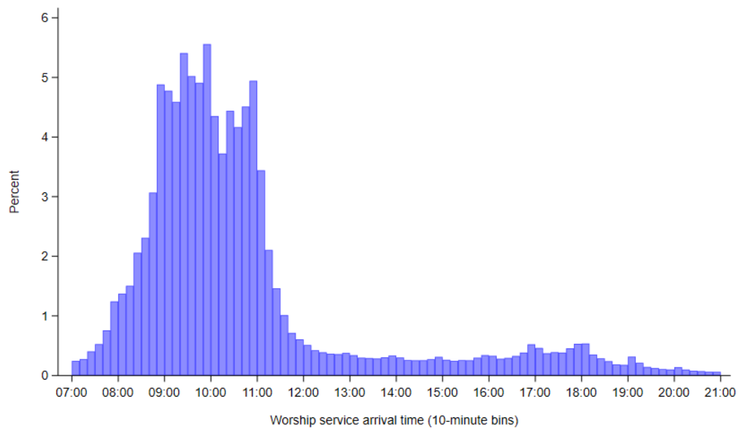

Panel B: Catholics

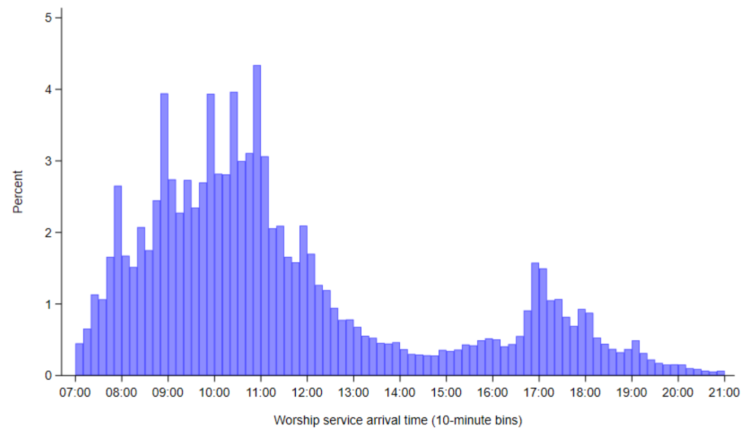

Panel C: Latter-day Saints (Mormons)

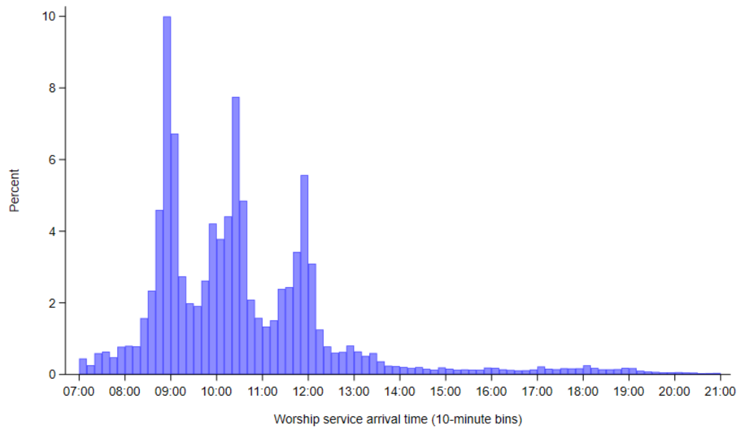

Panel D: Jehovah's Witnesses

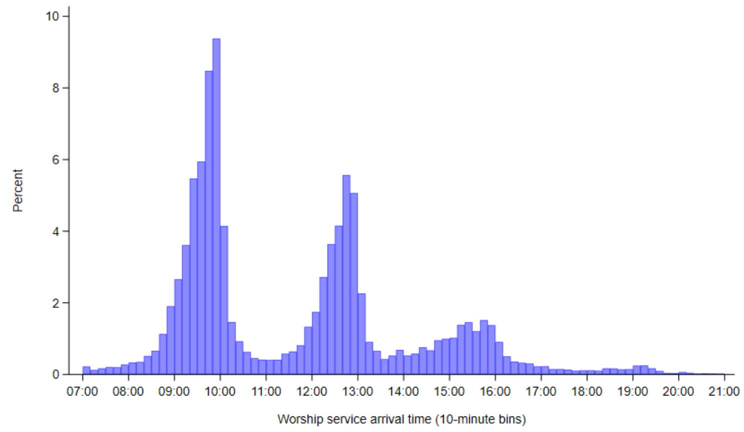

Notes: Each bar represents the percent of people within a given religion that arrive at a worship service for their respective religion within a given 10-minute time block.

Figure 5 (Continued): Worship Arrival Time by Religion

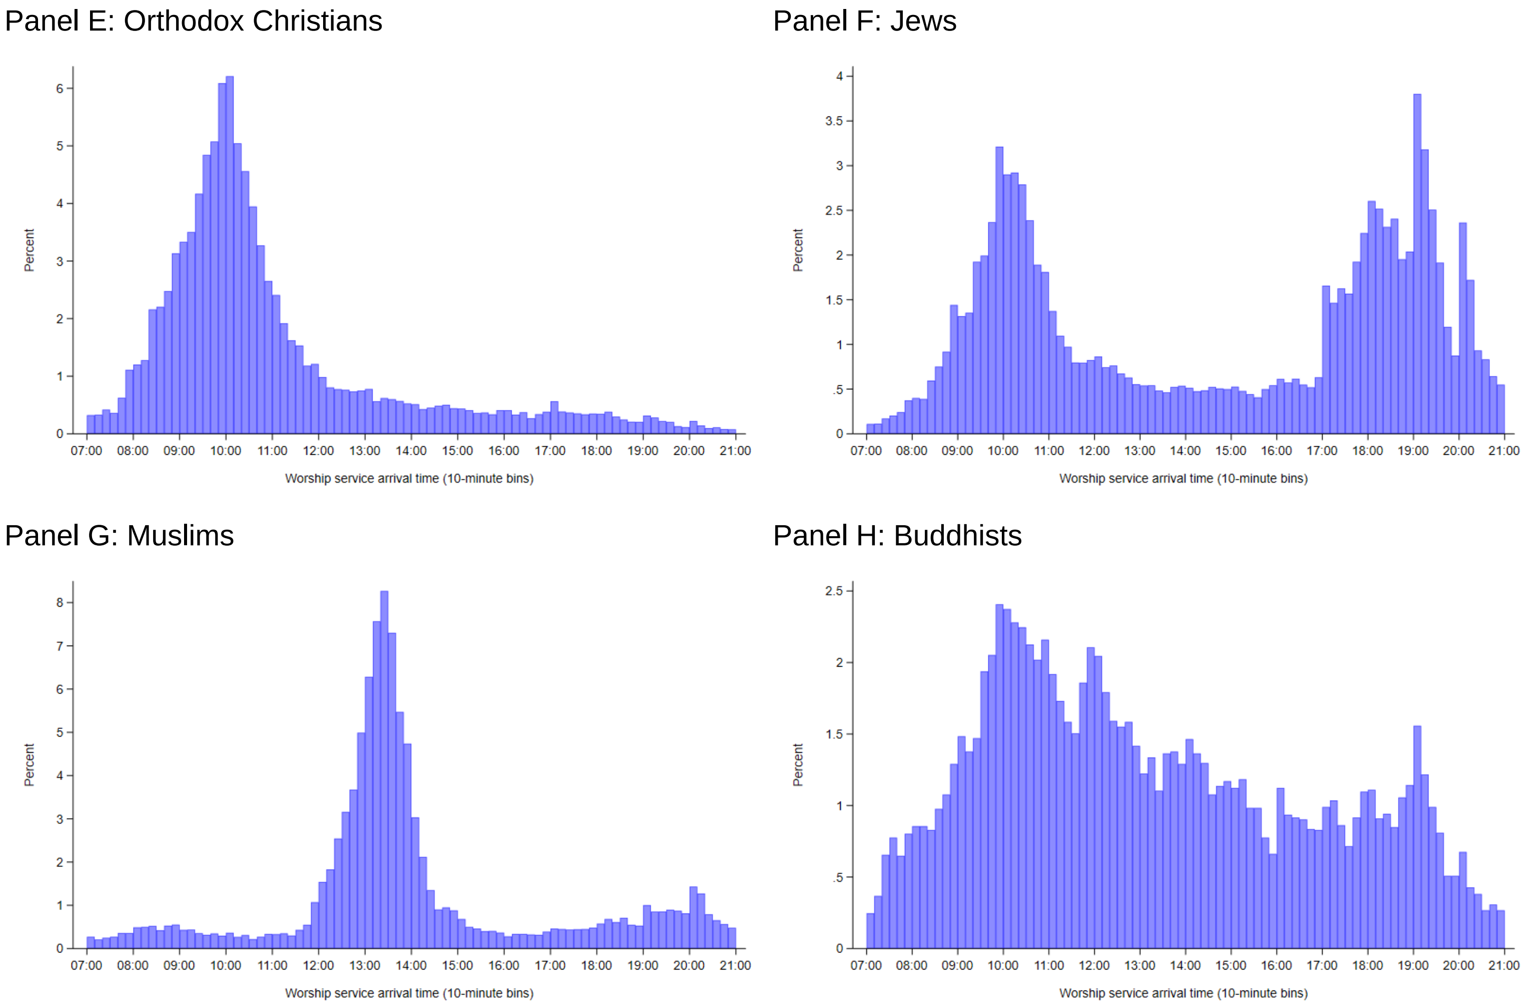

Notes: Each bar represents the percent of people within a given religion that arrive at a worship service for their respective religion within a given 10-minute time block.

Figure 6: Worship Duration by Religion

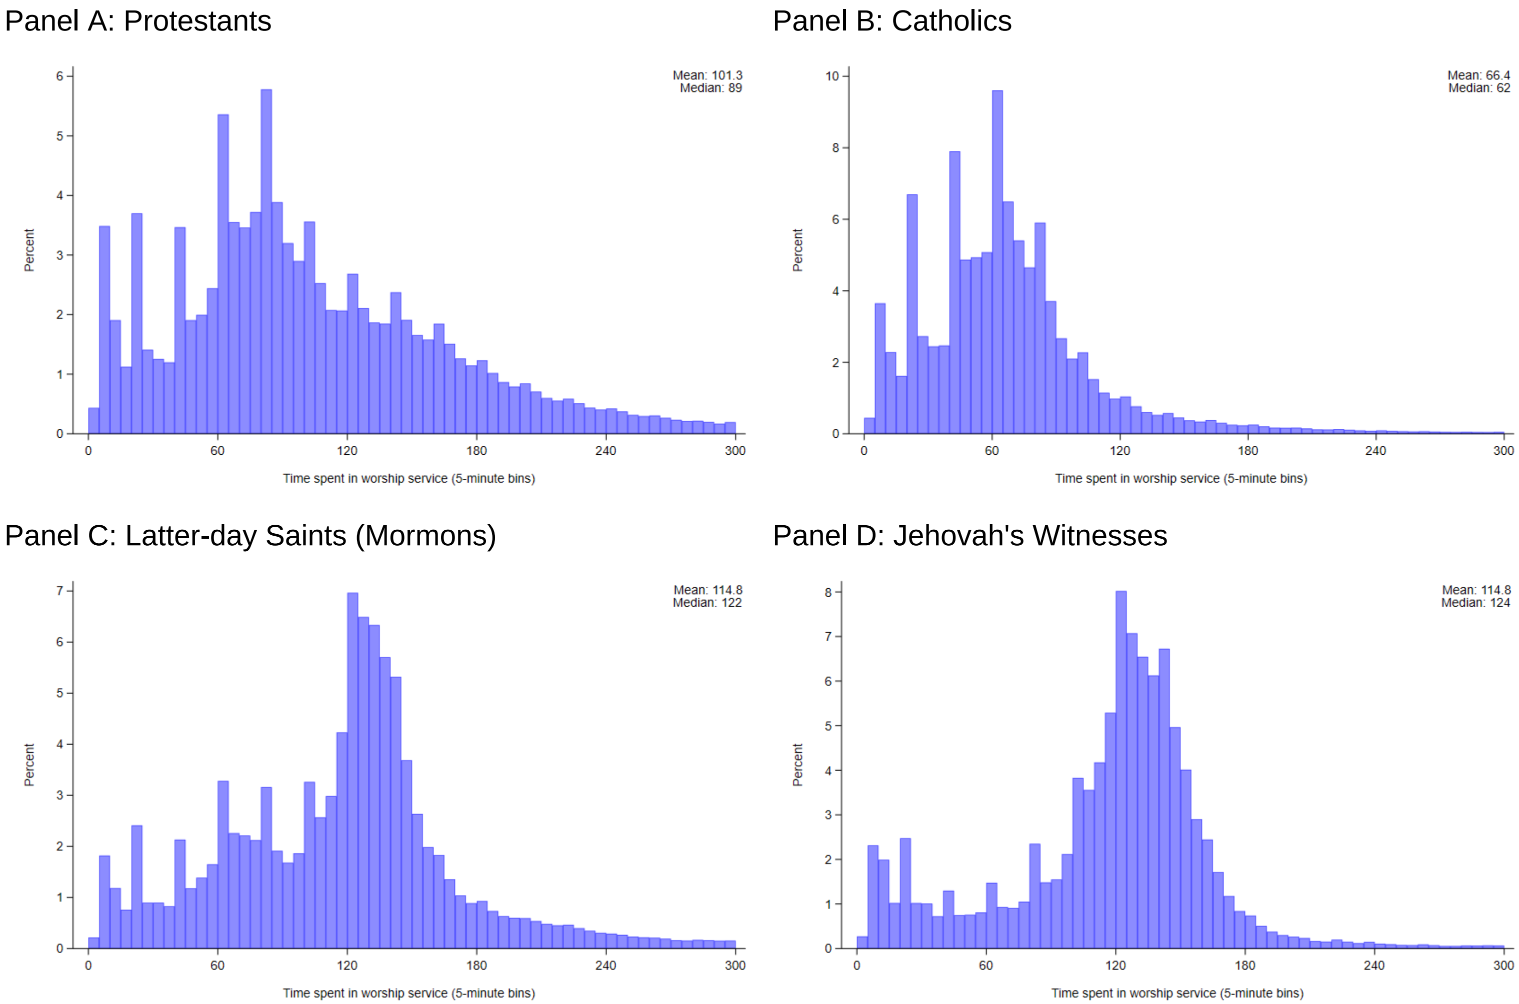

Notes: Each bar represents the percent of people in a given religion whose time spent at a wosrhip service falls within a given 5-minute bin of time.

Figure 6 (Continued): Worship Duration by Religion

Panel E: Orthodox Christians

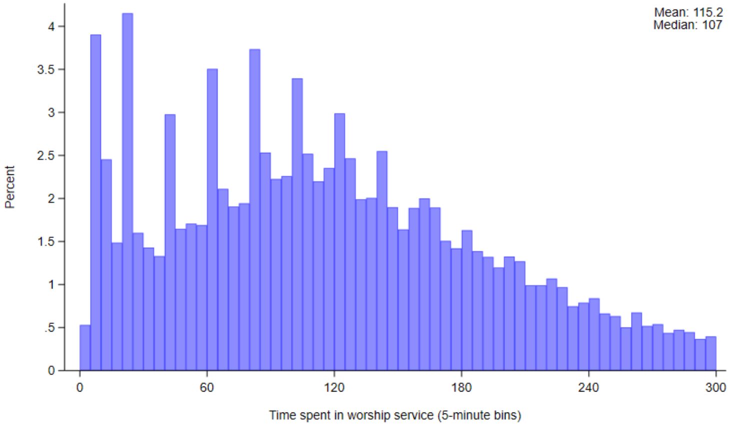

Panel F: Jews

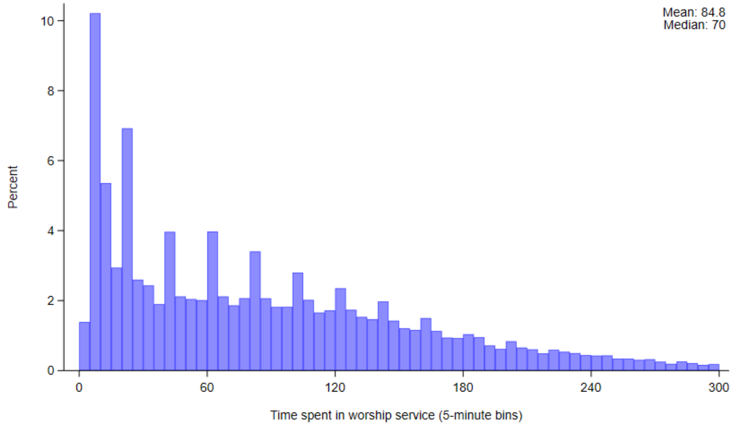

Panel G: Muslims

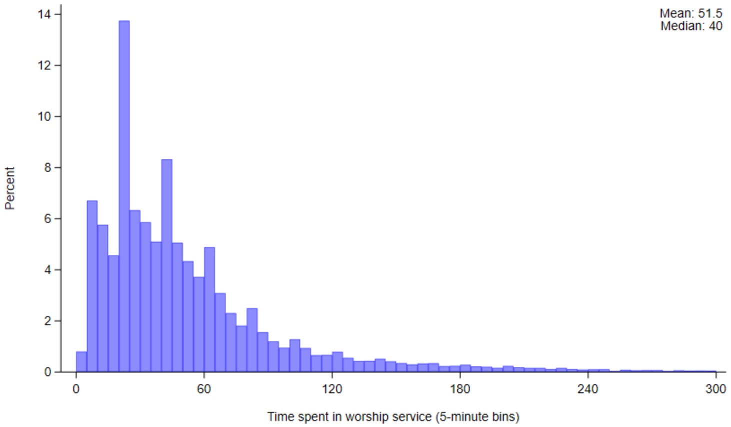

Panel H: Buddhists

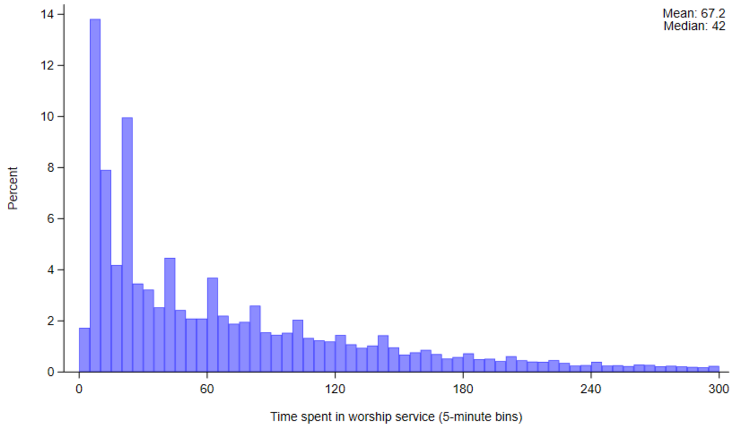

Notes: Each bar represents the percent of people in a given religion whose time spent at a wosrhip service falls within a given 5-minute bin of time.

Figure 7: Income Differences by Religious Attendance

Panel A: By Attendance

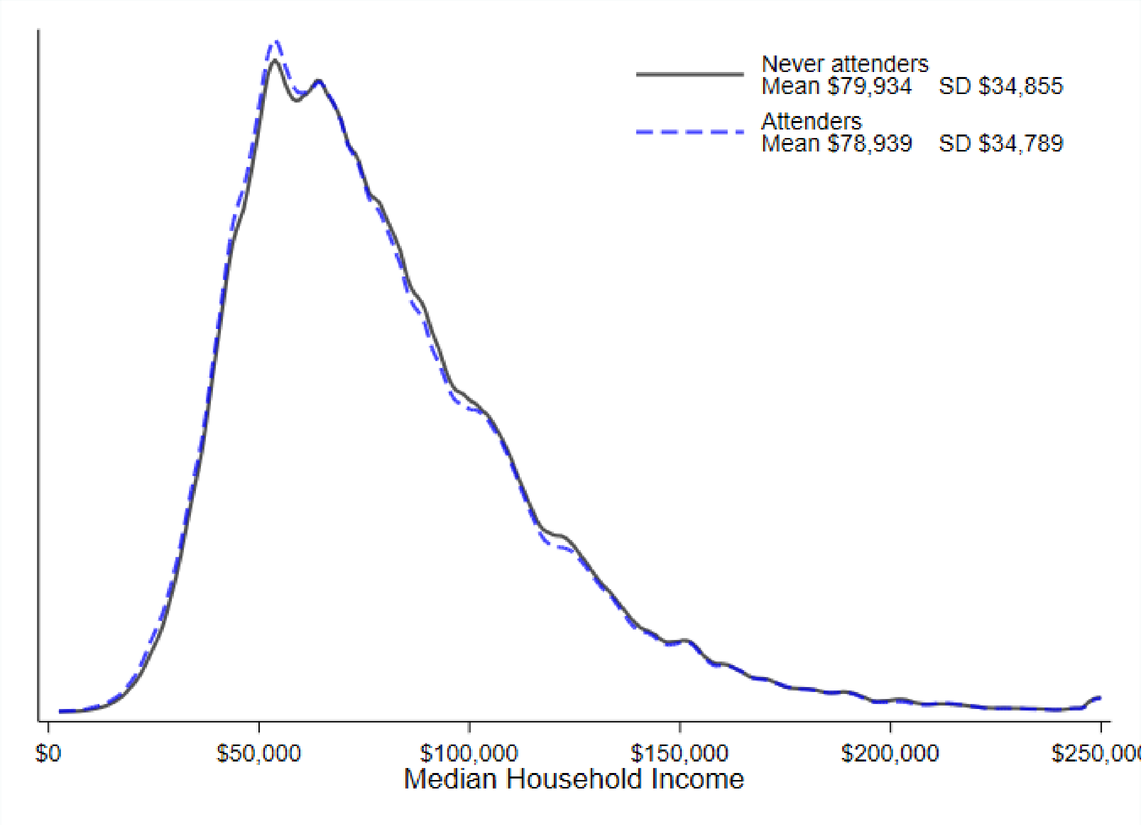

Panel B: By Frequency of Attendance

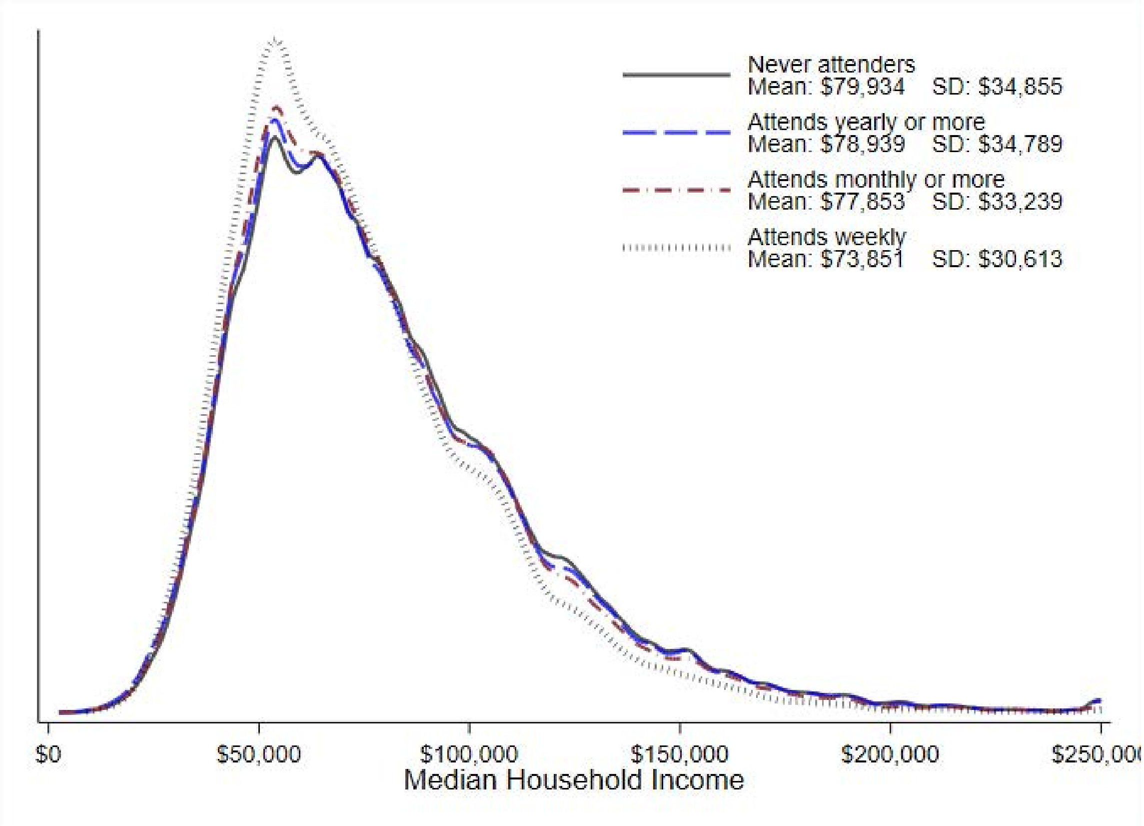

Notes: Panel A plots the distributions of median household income between those who attended a religious worship service at least once in the data and those who never did during the 47 week cellphone sample period. Panel B plots the distributions of median household income by the frequency of religious attendance, splitting the data into four mutually exclusive groups. Never attenders (0 religious worship visits during the sample period), yearly attenders (at least one religious worship visit during the sample period), monthly attenders (at least 11 religious worship visits during the sample period), and weekly attenders (at least 36 religious worship visits during the sample period).

Figure 8: Income Differences by Religion

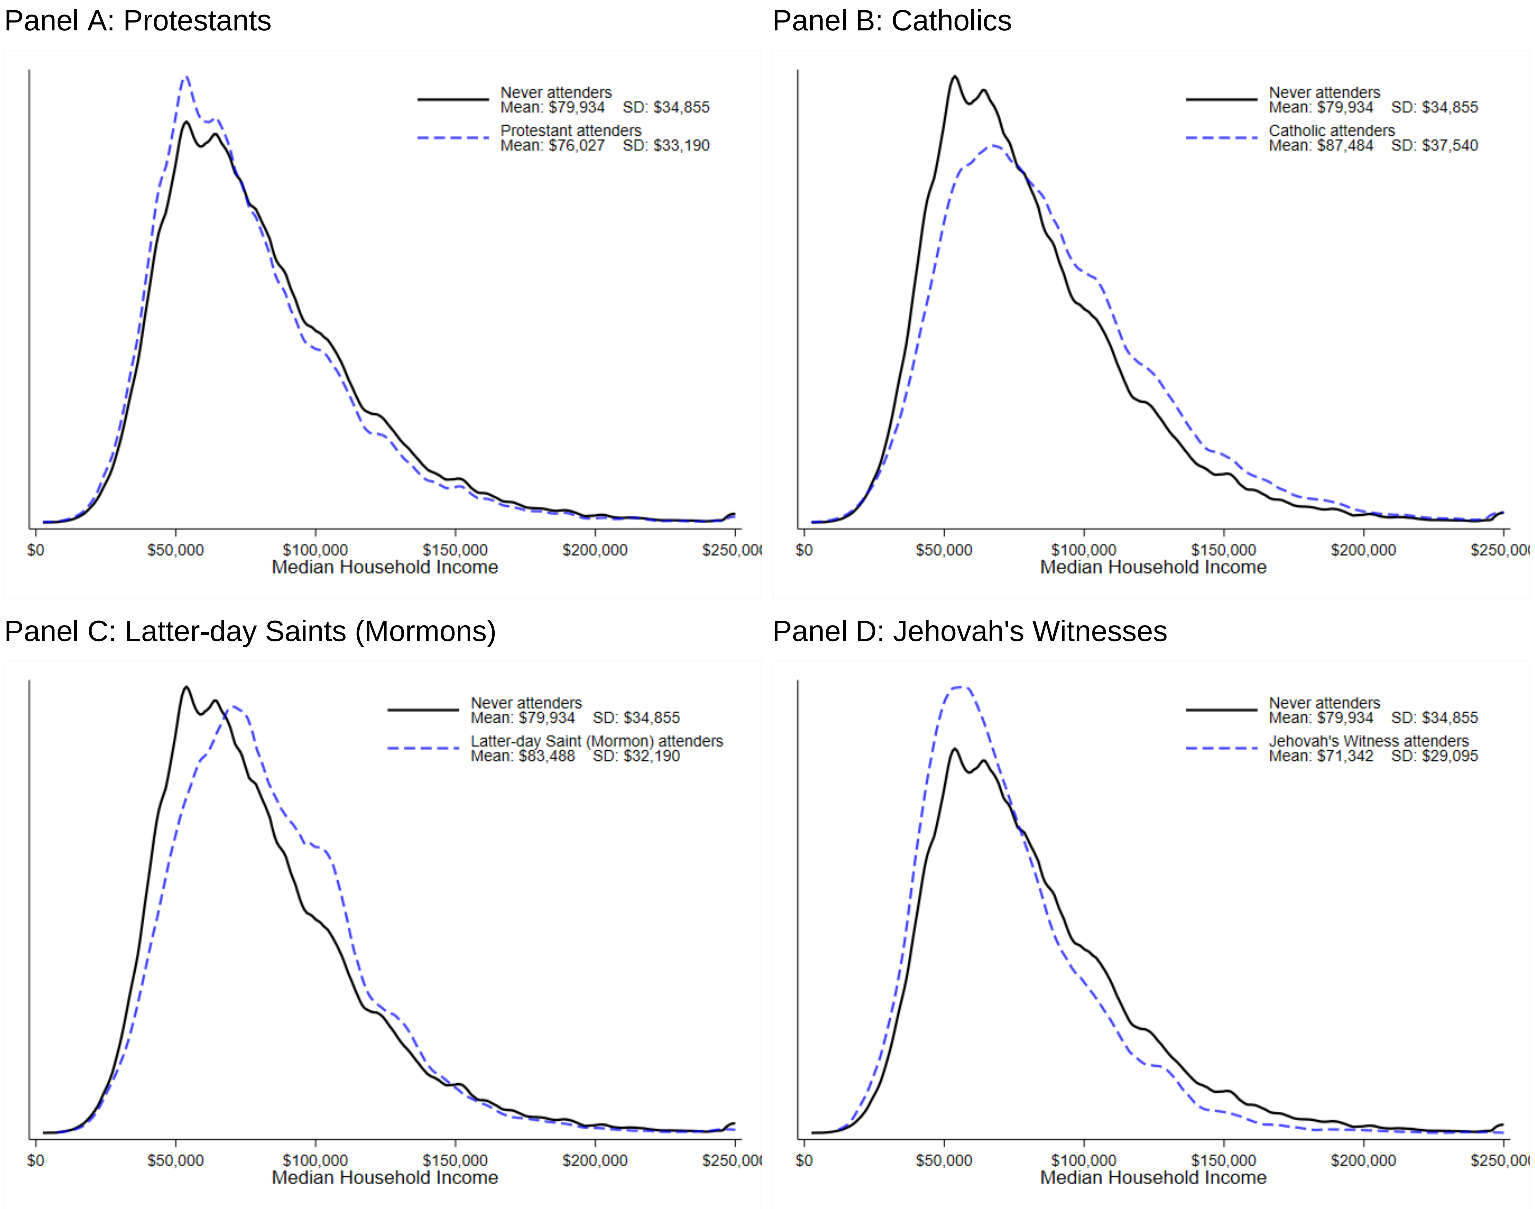

Notes: Each panel seperately plots the distributions of median household income among never attenders (0 religious worship visits during the cellphone sample period) and among those who attended at least one worship service for a given religion.

Figure 8 (Continued): Income Differences by Religion

Panel E: Orthodox Christians

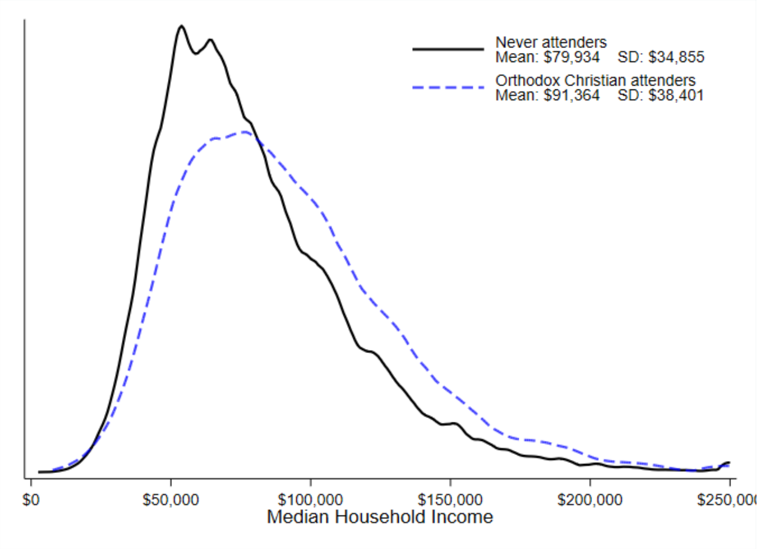

Panel F: Jews

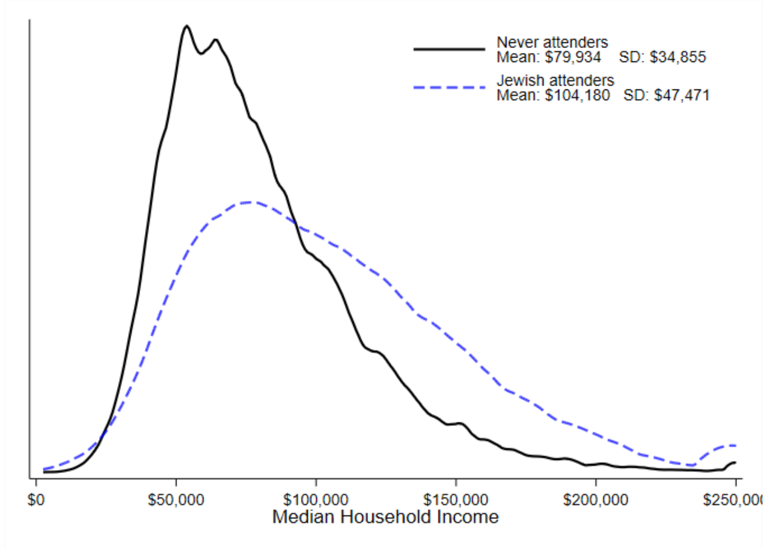

Panel G: Muslims

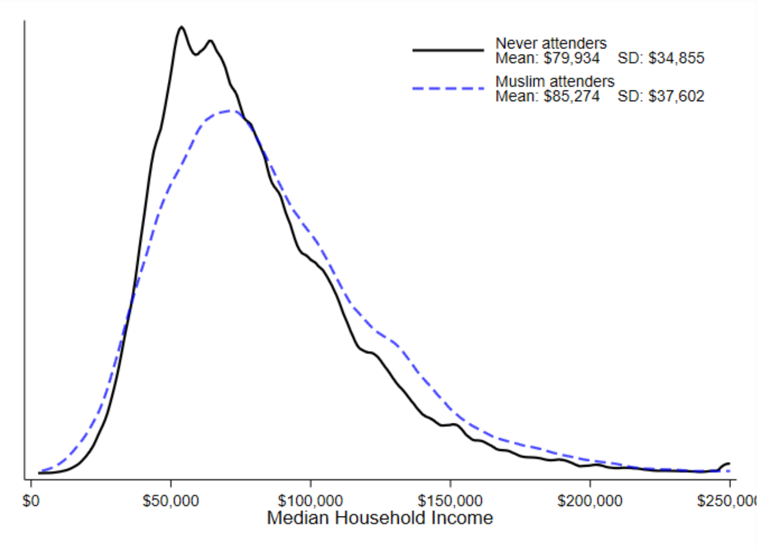

Panel H: Buddhists

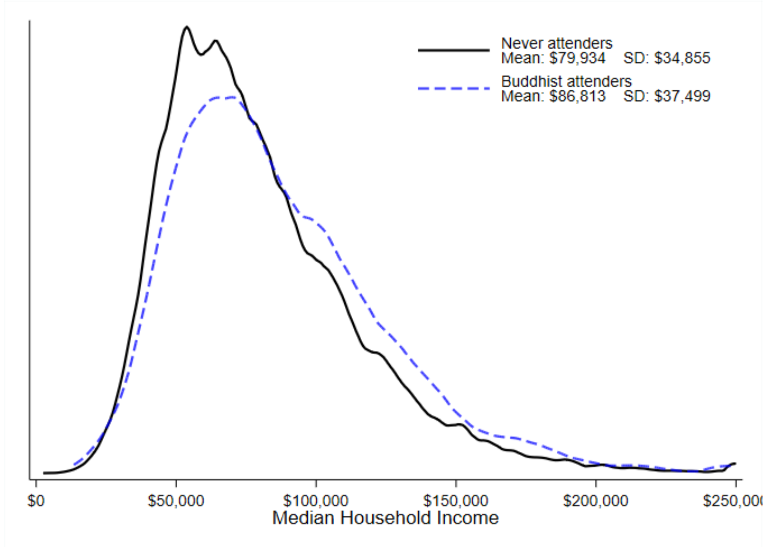

Notes: Each panel seperately plots the distributions of median household income among never attenders (0 religious worship visits during the cellphone sample period) and among those who attended at least one worship service for a given religion.

Figure 9: Income Difference for Co-worshippers

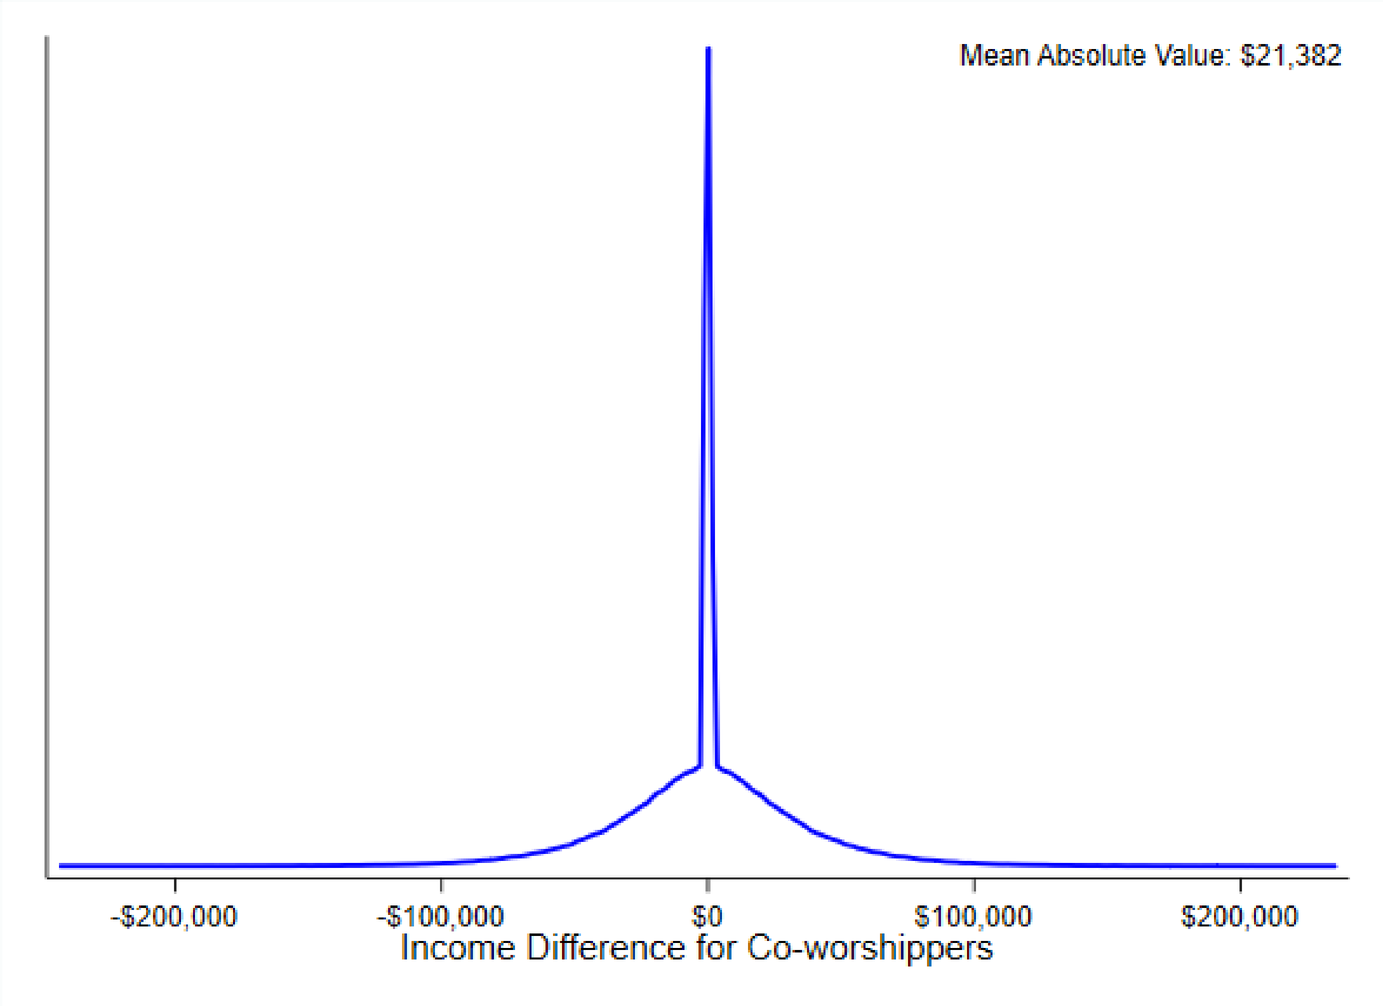

Notes: Income differences are calculated by randomly selecting pairs of co-worshippers (defined as two people that visit the same congregation within the same hour and day), and subtracting one co-worshippers income from the other's.

Figure 10: Income Difference for Co-worshippers by Religion

Panel A: Protestants

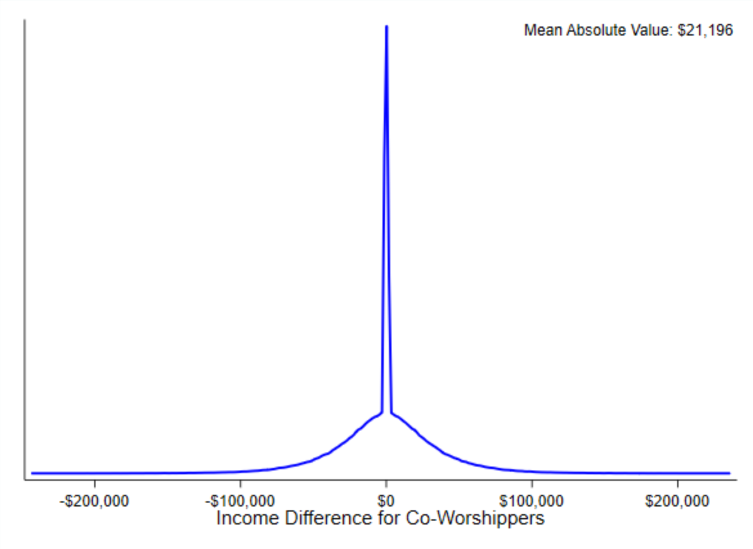

Panel B: Catholics

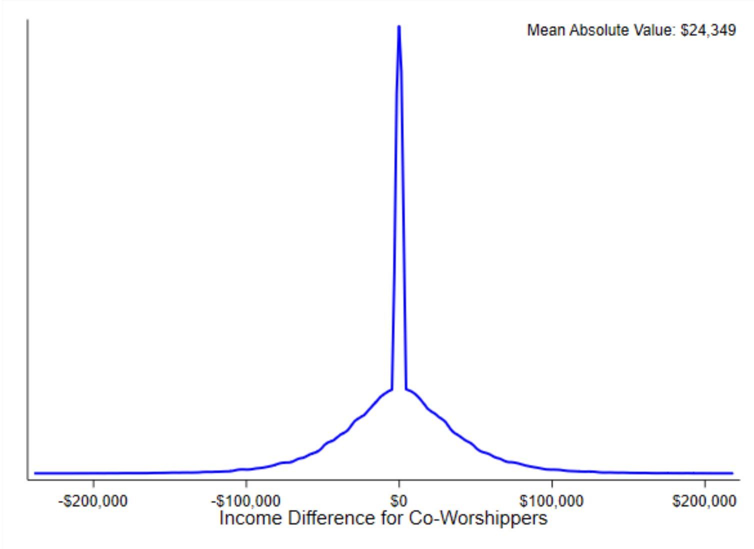

Panel C: Latter-day Saints (Mormons)

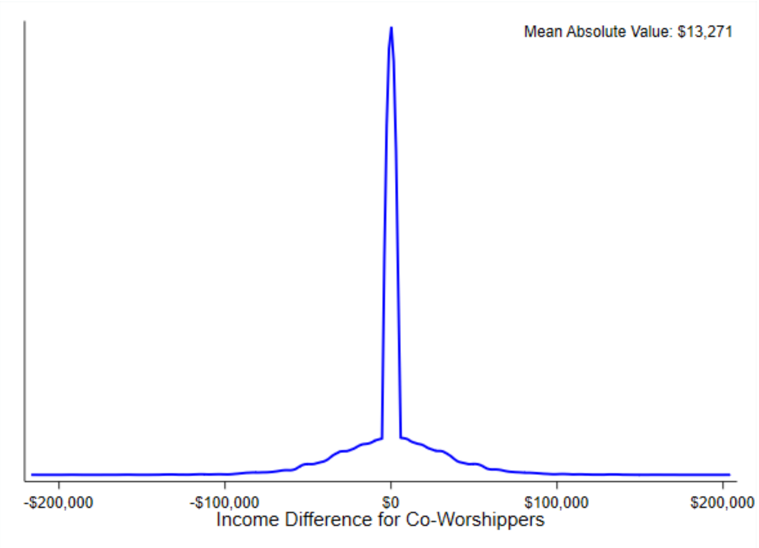

Panel D: Jehovah's Witnesses

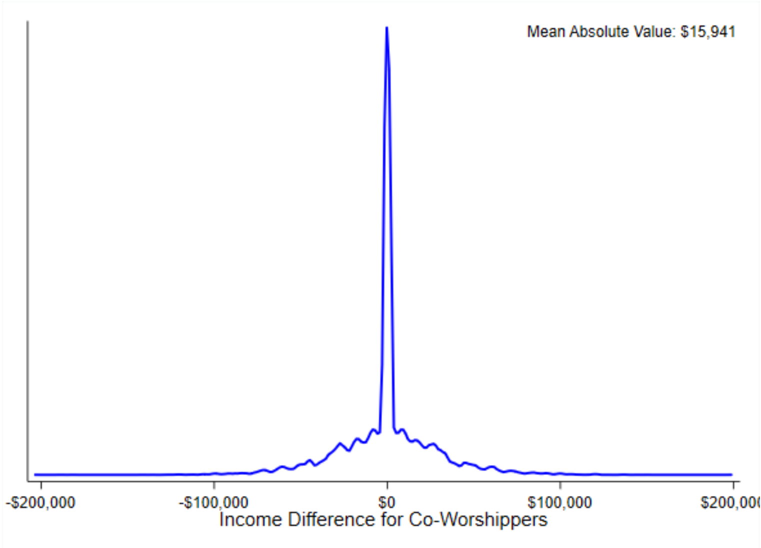

Notes: Each panel presents the distribution of income difference among co-worshippers (defined as two people that visited the same congregation within the same hour and day) for a given religion. Income differences are calculated by randomly selecting pairs of co-worshippers and subtracting one co-worshippers income from the other's.

Figure 10 (Continued): Income Difference for Co-worshippers by Religion

Panel E: Orthodox Christians

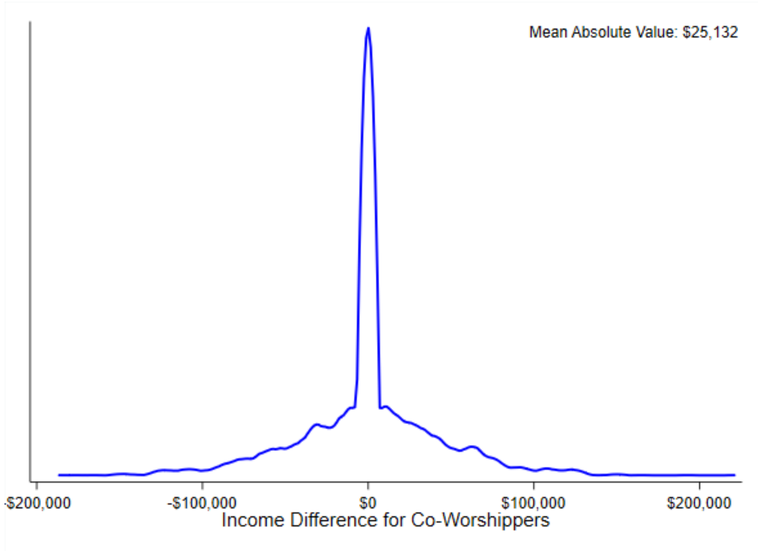

Panel F: Jews

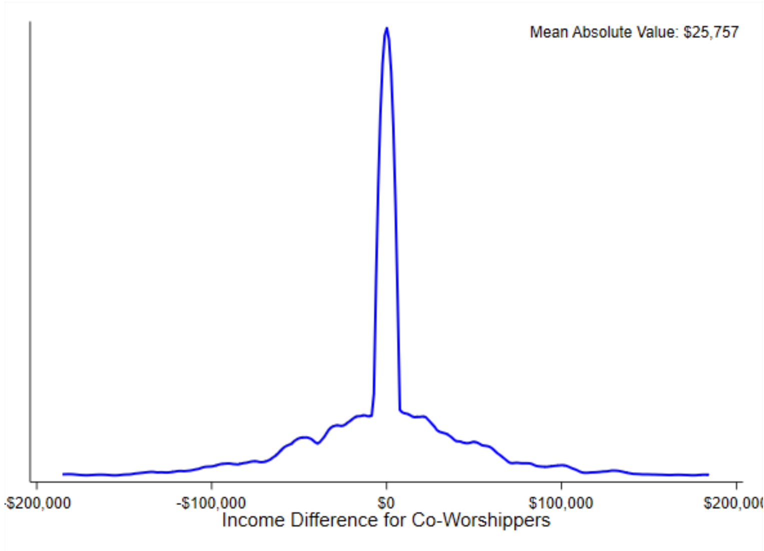

Panel G: Muslims

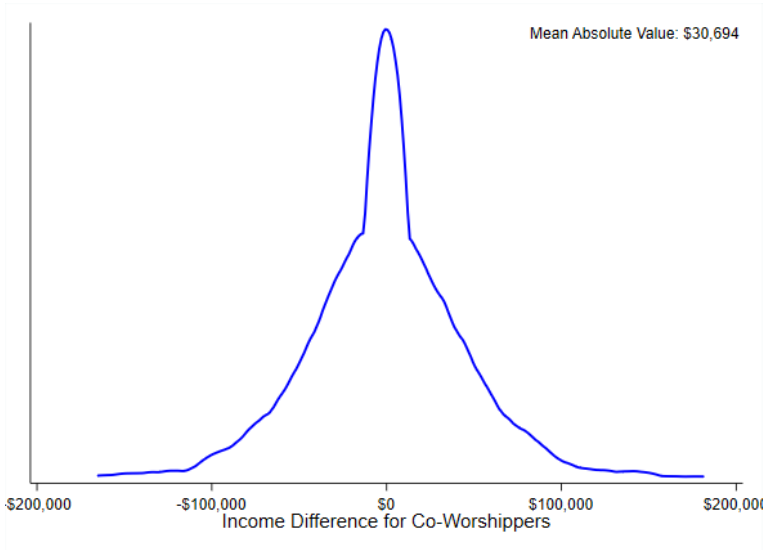

Panel H: Buddhists

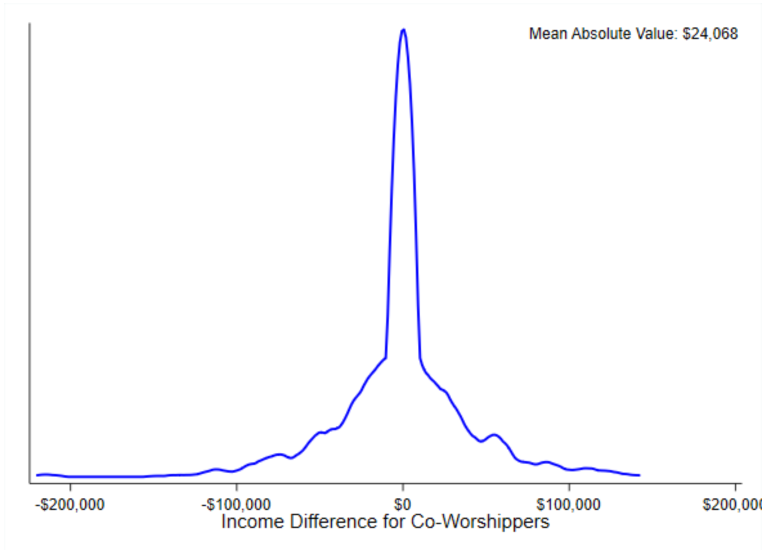

Notes: Each panel presents the distribution of income difference among co-worshippers (defined as two people that visited the same congregation within the same hour and day) for a given religion. Income differences are calculated by randomly selecting pairs of co-worshippers and subtracting one co-worshippers income from the other's.

Figure 11: Income Difference for Co-customers by Location

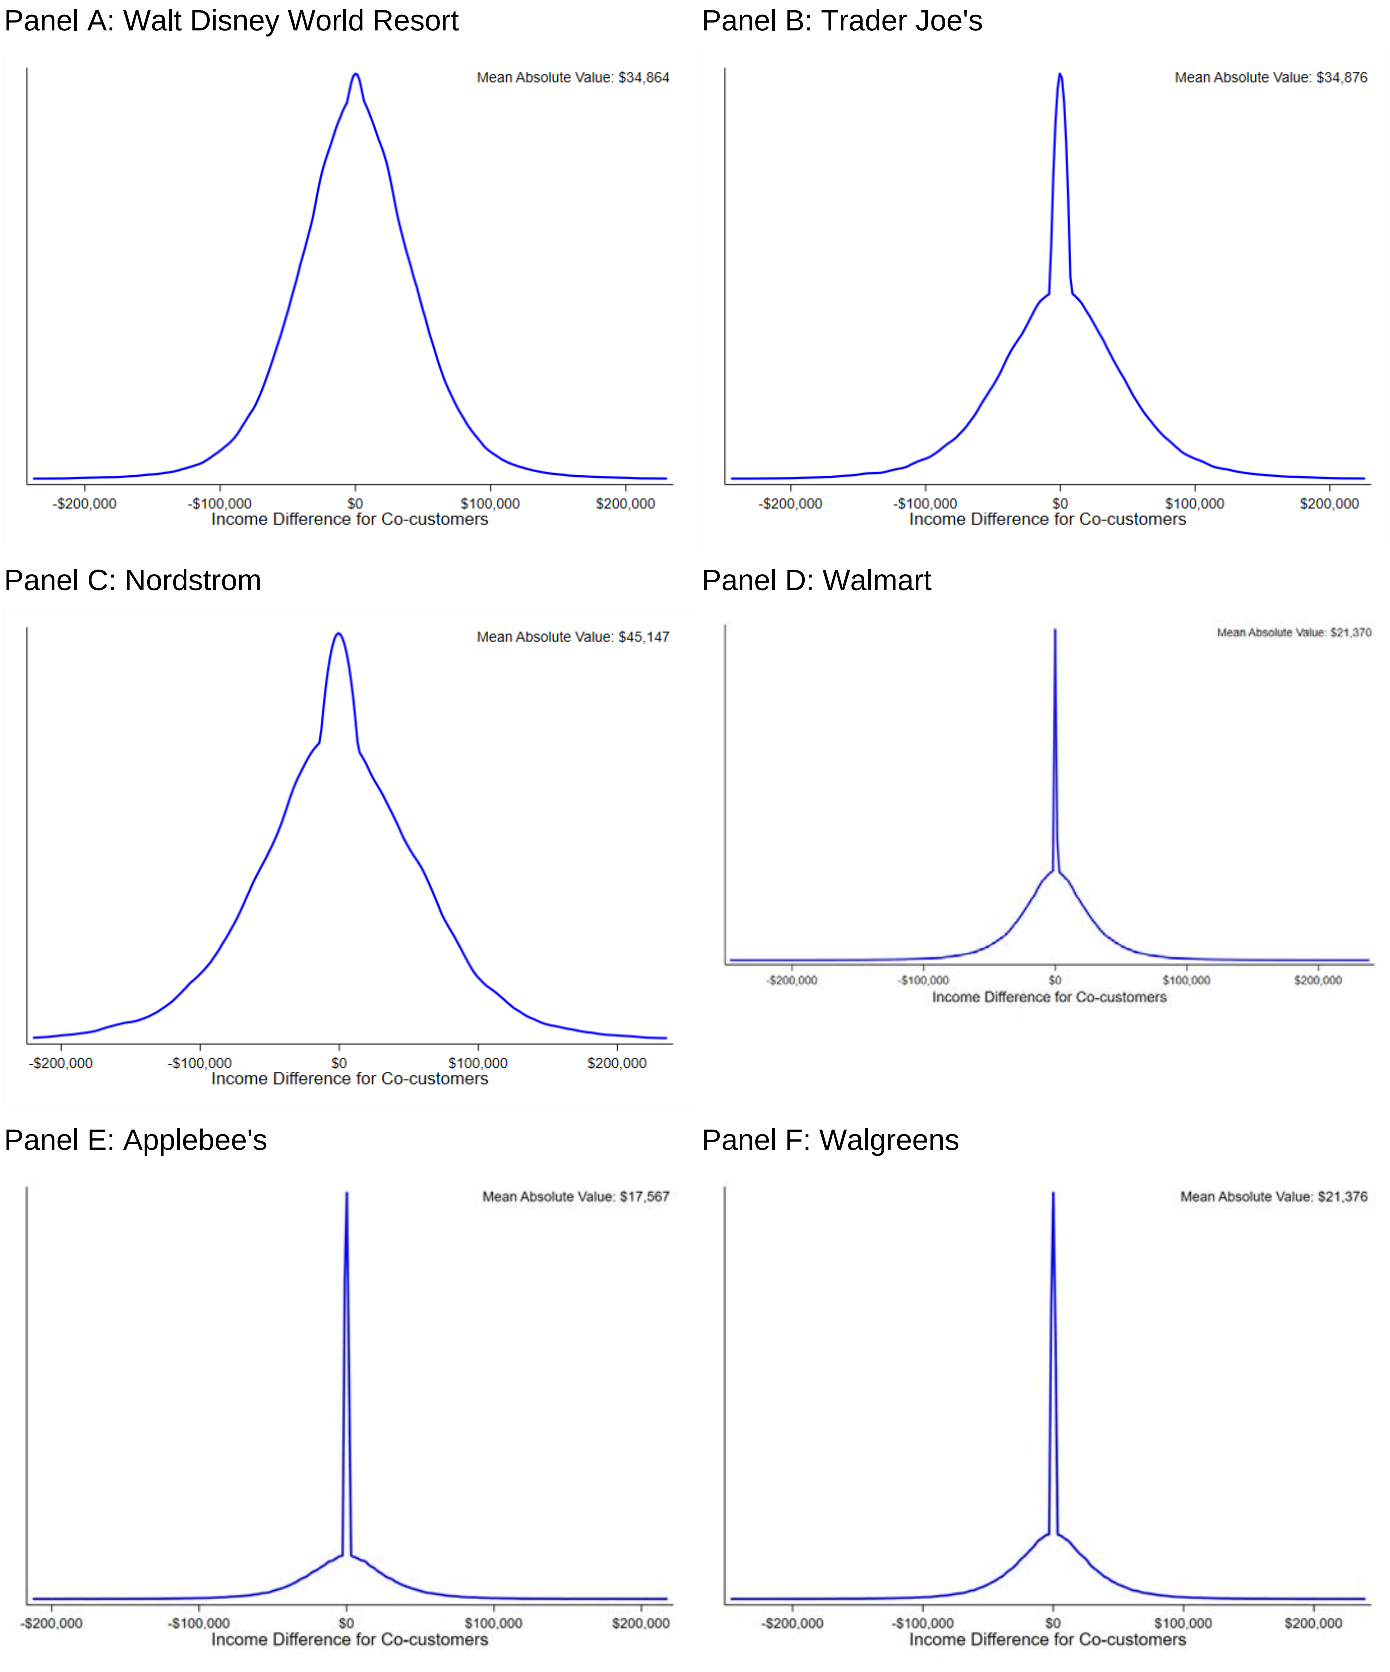

Notes: Each panel plots the distribution of income difference for co-customers, defined as two people that visit the same location on the same day and hour, for a given location. Income differences are calculated by randomly selecting pairs of co-customers and subtracting the median household income of one co-customer from the other.

Figure 12: Religious Attendance by Geography

Panel A: Yearly+ attenders by State

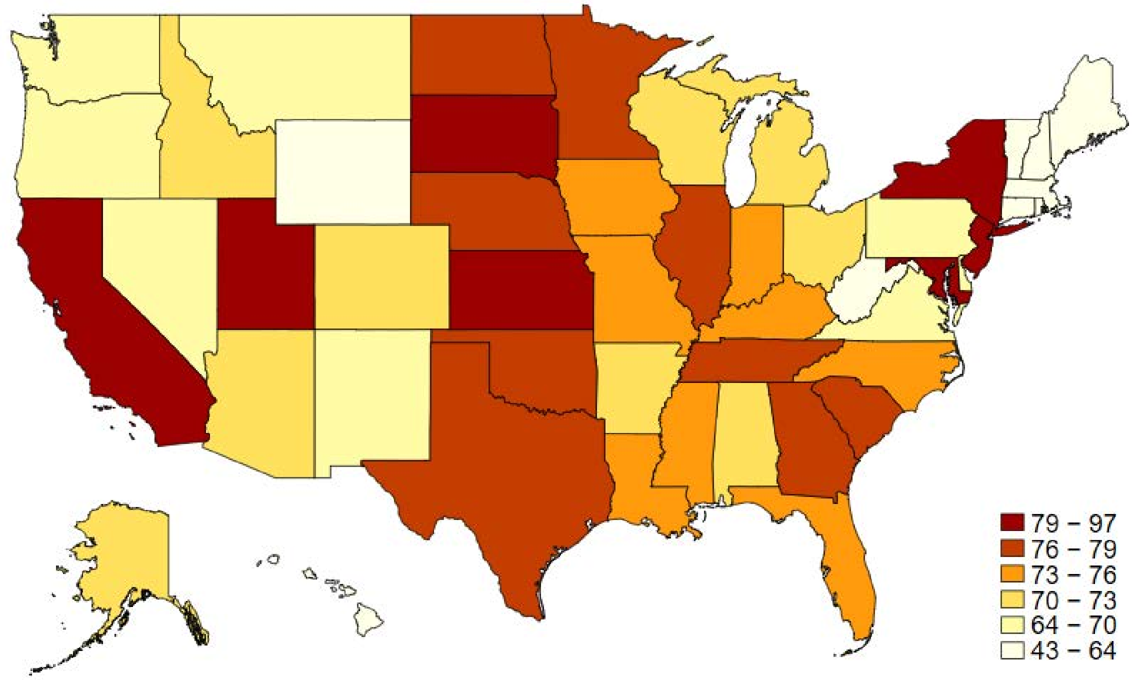

Panel B: Yearly+ Attenders by County

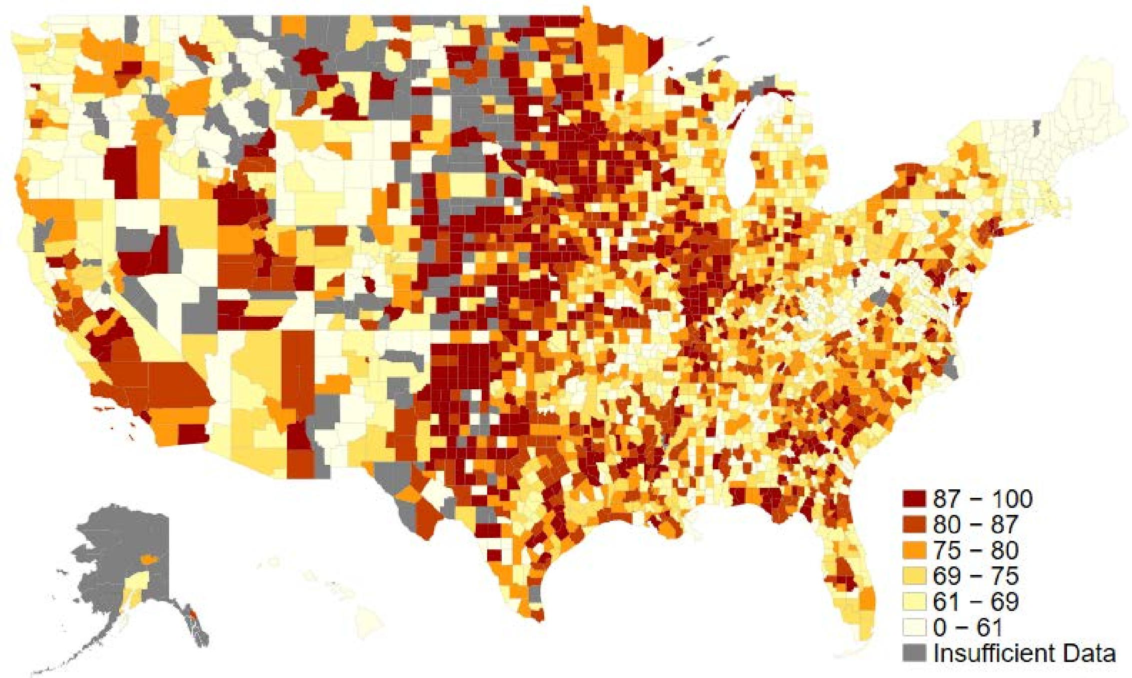

Notes: Panel A shows the percent share of yearly attenders (people who attend a religious worship service at least once during the 47 week study period) in each state's population. Panel B shows the percent share of yearly attenders on the county level.

Figure 13: Comparing Measures of Religious Attendance by State

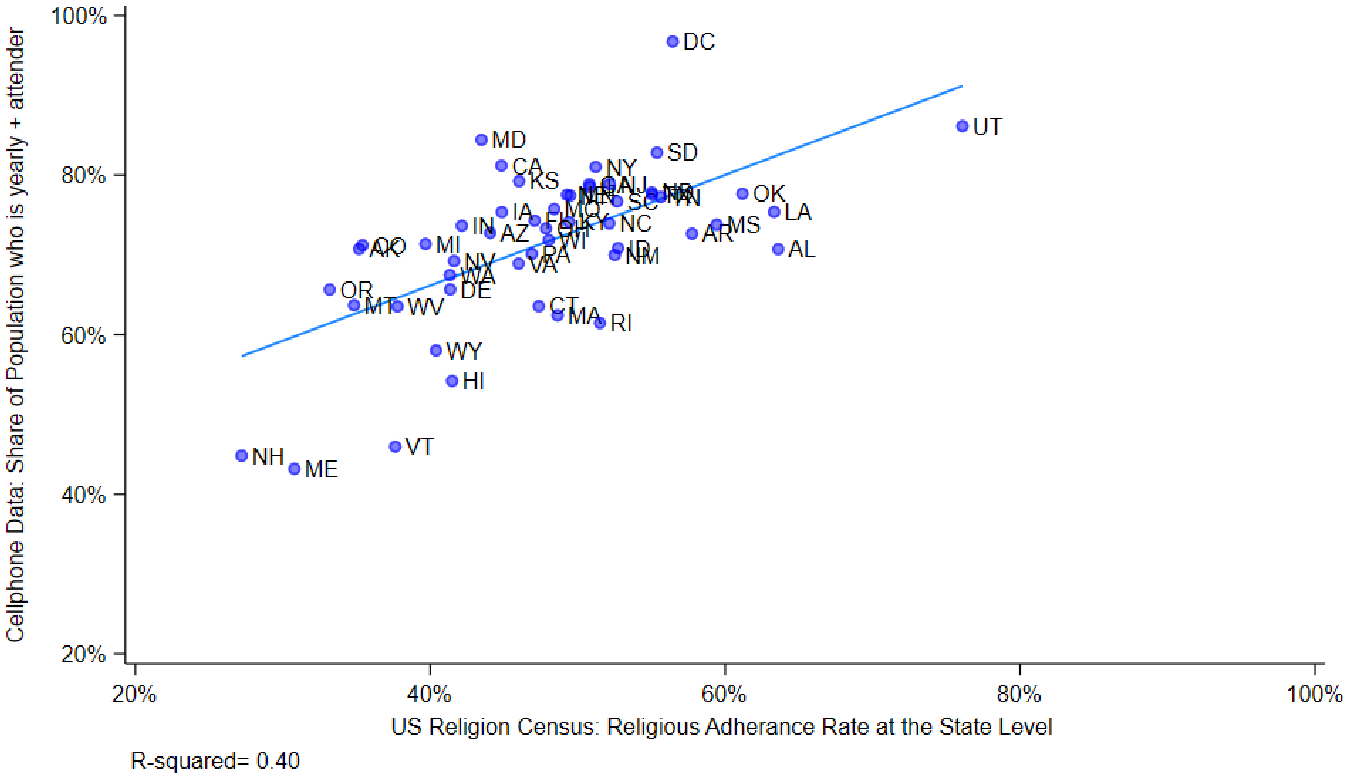

Notes: Each state is represented by a dot. The y-axis is the percent share of yearly attenders (people who attend a religious worship service at least once during the 47 week cellphone sample period) in each state's population as measured in the cellphone data. The x-axis is the fraction of the population that are considered religious adherents according to the US Religion Census.

Figure 14: The Effects of Weather on Religious Attendance

Panel A: Precipitation

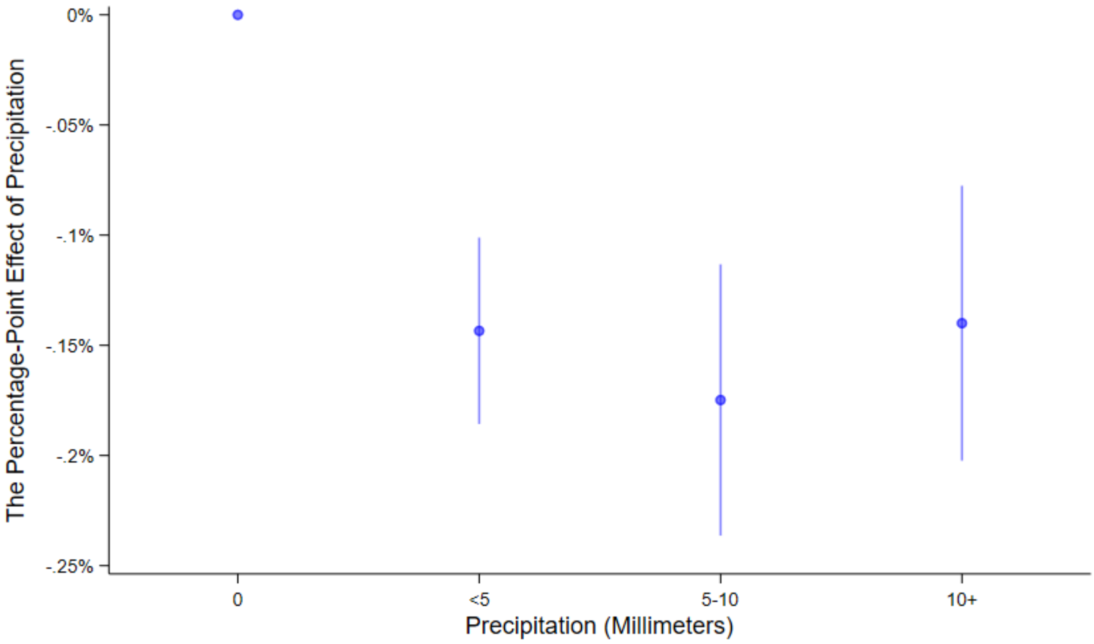

Panel B: Daily Maximum Temperature

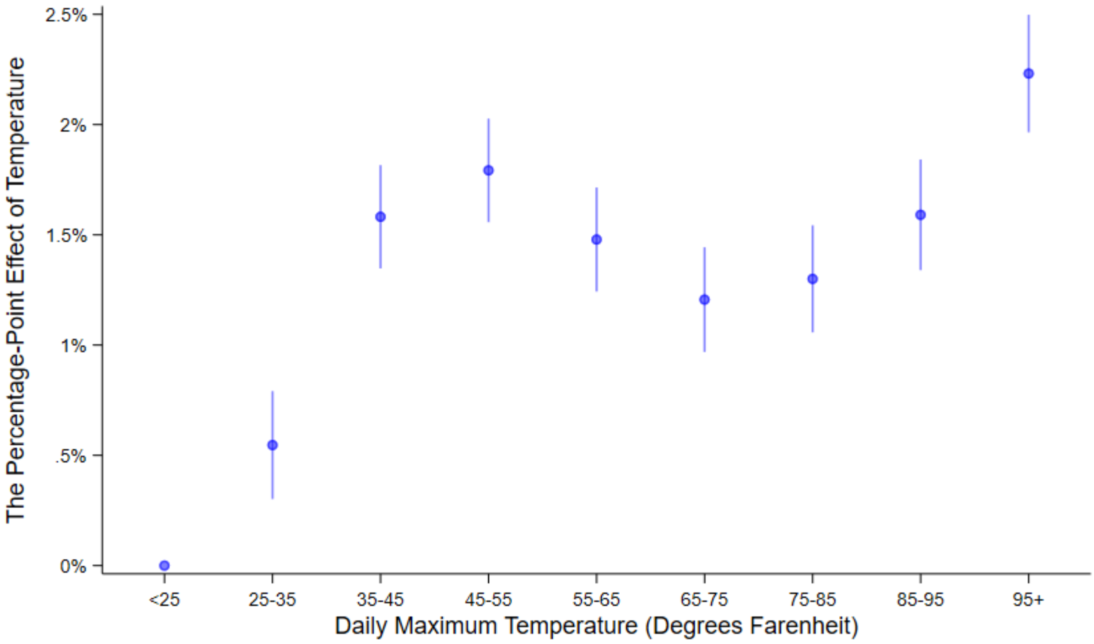

Notes: The analyses in Panel A and Panel B use data for religious worship services, precipitation, and maximum daily temperature for each Sunday within the cellphone sample period. Panel A indicates the percentage point effect of precipitation on religious attendance relative to a base group of zero milimeters of precipitation. There are 4 different bins of precipiation, each with it's own coefficient and 95% confidence interval: zero milimeters (no precipitation), less than five milimeters (but more than zero milimeters), 5-10 milimeters, and more than 10 milimeters. Panel B indicates the percentage point effect of daily maximum temperature on religious attendance relative to a base group of less than 25 degrees farenheit. Daily maximum temperature is split into the following bins: less than 25 degrees, 25-35, 35-45, 45-55, 65-75, 75-85, 85-95, and more than 95 degrees. The regressions in Panel A and Panel B are run at the individual level with date*religion and individual fixed effects. Precipiation and temperature data come from the National Oceanic and Atmospheric Administration (NOAA).

Figure 15: Visits to Other Locations by Religious Attendance

Panel A: Strip Clubs

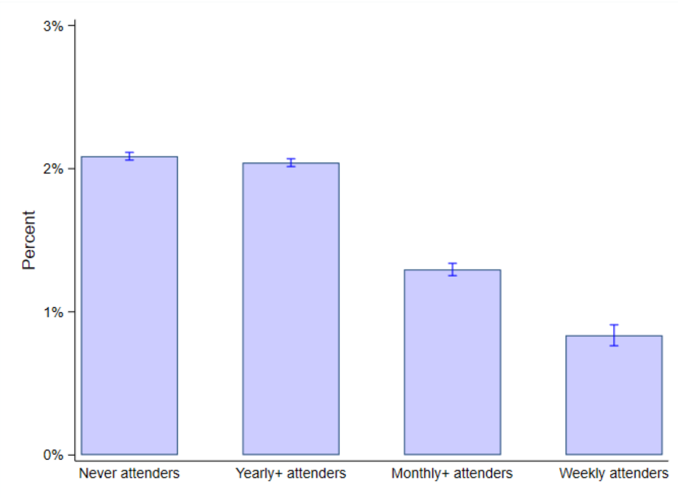

Panel B: Casinos

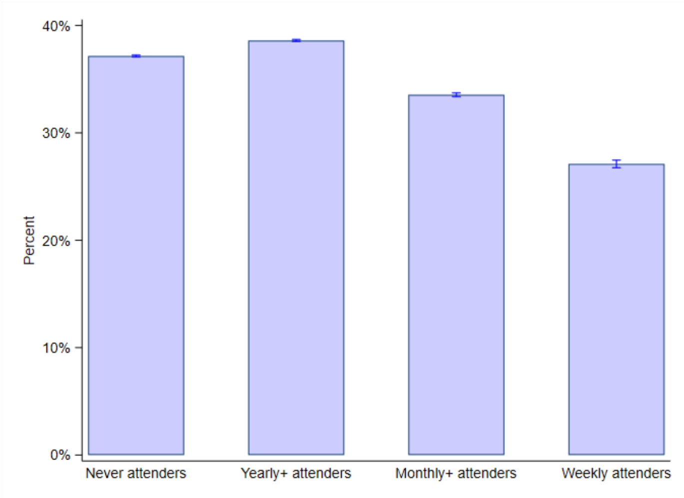

Panel C: Liquor Stores

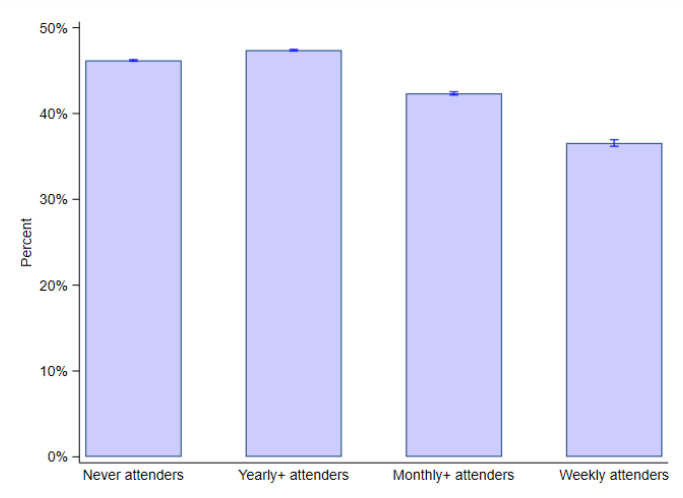

Panel D: Tobacco Stores

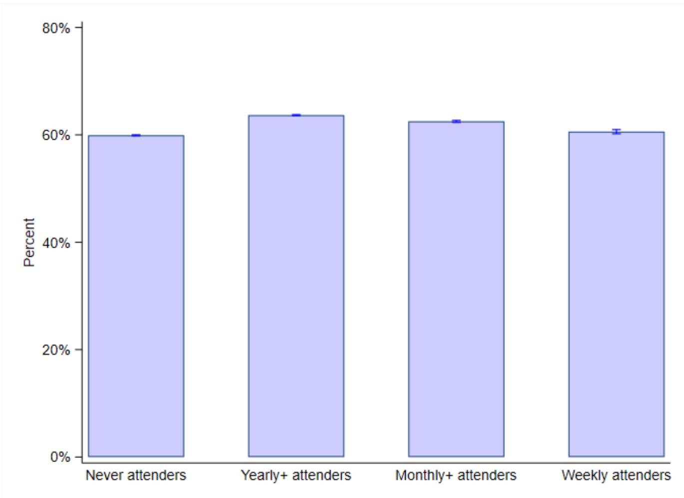

Panel E: Fitness Centers

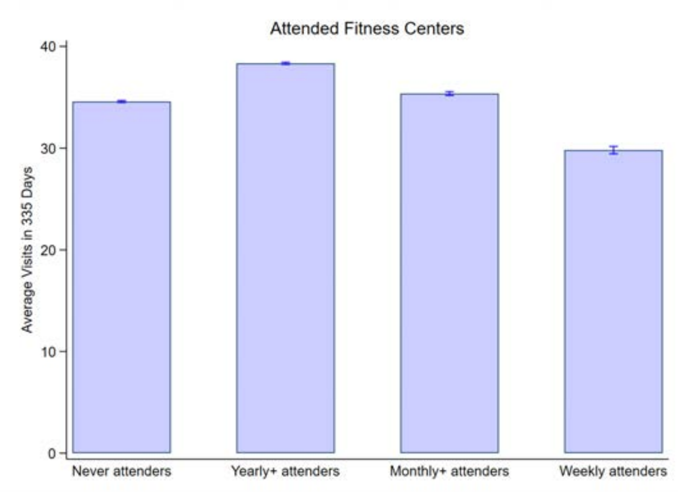

Notes: Panels A-D show the percent of people within each attendance category (never, yearly, monthly, or weekly) that also attend a given location (strip club, casino, liquor store, or tobacco store). Panel F differs from panels A-D in that it shows the average number of visits to fitness centers by religious attendance frequency, rather than the percent of people that visited at least once in the sample period. 95% confidence intervals are overlayed on top of each bar.

Figure 16: Visits to Other Locations by Religion

Panel A: Strip Clubs

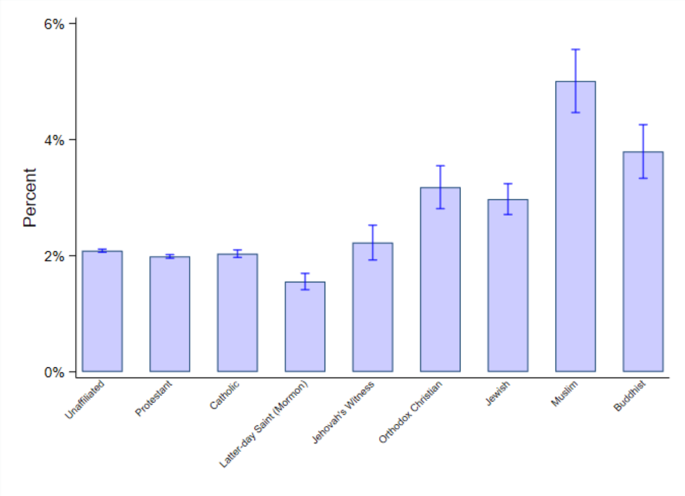

Panel B: Casinos

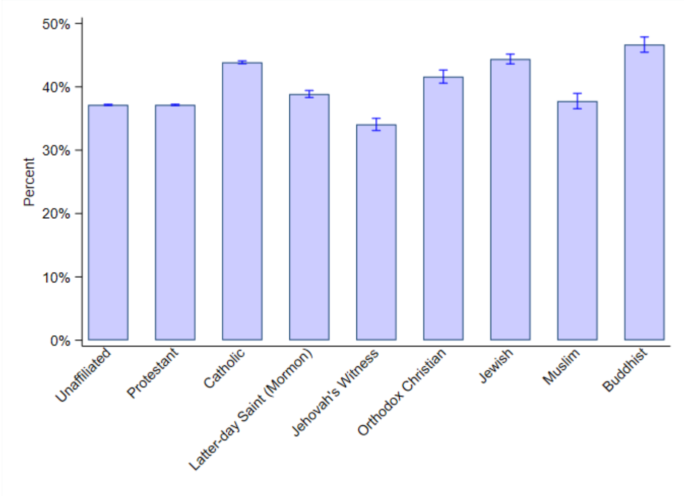

Panel C: Liquor Stores

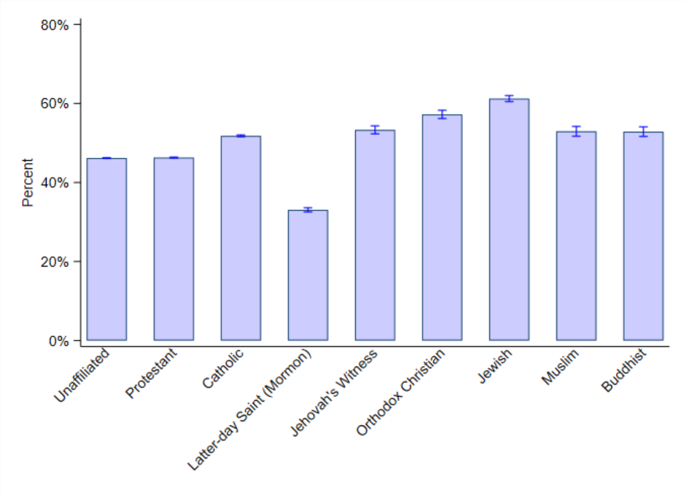

Panel D: Tobacco Stores

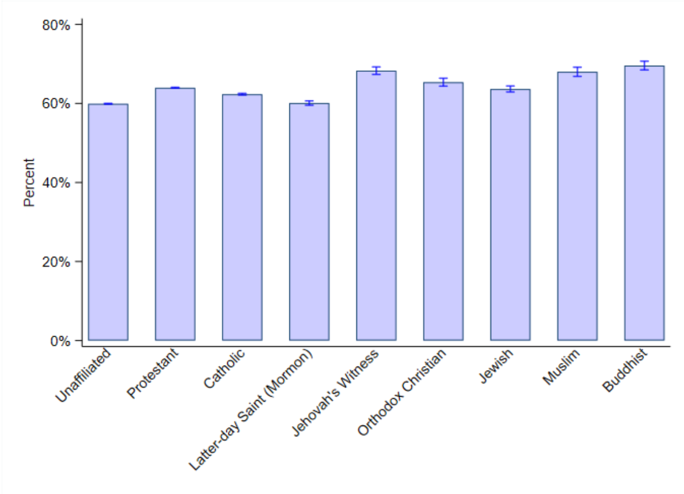

Panel E: Fitness Centers

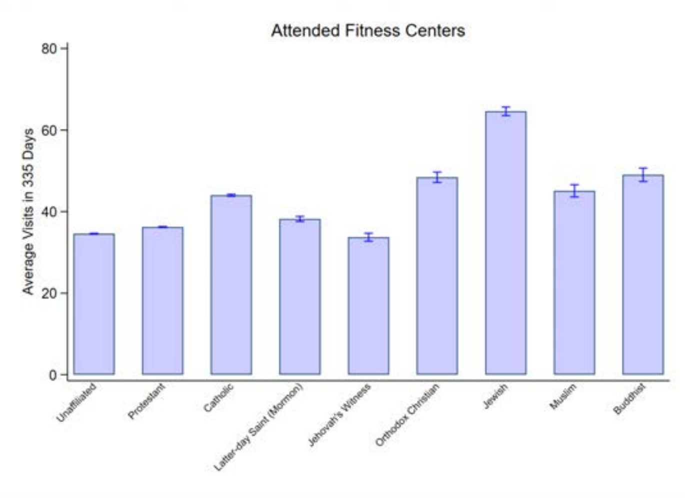

Notes: Panels A-D show the percent of people within each religious affiliation category that also attend a given location (strip club, casino, liquor store, or tobacco store) at least once during during the cellphone sample period. Panel F differs from panels A-D in that it shows the average number of visits to fitness centers by religious affiliation, rather than the percent of people that visited at least once in the sample period. 95% confidence intervals are overlayed on top of each bar.

Figure 17: Religious Attendance with Broader Set of Worship Days

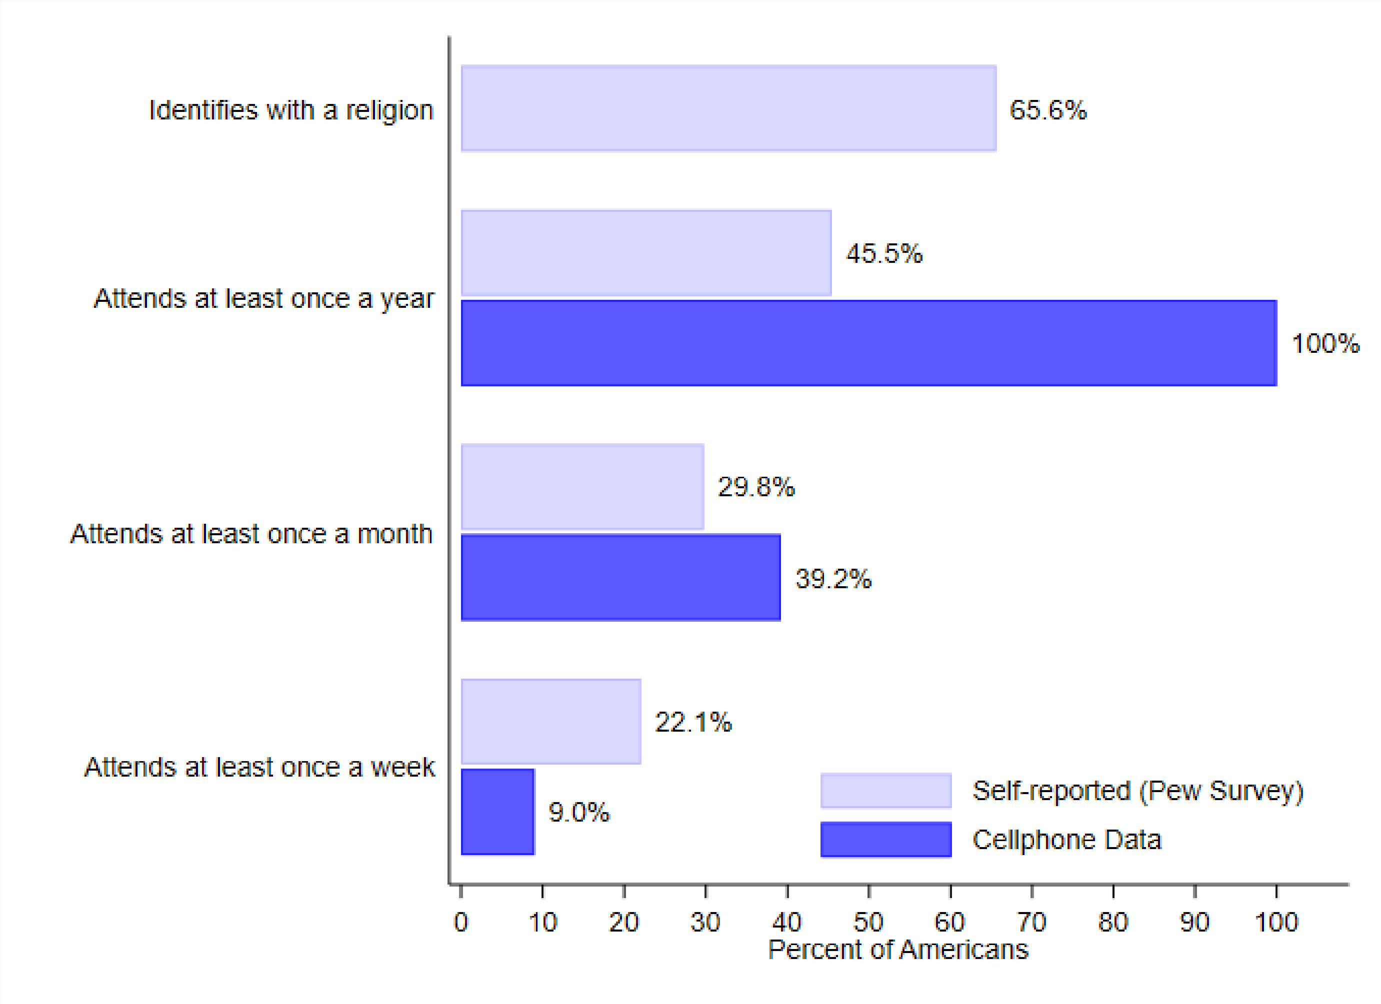

Notes: The light bars show self-reported religious attendance and religious affiliation among Americans based on data from the Pew Research. The dark bars show religious attendance in the cellphone data using an expanded definition of religious worship, which considers all visits to places of worship during the cellphone sample period, regardless of the day of the week.

Figure 18: Week-to-Week Variation in Religious Attendance (Different Measures)

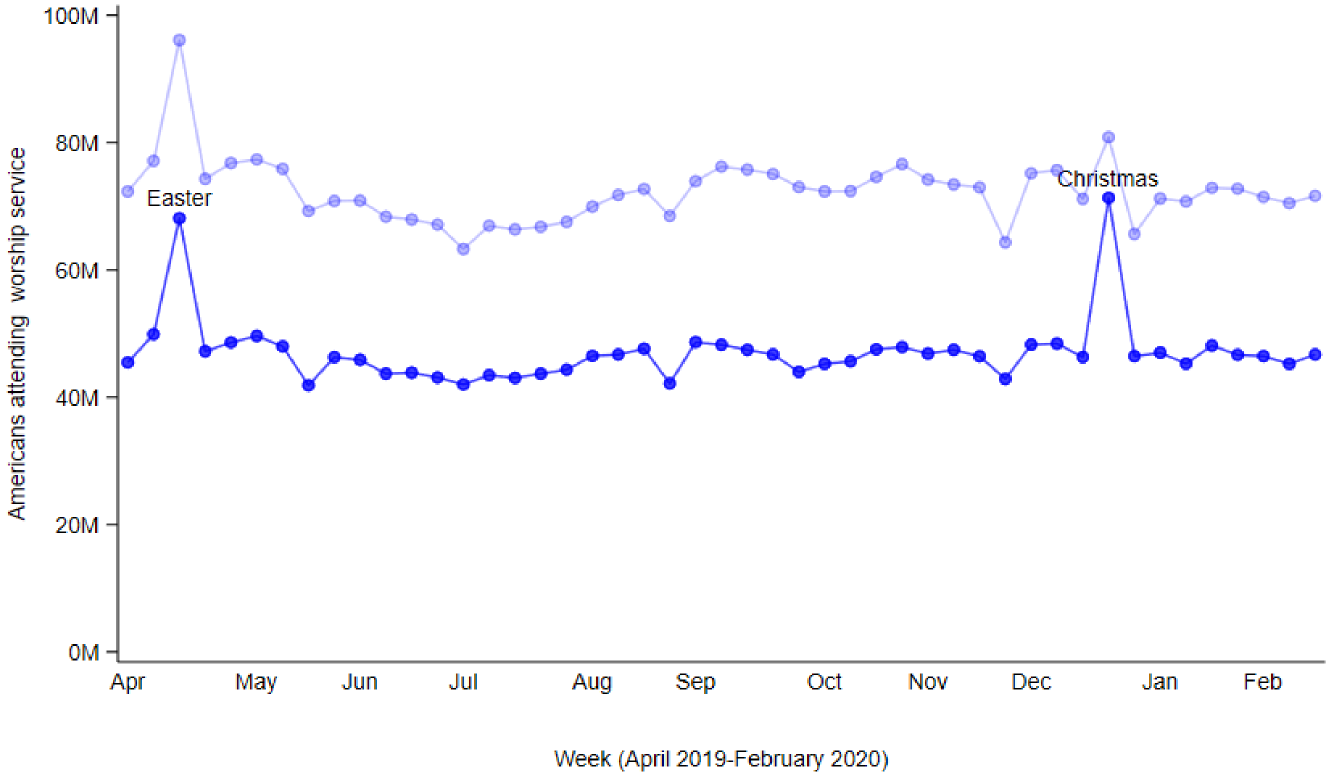

Notes: The darker line and dots indicate the total number of Americans attending a religious worship service during their religion's primary day(s) of worship from April 2019 to February 2020 using the main cellphone sample. The lighter line and dots also indicate total religious attendance, but consider all visits to places of worship during the cellphone sample period, regardless of the day of the week.

Figure 19: Week-to-Week Variation by Religion (Different Measures)

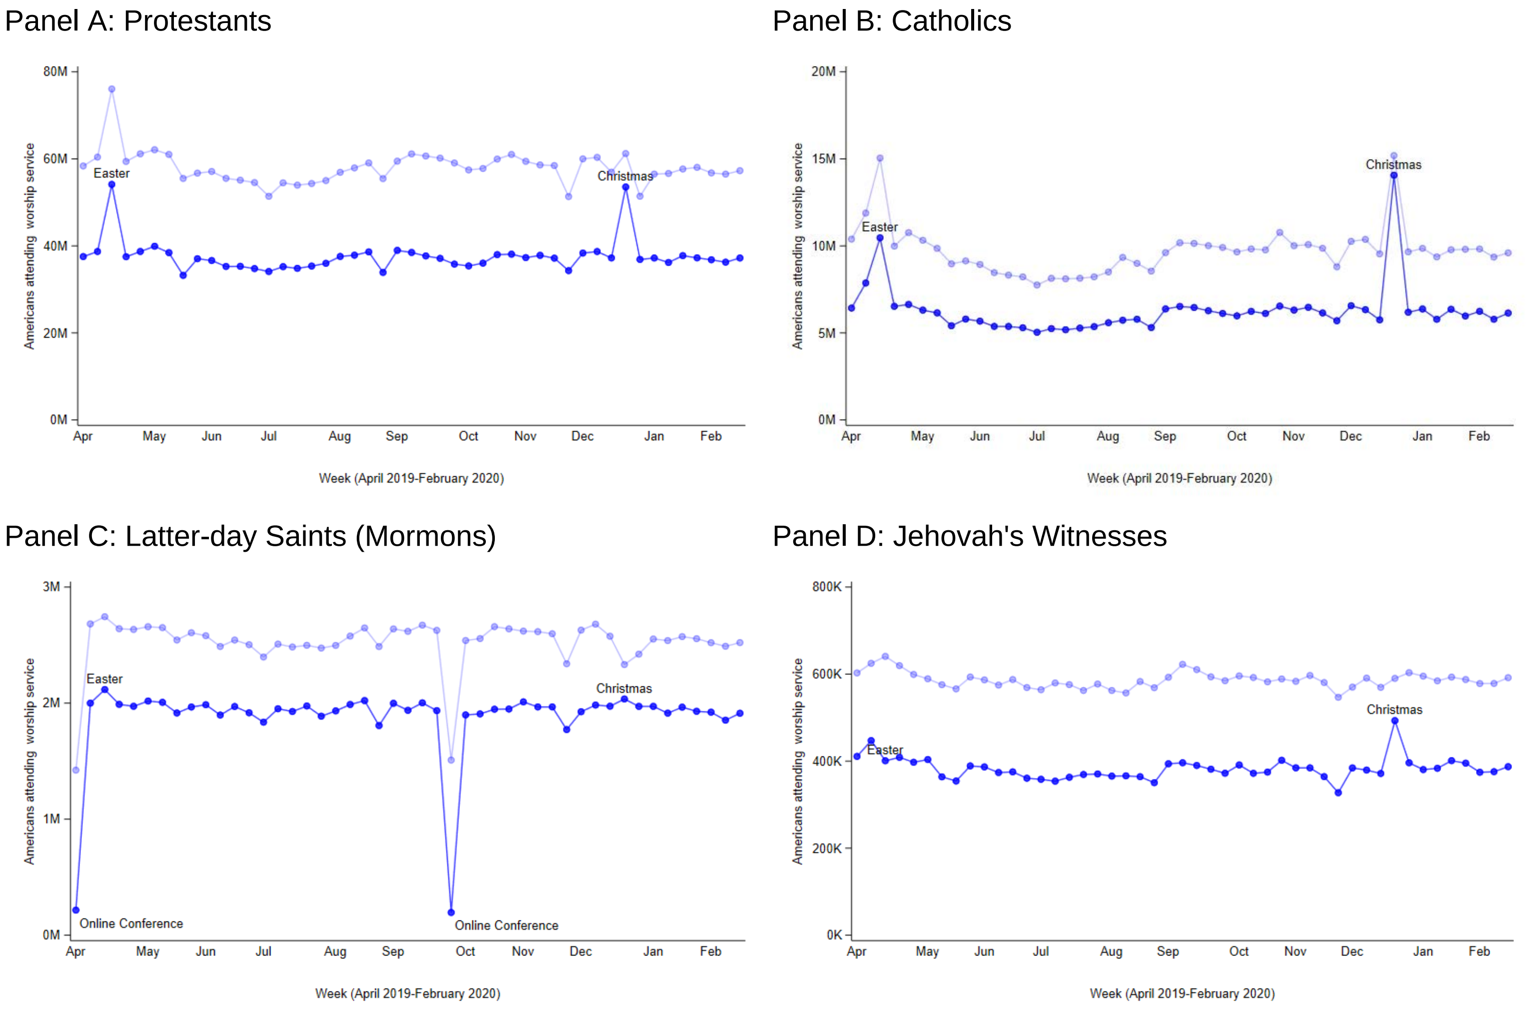

Notes: The darker line and dots indicate the total number of Americans attending a religious worship service during their religion's primary day(s) of worship from April 2019 to February 2020 using the main cellphone sample. The lighter line and dots also indicate total religious attendance, but consider all visits to places of worship during the cellphone sample period, regardless of the day of the week. Each panel presents these weekly totals for a given religion. Important religious holidays and other events that one might expect to significantly impact religious worship attendance in a given week are indicated by text (and shaded area if the holiday occurs over multiple days e.g. Ramadan). The y-axis is scaled by either millions (M) or thousands (K) depending on the size of that religious constituency in the United States.

Figure 19 (Continued): Week-to-Week Variation by Religion (Different Measures)

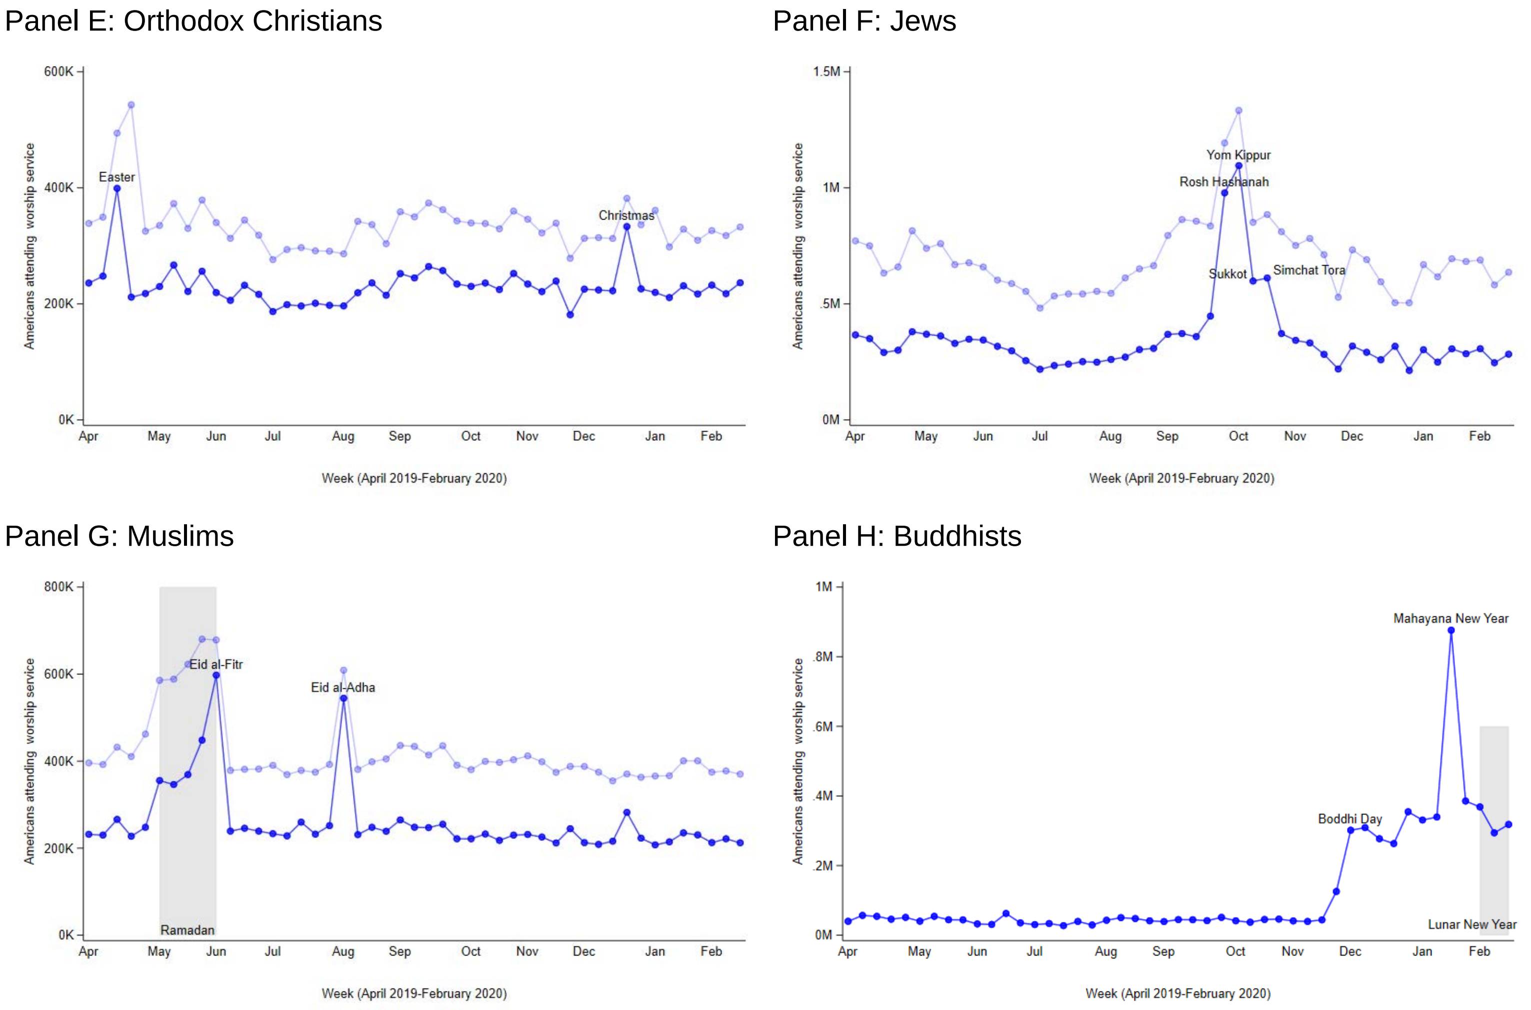

Notes: The darker line and dots indicate the total number of Americans attending a religious worship service during their religion's primary day(s) of worship from April 2019 to February 2020 using the main cellphone sample. The lighter line and dots also indicate total religious attendance, but consider all visits to places of worship during the cellphone sample period, regardless of the day of the week. Each panel presents these weekly totals for a given religion. Important religious holidays and other events that one might expect to significantly impact religious worship attendance in a given week are indicated by text (and shaded area if the holiday occurs over multiple days e.g. Ramadan). The y-axis is scaled by either millions (M) or thousands (K) depending on the size of that religious constituency in the United States.

Figure 20: Smartphone Habits Related to Religious Worship Services and Other Activities (Survey Results)

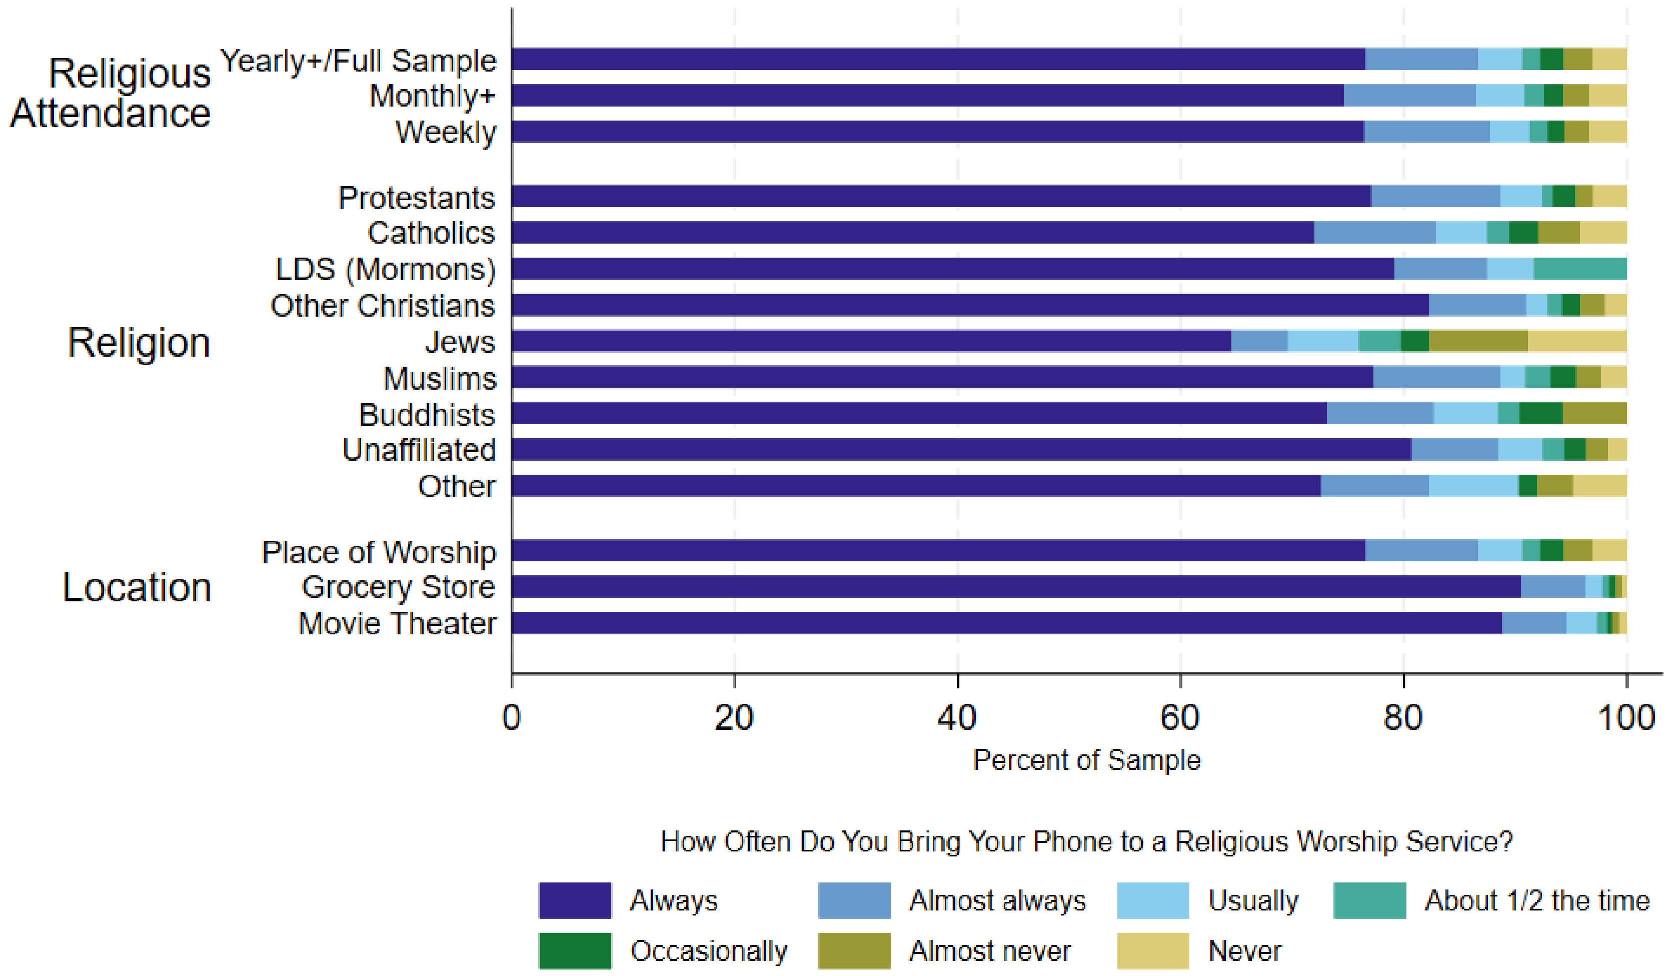

Notes: The bars labeled 'Yearly+/Full Sample' through 'Place of worship' show online prolific survey responses to the following question: 'How often do you bring your phone to a religious worship service?'. Those who do not attend religious worship services at least once per year are not asked this question and are by default excluded from analysis. Therefore, the yearly+ category (those who attend at least once per year) is identical to the 'Place of Worship' category and effectively makes up the entire sample of interest. Answers to this question are broken down by self-reported religious attendance frequency and by religious affiliation. The bars labelled 'Grocery Store' and 'Movie Theater' present responses to questions that are identical to the question asked for the other bars except that 'religious worship service' is replaced by the appropriate location name (either 'grocery store' or 'movie theater'). These two questions are only asked among respondents who report visiting that given location at least once yearly.

Table 1. Survey Measures of Religiosity in America

Survey Name	Data Source	Year(s) Survey Runs	Year(s) Used for Comparison	Observations	% Affiliated	% Yearly+	% Monthly+	% Weekly+	Any Given Sunday
Gallup Poll Social Series	GALLUP	1999(1)2022	2019	29,525	79	45	33	23	-
The General Social Survey	NORC	1972(1)1978, 1980, 1982(1)1991, 1993, 1994(2)2018, 2021, 2022	2018	2,348	76	64	41	23	-
U.S. Religion Census	Association of Statisticians of American Religious Bodies	1980(10)2020	2020	217 religious bodies and 155 congregations	49	-	-	-	-
Baylor Religion Survey	Baylor University	2005, 2007, 2010, 2013, 2017, 2021	2017	1,501	78	65	42	30	-
American Values Survey	Public Religion Research Institute	2010(1)2022	2019	2,527	78	54	37	28	-
Nationscape	Voter Study Group	2021	2021	318,736	74	-	-	-	-
Cooperation Election Study	Harvard University	2006(1)2022	2019	18,000	65	45	32	24	-
American Trends Panel	Pew Research	2014(1)2022	2020 (February)	6,395	63	45	30	22	-
American Time Use Survey	US Bureau of Labor Statistics	2003(1)2023	2019	2,512	-	-	-	-	19
This Paper: Cellphone Data	Veraset	2019, 2020	Apr 2019 - Feb 2020	2,131,844	-	73	21	5	14

Notes: Religious affiliation and attendance in the United States as reported by various surveys. For each survey, data is drawn either from 2019 or the nearest year available.

Table 2: Actual Attendance vs Predicted Attedance from Cellphone Sample

Location Name	Attendance Measure	Cellphone Prediction	Actual Attendance	Locations in Cellphone Data	Actual Number of Locations
AMC Entertainment	Ticket Sales	135,977,120	195,391,000	539	636
Lowe's	Estimated Transactions	1,144,814,720	794,639,053	1,726	1,728
Lowe's	Transactions	1,144,814,720	790,643,077	1,726	1,728
Mcdonald's	Estimated Transactions	4,025,794,560	4,116,138,889	13,549	13,837
Olive Garden	Estimated Number of Entrees	180,133,152	188,682,906	875	868
Planet Fitness	Estimated Workouts	364,002,304	329,000,000	2,040	2,086
Six Flags	Reported Attendance	23,190,272	26,158,000	21	23
Target	Estimated Transactions	2,009,633,152	1,271,373,626	1,981	1,868
The Home Depot	Estimated Transactions	1,272,200,448	1,396,091,848	1,955	1,984
Walmart	Estimated Transactions	7,109,920,768	6,588,754,695	4,674	4,769

Notes: Estimated visits to certain locations based on cellphone data compared to estimates based on official reported figures. The "Attendance Measure" column indicates how the "Actual Attendance" column is calculated. The "Locations in Cellphone Data" column shows the number of locations for a given "Location Name" that can be identified in the cellphone data. Actual attendance and location counts are based on information provided in official documents (e.g. SEC filings, annual investor reports) when available, and found through google searches otherwise.

Table 3. Robustness to Sample Restrictions

Imposed Restrictions		Results			
Days in Sample	Average Hours Per Day	% Yearly+	% Monthly+	% Weekly	Observations
325+	0+	72.7	20.7	4.6	2,131,844
327+	0+	72.6	20.7	4.8	1,909,179
329+	0+	72.2	20.7	4.9	1,639,821
331+	0+	71.6	20.6	5.0	1,326,773
333+	0+	70.5	20.2	5.1	936,364
325+	12+	72.7	20.7	4.7	2,104,296
325+	14+	72.3	20.8	4.8	2,014,766
325+	16+	70.8	20.6	4.9	1,744,033
325+	18+	67.1	19.7	5.1	1,212,401
325+	20+	58.7	16.9	4.9	521,689
330+	12+	72.0	20.7	5.0	1,478,217
330+	14+	71.7	20.7	5.0	1,433,391
330+	16+	70.5	20.6	5.1	1,277,506
330+	18+	67.2	19.8	5.3	941,381
330+	20+	59.2	17.1	5.1	444,385

Notes: Robustness of religious attendance rates to varying sample restrictions. "Days in Sample" refers to the minimum number of days during the study period that a cellphone must appear in the raw data to be included in a given sample. "Average hours Per Day" refers to the minimum number of hours per day that a cellphone must appear in the raw data to be included in a given sample.

Appendix Figure 1: Ping Frequency for All Veraset Cellphones

Panel A: Full Distribution

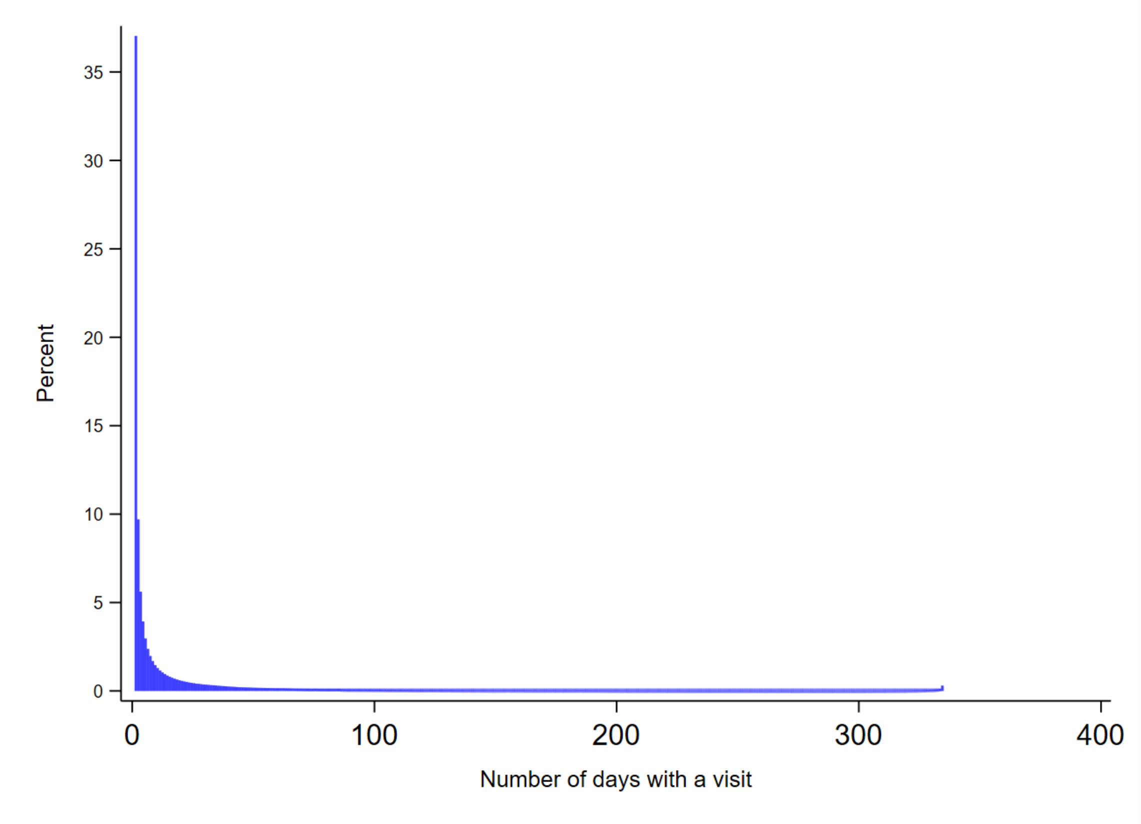

Panel B: Distribution for Cellphones with >200 Days

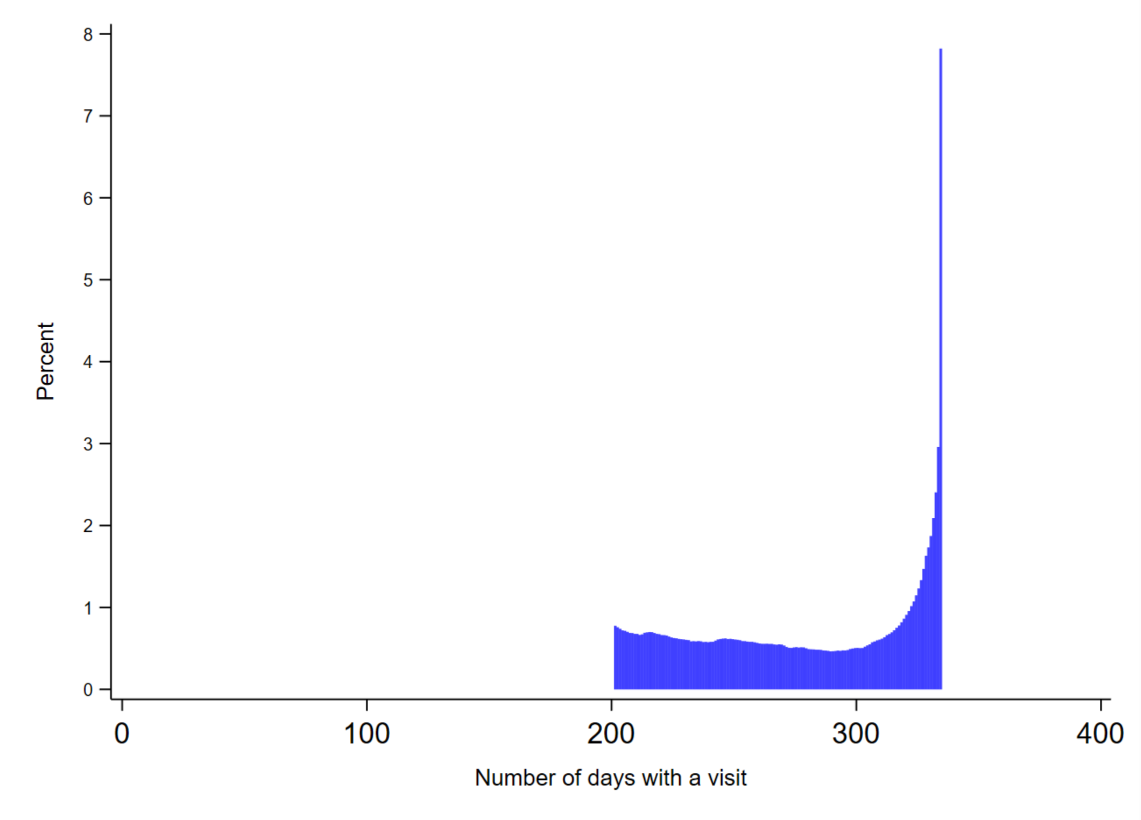

Notes: Panel A shows the full distribution of the frequency of cell phone pings in the sample. Panel B shows the same distribution but only among cellphones that record a visit on at least 200 days in the cellphone sample period.

Appendix Figure 2: Average Hours Each Day with a Known Phone Location - Main Sample

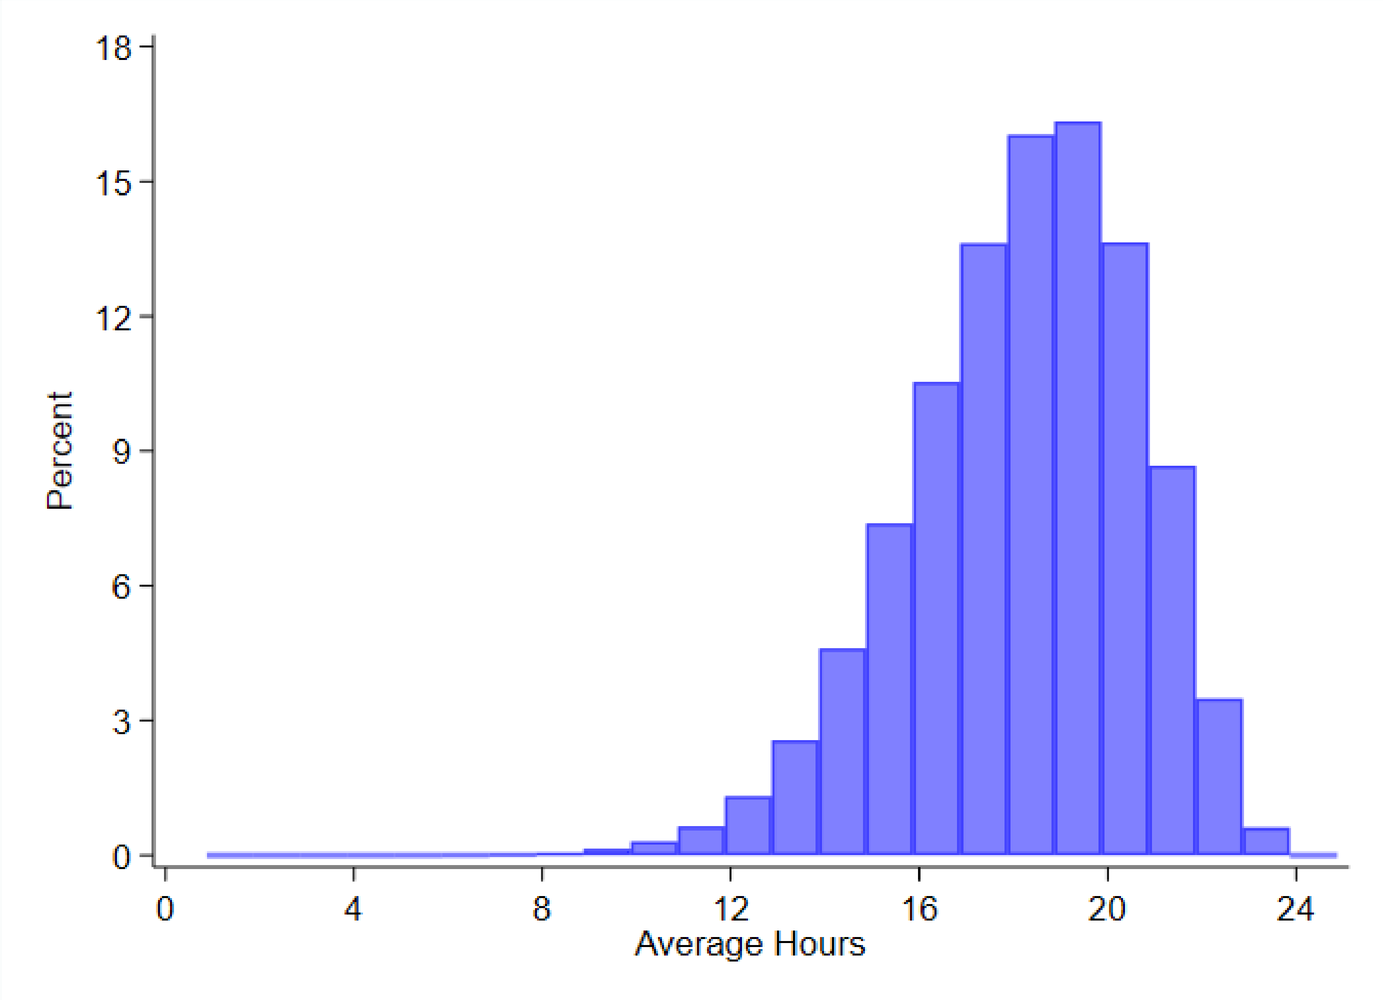

Notes: Distribution of average hours accounted for in the main sample of cellphones with 1-hour bins.

Appendix Figure 3: Geographic Representation of Cellphone Sample

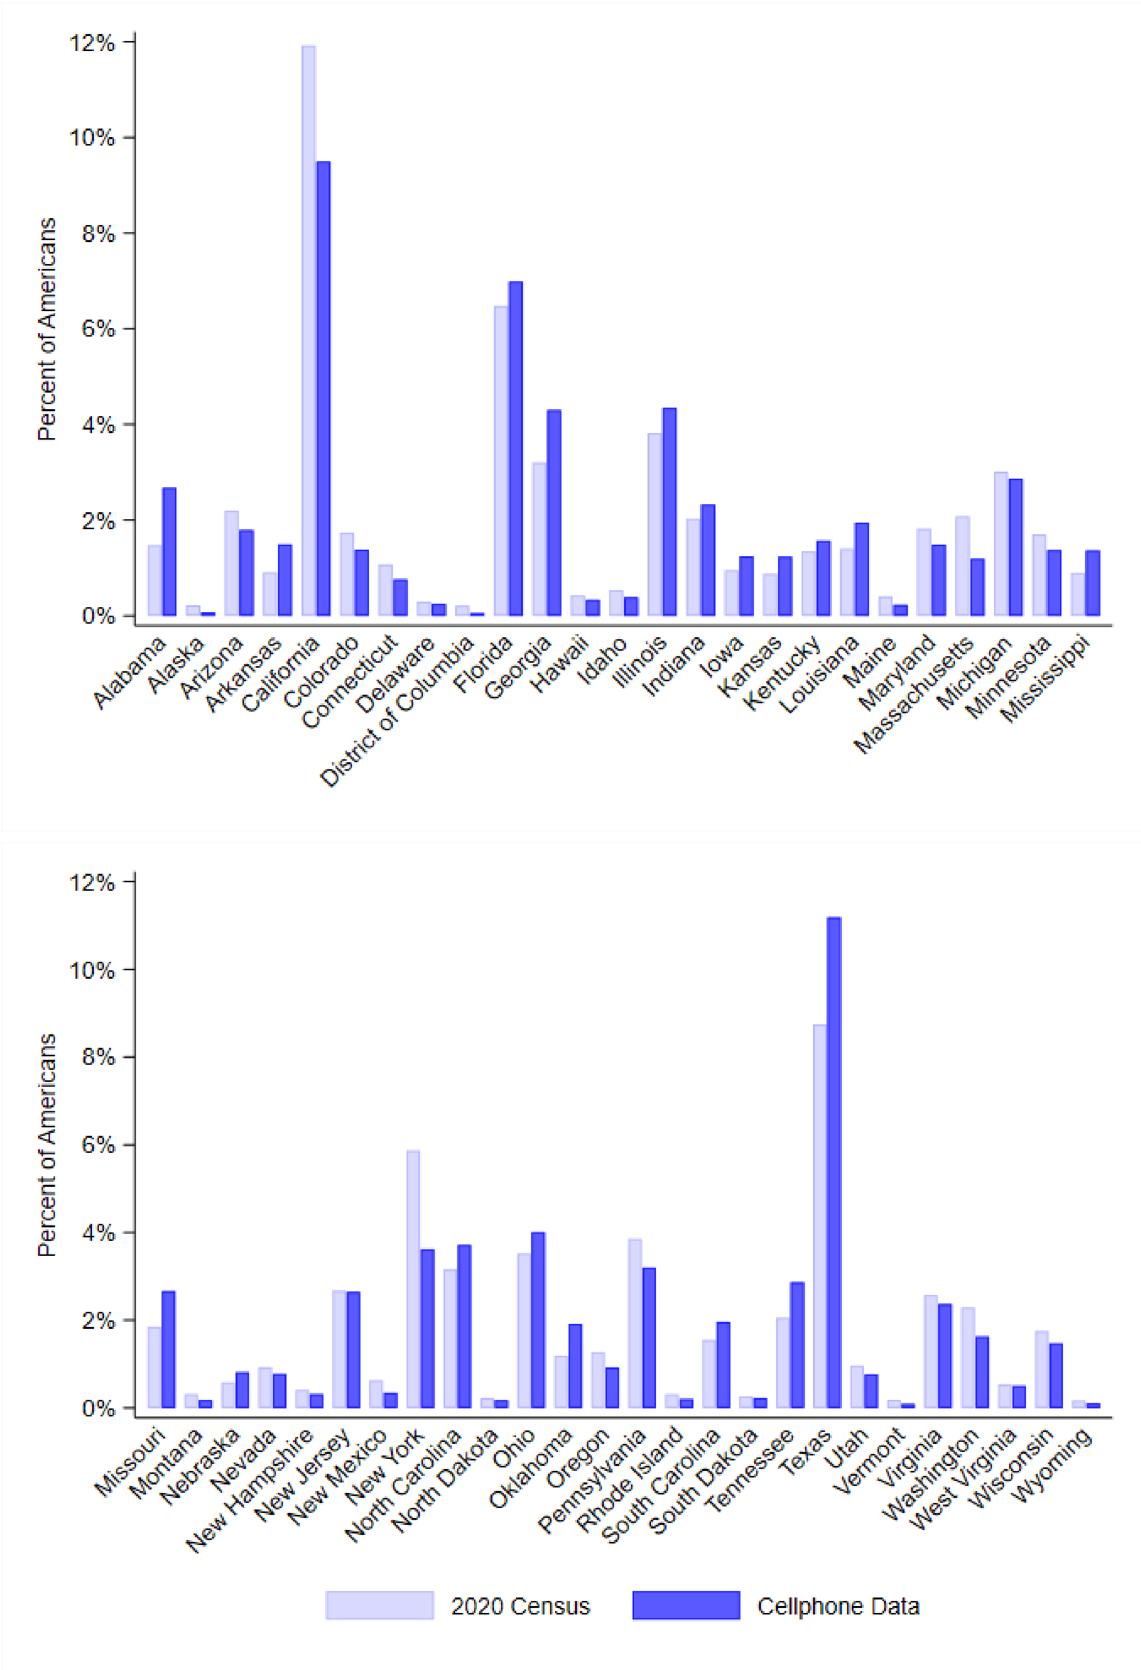

Notes: The lighter bars show the population of each US state as a percentage of the entire US population. The darker bars show the number of cellphones in each US state as a percentage of the cellphone sample.

Appendix Figure 4: Median Household Income Representation of Cellphone Sample

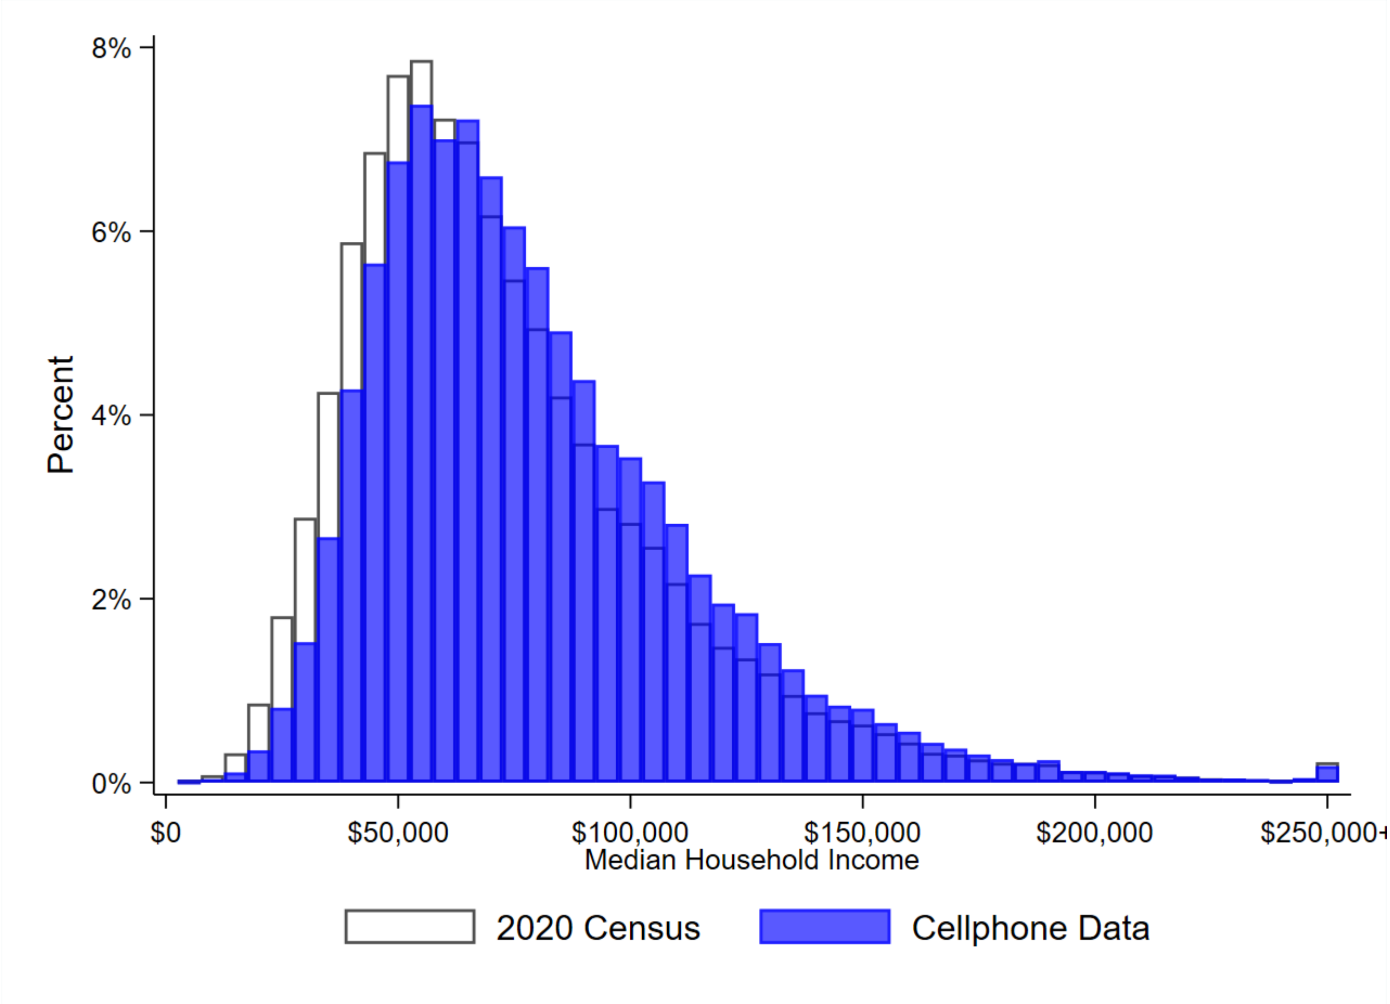

Notes: The distribution of median household income in the US according to the 2020 US census is indicated by the white bars. The distribution of median household income across the cellphone sample is indicated by the blue bars.

Appendix Figure 5: Population Density Representation of Cellphone Sample

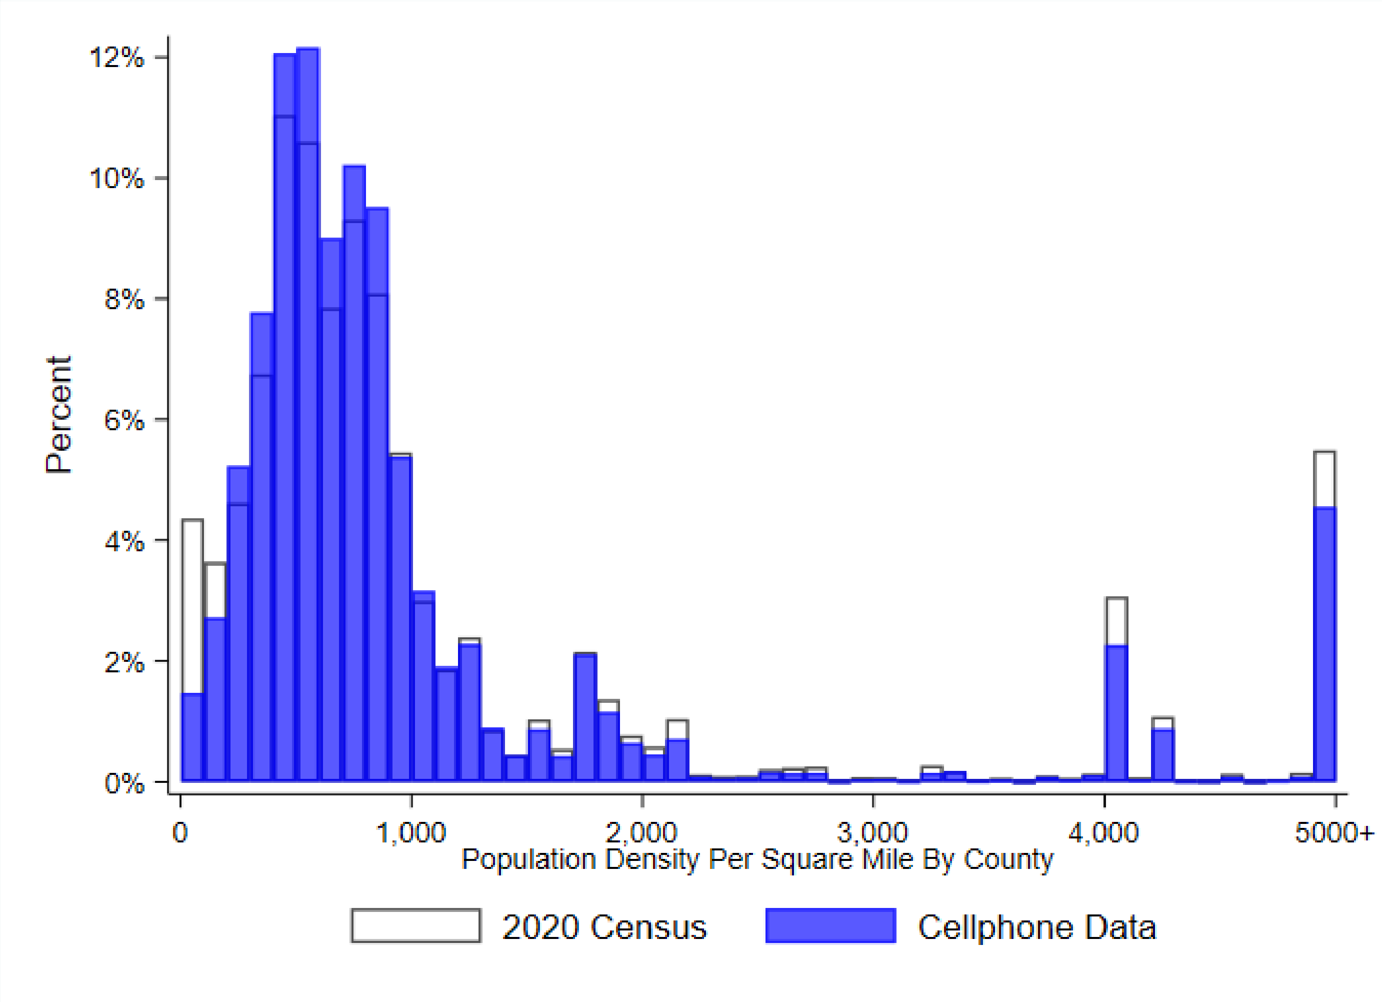

Notes: The distribution of population density per square mile by county in the US according to the 2020 US Census is indicated by the white bars. The distribution of population density per square mile by county according to the cellphone sample is indicated by the blue bars.

Appendix Figure 6: Predicting Baseball Game Attendance with Cellphone Sample

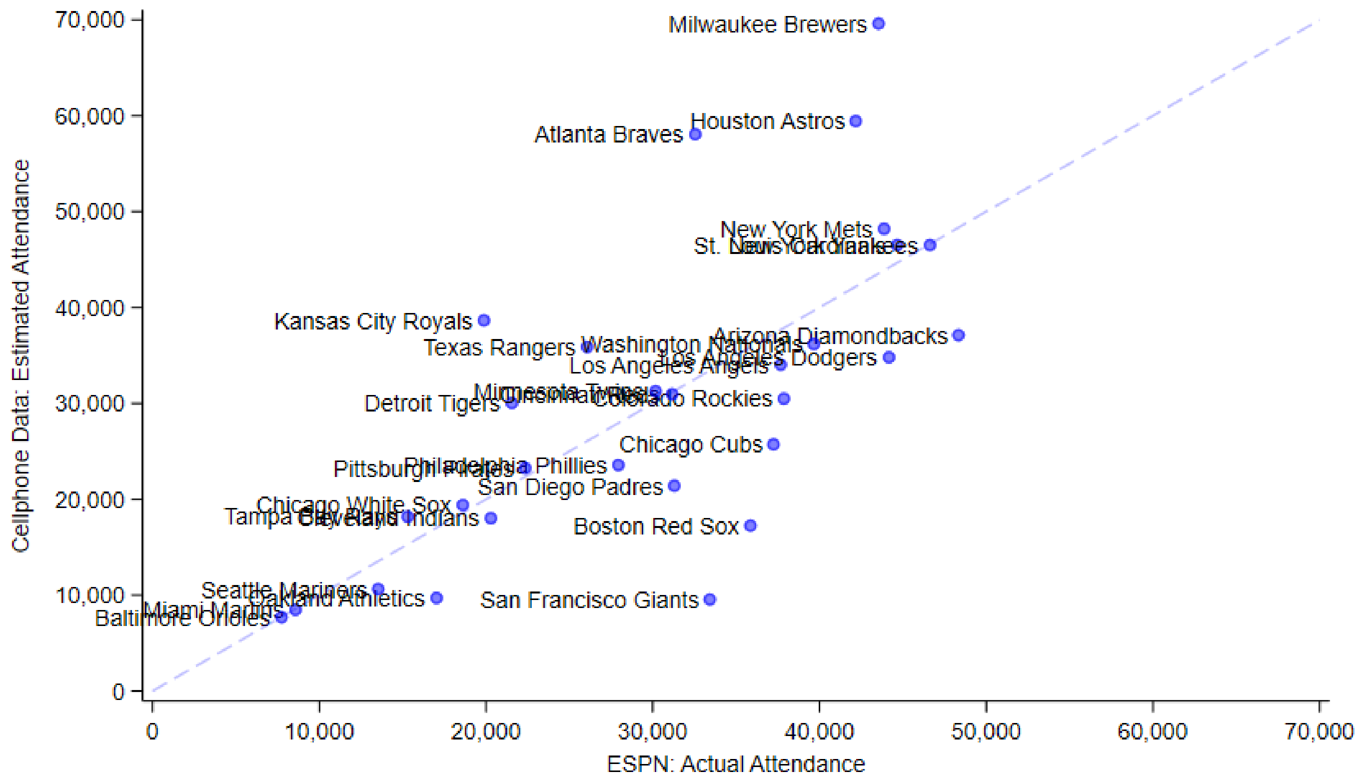

Notes: The y-axis is the attendance of a randomly selected home game for each Major League Baseball (MLB) team estimated using the cellphone sample. The x-axis is the actual attendance for that game according to ESPN. The dotted blue line is a 45 degree line. Any points above (below) the line indicate an over(under)-estimation of attendance.

Appendix Figure 7: Predicting Basketball Game Attendance with Cellphone Sample

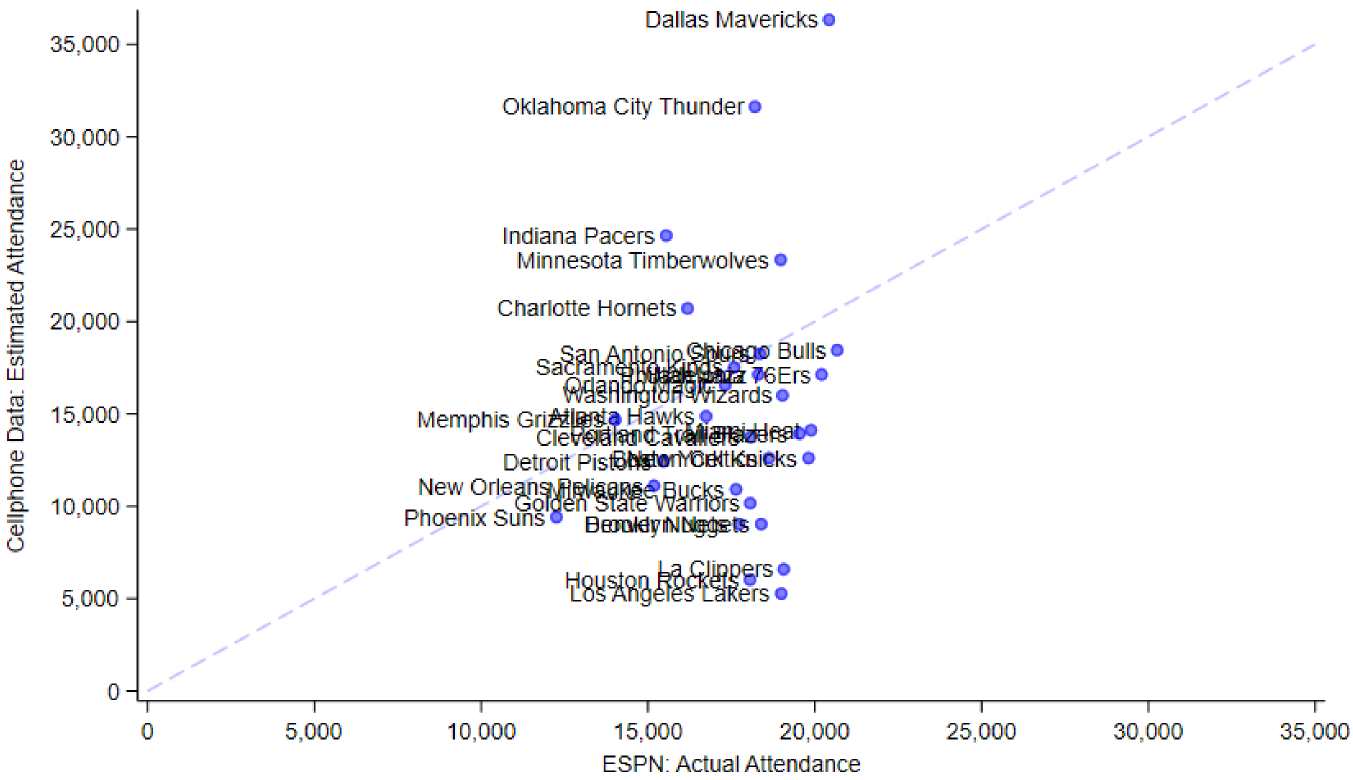

Notes: The y-axis is the attendance of a randomly selected home game for each National Baskebtall Association (NBA) team estimated using the cellphone sample. The x-axis is the actual attendance for that game according to ESPN. The dotted blue line is a 45 degree line. Any points above (below) the line indicate an over(under)-estimation of attendance.

Appendix Figure 8: Income Differences by Religion and Frequency

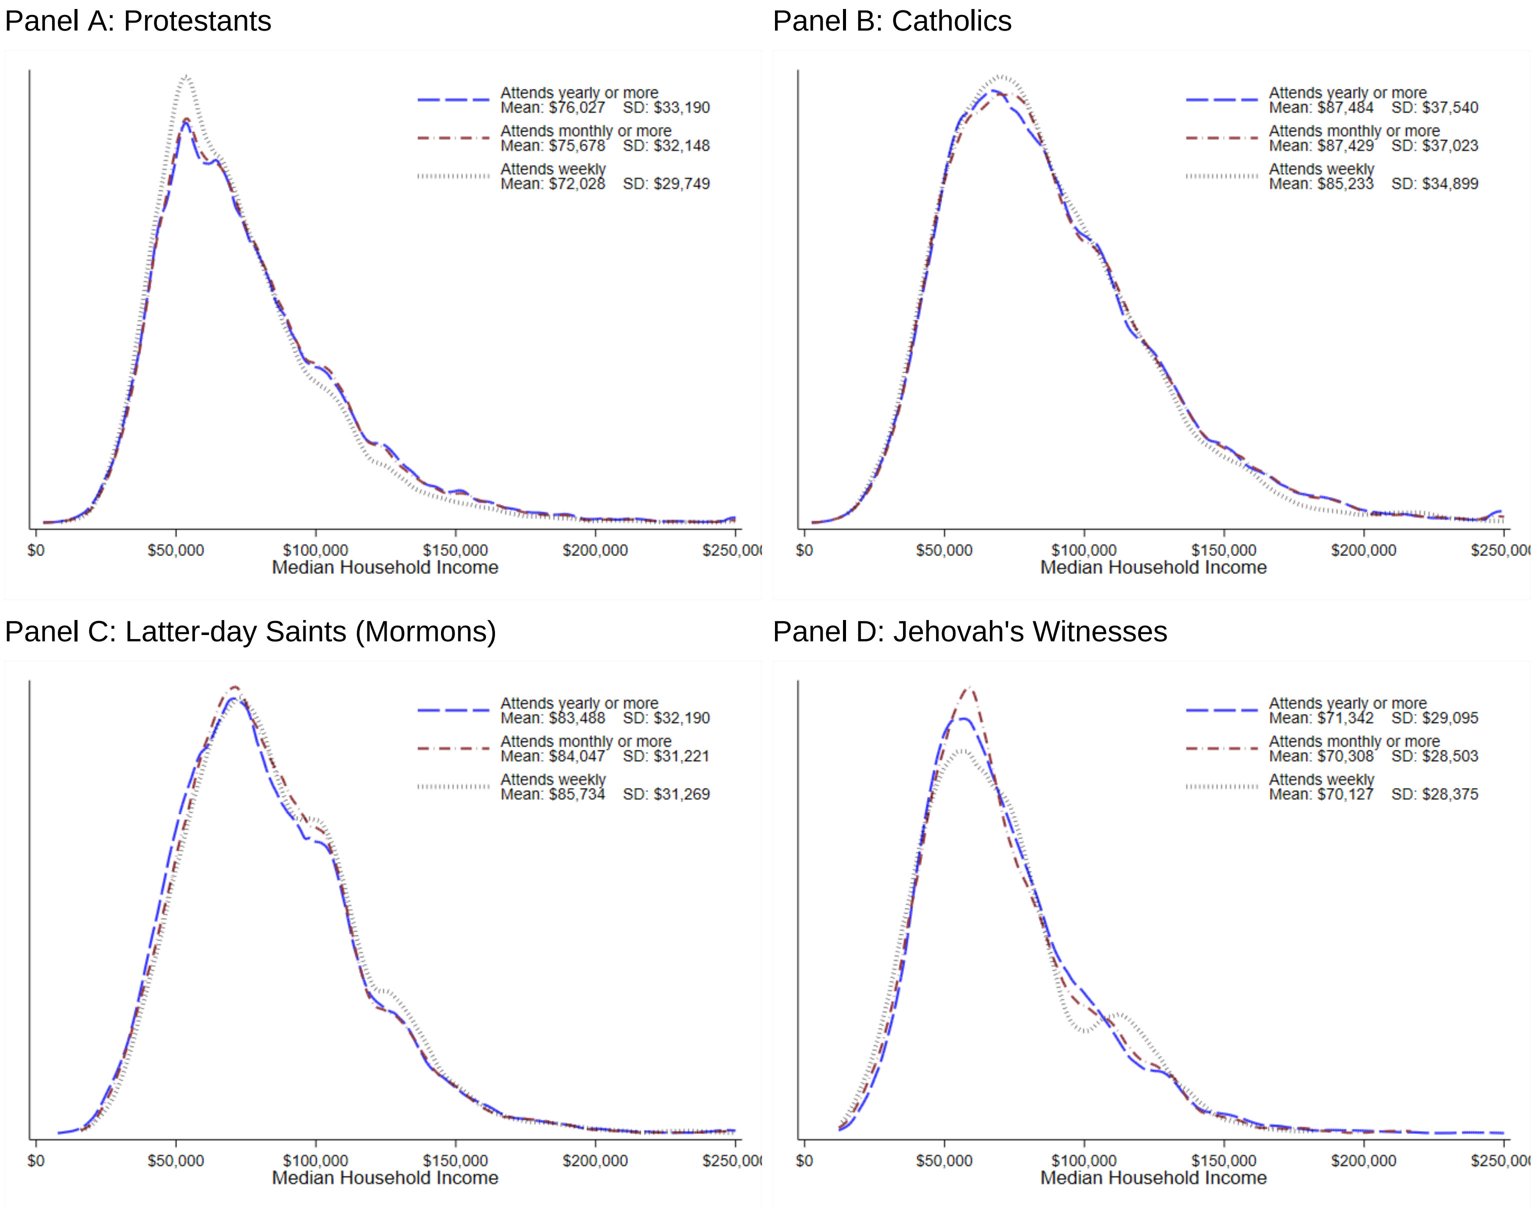

Notes: Each graph plots the distribution of median household income for yearly attenders (at least one religious worship visit during the sample period), monthly attenders (at least 11 religious worship visits during the sample period), and weekly attenders (at least 36 religious worship visits during the sample period). Each panel presents these graphs for a given religion.

Appendix Figure 8 (Continued): Income Differences by Religion and Frequency

Panel E: Orthodox Christians

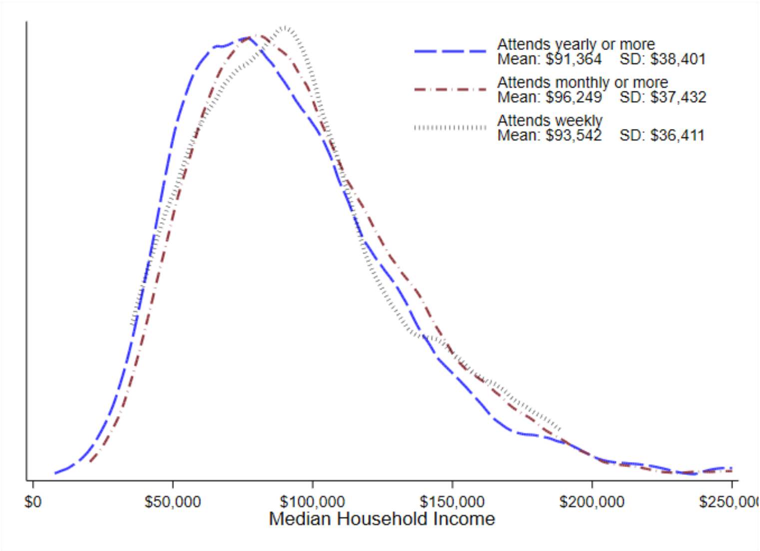

Panel F: Jews

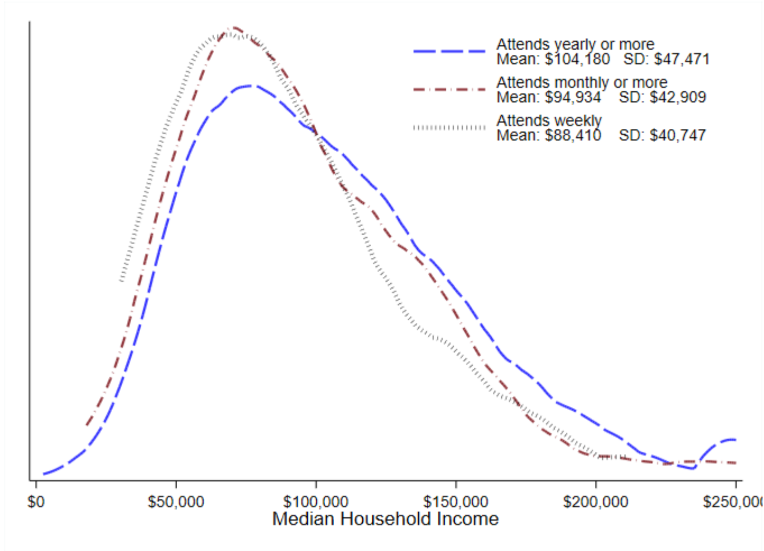

Panel G: Muslims

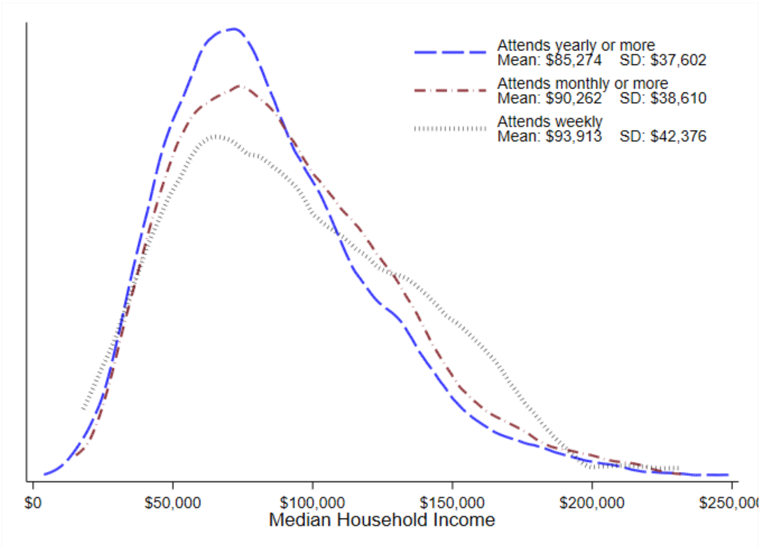

Panel H: Buddhists

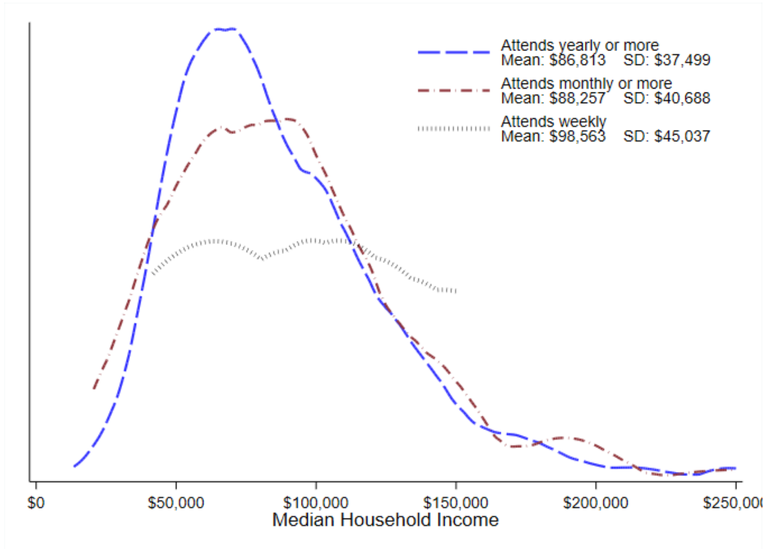

Notes: Each graph plots the distribution of median household income for yearly attenders (at least one religious worship visit during the sample period), monthly attenders (at least 11 religious worship visits during the sample period), and weekly attenders (at least 36 religious worship visits during the sample period). Each panel presents these graphs for a given religion.

Appendix Figure 9: Characteristics of Online Survey Sample

Panel A: Self-Reported Religious Attendance

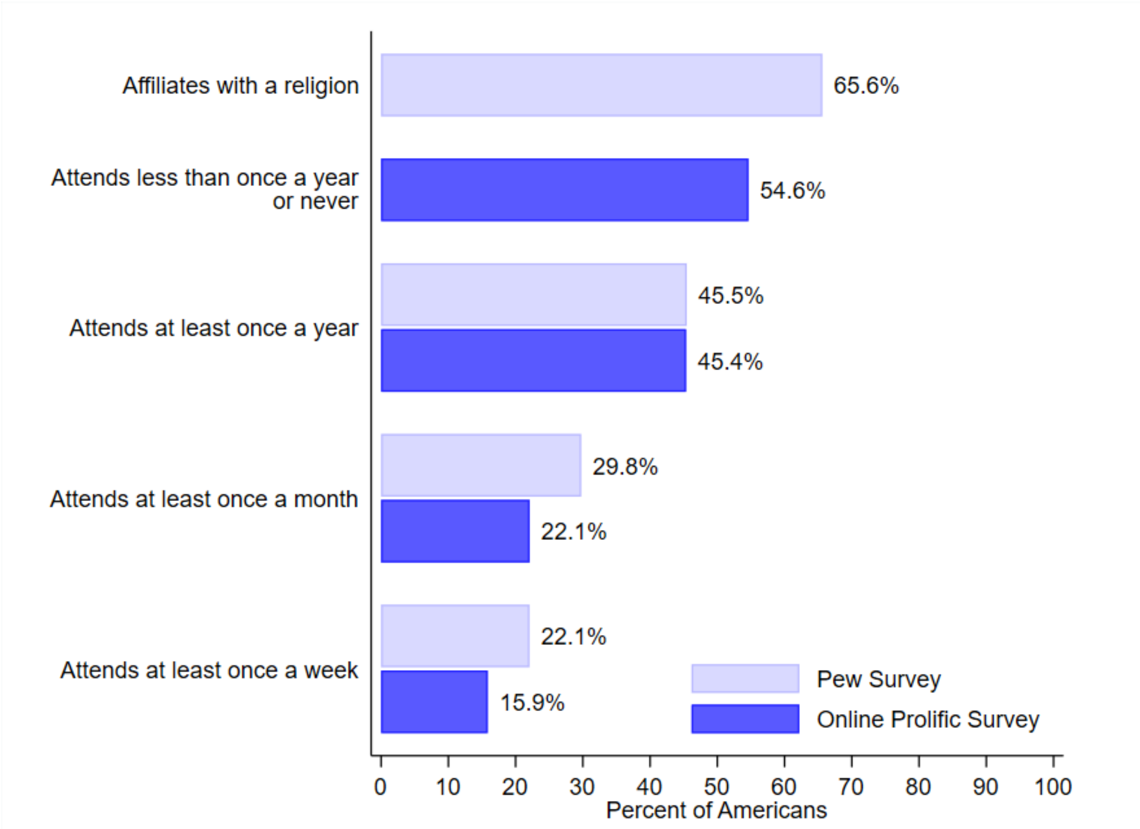

Panel B: Religious Affiliation

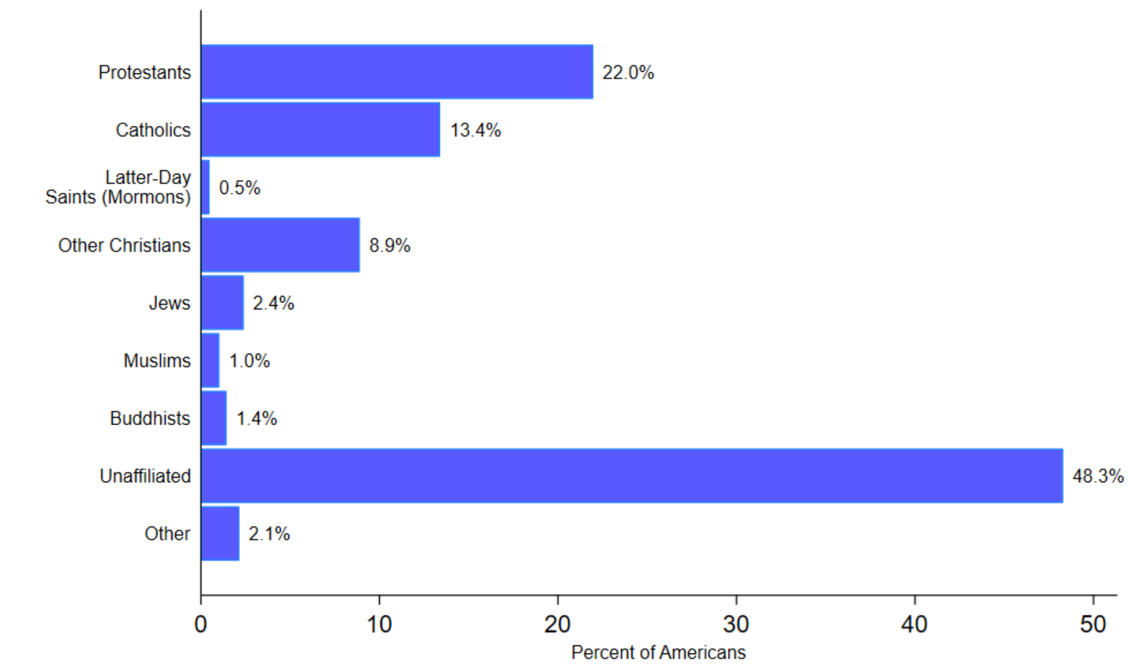

Notes: In Panel A the light bars show self-reported religious attendance and affiliation among Americans based on data from Pew Research. The dark bars show self-reported religious attendance among online prolific survey respondents. Panel B breaks down the percent share of different religious affiliations among online prolific survey respondents. For the prolific survey, Jehovah's Witnesses are part of the 'Other' category and Orthodox Christians are part of the 'Other Christian' category.