

Supplementary Materials for
Functional structures of US state governments

November 6, 2018

Contents

1	Data Collection Process	3
1.1	Listing Web Presence	3
1.2	Web Crawling	4
1.3	Alexa	6
1.4	Internet Wayback Machine	6
2	Known Dataset Issues	8
2.1	Web Malpractices	8
2.2	Web Dynamism	9
2.3	State Technological Level	11
2.4	Agency Outreach	11
2.5	Alexa	11
2.6	Intranet Activities	12
3	Government Activity Classification	13
3.1	The Procedure	13
3.2	Crowdsourced Validation	17
3.2.1	Setup	17
3.2.2	Results	18
3.3	Ubiquity of Government Functions	18
4	Building the Government Map	20
5	Hierarchies in Government Networks	24
6	Significance	28
6.1	In-Depth: Massachusetts	28
6.2	High Level: North Carolina, Arkansas, and Idaho	31
7	Variation Across States	32
7.1	Quantitative Variation: Scaling	32
7.2	Qualitative Variation: Function Implementation	34
7.2.1	Diversity Measure	35
7.2.2	Alternatives	36

8	State Difference Regression	38
8.1	Dependent Variables	38
8.2	Independent Variables	39
8.2.1	Ideology	39
8.2.2	Economic Structure	40
8.2.3	Income	40
8.2.4	Geographical Distance	40
8.2.5	Network Position	41
8.3	Robustness Checks	42
8.3.1	Statistical Tests	42
8.3.2	Alternative Ideology Measures	44
A	Government Function Classifications	49
A.1	US Census	49
A.2	Web Classification	50

Chapter 1

Data Collection Process

The data on which this paper is based, along with scripts allowing for the reproduction of all our main results, is available on the official paper website <http://govmaps.cid.hks.harvard.edu/> and on the Harvard Dataverse at <https://dataverse.harvard.edu/dataverse/govmaps>.

For each state we collect a list of public agencies including their website. For each agency, we: crawl its website, storing all the hyperlinks and the website’s textual content; report the county where its headquarters are located via Google Maps; collect from Alexa information about the number of web visitors; and obtain the year of creation of the website using the Internet Wayback Machine. Table 1.1 summarizes the basic information at the agency level in our dataset. The nature, acquisition, and pre-processing of the datasets is explained further below.

1.1 Listing Web Presence

The first step to download the web presence of a government is to list all the known public websites. We avoid relying exclusively on *.gov or state.* websites, as many government websites belong to a variety of different generic top level domains. Instead, we rely on the fact that every state in the US has a central government website¹. Each of these websites has a state agency directory,

¹E.g. <http://ohio.gov/>, <http://www.mass.gov/portal/> and <http://www.ncgov.com/>, date of access Oct 30th, 2014

ID	Agency Name	Agency URL	County	...
27438	Mass. Division of Local Mandates	http://www.mass.gov/auditor/	Suffolk	...
27510	Mass. Supreme Judicial Court	http://www.mass.gov/courts/sjc/	Suffolk	...
...	...	...	...	...

Table 1.1: Statistics about the data collection at the agency level.

where the central government lists all its agencies². We created a collection of Python scripts scraping these lists from all state directories, using the open source `BeautifulSoup` library. We integrated these lists to include:

- County and city governments: some states report the list of all their counties and cities in their main website³, while others do not. For the states which do not, we relied on the information collected in `http://www.statelocalgov.net/`.
- Public schools and school districts: some states list all their schools in their education department website⁴, while others do not. For the states which do not, we relied on the information provided by `http://elementaryschools.org/` and `http://high-schools.com/`.
- Higher education: we integrated the publicly available lists of public universities and community colleges with manual searches on Wikipedia⁵.
- Public libraries: we relied on the information provided in `http://www.publiclibraries.com/`.

During the collection, we created a unique ID for each collected URL. We also associated, where possible, an organization name to the URL using the information present in the web directory. We collected a total of 61,584 URLs. To verify the list, we created another Python script to check if a connection could be made to each of these websites. The number of URLs returning a non-empty non-404 response, or a valid redirect used to update the bookmark, were 48,044. This list was fed as input to our web crawler. For known bias in this pre-process procedure see Chapter 2.

1.2 Web Crawling

The main data collection phase was carried out with a Python web crawler, developed internally based on the `Scrapy` library. The crawler ran from November 5th, 2014 to December 15th, 2014, downloading a total of 33,768,263 pages.

The crawler takes as input a list of URLs. It downloads the content present at URLs up to 4 levels deep. We define this threshold because we are interested in capturing all the information all citizens can reasonably reach, and pages beyond the fourth level are usually hard to find. A page belongs to the fifth

²E.g. `http://ohio.gov/agencies/`, `http://www.mass.gov/portal/global-agency-list.html` and `http://www.ncgov.com/government/agencies/index.aspx`, date of access Oct 30th, 2014

³E.g. `http://www.delaware.gov/topics/municipalities.shtml` and `https://www.idahogov/counties/`, date of access Oct 30th, 2014

⁴E.g. `http://www.k12.wa.us/maps/SchoolURL.aspx`, `http://www.ode.state.or.us/pubs/directory/schooldirectory_20120904.pdf` and `http://w20.education.state.mn.us/MdeOrgView/search/tagged/MDEORG_DISTRICT_SCHOOL`, date of access Oct 30th, 2014

⁵`https://en.wikipedia.org/wiki/List_of_state_universities_in_the_United_States`, `https://en.wikipedia.org/wiki/List_of_community_colleges#United_States`




Figure 1.1: The schematics of the web crawler behavior.

level if the user needs to click on four consecutive links to access it, without any link in the chain being available at any previous level of the hierarchy.

The crawler is instructed also to download the pages to which the URL directly links. Aside from these direct links, the crawler is instructed not to leave the domain it is currently crawling; it will not crawl deeper into the linked site or follow any of its links to other sites. In practice, the crawler stores two different sets of pages: the *internal pages* are the pages belonging to the domain of the URL provided as input; and the *external pages* which are the pages directly linked by at least one internal page. All links present in an internal page are followed (up to the fourth depth level), while no link in an external page is followed. A depiction of the crawler behavior is provided in Figure 1.1.

The crawler stores two different types of information: state internal links and state external links.

State Internal Links

The aim of this information extraction phase is to have, for each website, the complete set of hyperlinks connecting it to the other websites of the state government.

Any web crawler has to parse the pages it downloads to find all the hyperlinks included in the page. It does so to find the other pages it has to download to complete its job. We customize this phase to classify the hyperlinks into three categories: website internal, state internal and state external.

The website internal are all relative and absolute hyperlink paths that point to a page in the same website that it is currently being crawled.

The state internal hyperlinks are those absolute paths (starting with `http`)

that point to a URL recognized as a state agency’s website. This is achieved by letting the crawler know what are all the state agencies URLs. The crawler tries to find a match in this list. If it does, it adds one to the counter of the edge going from the currently crawled website to the pointed website.

The state external hyperlinks are all the hyperlinks that fail the previous two checks. They are passed to the next data processing phase described below.

State External Links

The expected output of the last main information extraction phase is the complete set of hyperlinks of each website that connect it to websites outside the state government.

As described in the previous section, the crawler returns the hyperlinks that are not pointing to a page inside the currently crawled website, nor to a website that we know is part of the state agency directory. Our interest is in the website that is being pointed to, not the specific page or subdomain. Therefore, we aggregate the link to its top level domain. For instance, the link `http://www.michelecoscia.com/?p=888#respond` is reduced to just `www.michelecoscia.com`. We do so by relying on the `urlparse` Python library.

The state external hyperlinks are then counted with the same logic as the state internal hyperlinks.

1.3 Alexa

We also collect two other kinds of information on these websites. First, we evaluate how much page visits (traffic) the state government websites are generating over the web. We rely on the service provided by `alexa.com`. Alexa collects traffic information from a variety of toolbars installed in the browsers of many people; this information is available through the Alexa APIs⁶.

For each website, we download the daily traffic information from September 1st, 2007 to August 31st, 2014. We then aggregate this information monthly. The traffic information provided by Alexa is the “page views per million”. For example, if a website scores 42 for a particular day, it means that in that day it generated 0.0042% of the global worldwide page views.

1.4 Internet Wayback Machine

The second additional data collection phase has the objective of assigning to each website the year in which it appeared on the web. We cannot rely on Alexa information because it might include incomplete information for two reasons. First, few people had an Alexa-powered toolbar installed in the early days of the web. Second, it is difficult to disentangle a “zero” result from Alexa from a

⁶<http://aws.amazon.com/awis/>

“website not found” result. Below a certain traffic threshold Alexa considers a result equivalent to zero.

Instead of using a traffic-related measure, we rely on a direct visit measure. The Internet Archive⁷ provides a service called the Internet Wayback Machine. Using this service, it is possible to access to older versions of a website. The Wayback Machine provides an API interface too⁸. By specifying a particular timestamp, the API will return the snapshot taken of a given url that is the closest in time to that timestamp. In the response, it will report the timestamp of that particular view. We then query the API asking for the closest timestamp to January 1st, 1996. We extract from the response the year of the timestamp, which is the earliest version of the website (back to 1996).

⁷<https://archive.org/index.php>

⁸http://archive.org/help/wayback_api.php

Chapter 2

Known Dataset Issues

Obtaining data via web scraping necessarily introduces biases. Here we provide an overview of the biases we encounter in our data collection phases. For most biases we report our solution to support our view that none of these biases fundamentally questions our results.

2.1 Web Malpractices

Web design is not a trivial process. People started designing websites before there were widely accepted standards. New technologies might introduce unexpected or unparsable elements in the html code of a page. Finally, the burden of showing the information correctly on the web is usually mostly on the client side. This makes browsers really good in interpreting nonsense html markup and frees the webdesigner to be consistent in its usage of the available technology. In this section we list the major hurdles our crawler has to face and how we address them.

First, not all agencies have their own url domain. Many completely independent websites are hosted in a subpath of a parent agency, or even the central government itself. This makes the crawling difficult because we have to correctly detect if a subpath is part of a different agency website or not. We address this issue by checking the full path – not only the domain – of the direct link to the agency website with regular expressions. We consider a hyperlink to be internal to the website if both domains match and the full hyperlink path matches the regexp `.*main_path.*`. For instance, the page `http://www.mass.gov/mdaa/policy-and-legislation/where-we-stand.html` will match `mass.gov/mdaa` instead of `mass.gov`, because the `.*mass.gov/mdaa.*` rule has priority over `.*mass.gov.*`.

Second, there are a number of issues with the html markup of many pages. A collection of issues includes, but it is not limited to: multiple end tags (`</html>`) in the same page, tags not nested according to the XHTML standard¹, usage

¹<http://www.w3.org/TR/2010/REC-xhtml11-20101123/>

of deprecated tags to specify hyperlinks such as `<area>` images instead of the standard `<a>`, usage of the deprecated `<frame>` page structure that effectively hides the html code, using Javascript links instead of the href attribute of `<a>` tags, and so on.

The parsing issues are solved by using the computationally expensive but very lenient parsing library `html5lib`. `html5lib` can handle multiple end tags and non-standard markup. The non-standard linking tags are solved by customizing the hyperlink extraction logic of the crawler. In this way, we do not miss hyperlinks that are included in `<area>` and `<frame>` tags.

Finally, we note that it is theoretically possible to use `Scrapy` to extract Javascript links, via the `process_value` argument of its `LinkExtractors`. However there is no standard Javascript way to create hyperlinks, so we would have to develop one ad hoc solution for each case. As these links are very infrequent, we decide to ignore them.

2.2 Web Dynamism

The WWW is a highly dynamic environment. There is no guarantee that a piece of content accessed today is going to be available tomorrow. This issue affects our study in two ways: our results could be not reproducible and our results could contain decisive biases.

The result reproducibility is the most important issue. Given that websites can come and go much more rapidly than agencies, how sure can we be that the government's internet footprint is an accurate approximation of its actual structure? In the main paper we examine the evolution of the government web presence extracted from the Internet Wayback Machine, noting that its growth rate has slowed to less than 2% annually. This suggests that the government's web presence is saturating, rather than evolving steadily or at an accelerating rate: most agencies that are going to represent themselves online have already done so. In Figure 2.1 we supplement this conclusion with timelines for all 50 governments. From the figure, we can see a universal tendency for state governments' Internet presence to grow at a decreasing rate. Undoubtedly there is still change in the government's web presence over time, with websites coming and going, such that at the very micro level our results are specific to the timing of the crawl and cannot be precisely replicated later. But the overall general patterns are very likely to be robust to these local errors.

Web dynamism could also introduce some additional biases. Websites can move location: it is possible that a hyperlink in a website is pointing to an old URL. This could introduce a bias because a state internal hyperlink pointing from one government agency to another might be misclassified as a state external link. Since domain changes are unpredictable, we are only able to catch this event if the pointed URL has a particular structure (namely, it is a `.gov`, or it belongs to a `.state.*` domain). In all other cases, we are unable to catch it. However, most broken links belong to unmaintained and usually difficult to find pages at the peripheries of the government's web footprint. In addition,




Figure 2.1: The timelines of each state government as extracted from the Internet Wayback Machine. In each plot, the x axis is the year and in the y axis we report the cumulative number of websites.

this bias would lead us to undercount state internal links and overcount links external to the state; in the paper we note that 99.09 percent of all hyperlinks are between agencies in the same state, suggesting that this source of bias is minimal.

The change in location affects also our wayback machine results. We know the most recent url for a state agency website, but that agency could have had an earlier website at a different url. But again, this problem does not affect our previously exposed argument about the government growth rate. In fact, if there is a bias, it is operating against our argument increasing the growth rate, because the moved url will appear in our data as a newer agency.

2.3 State Technological Level

The level of technical competence in a state could also influence our results. A state with highly competent technical personnel may be more likely to have a website for each of its agencies, including smaller schools and libraries. Computer usage also varies, and states with fewer connected citizens may not perceive as great a need to create websites in the first place. While we note these potential sources of bias, our earlier results suggest that the growth of government web presence has slowed so dramatically that there appear to be few agency late-comers, and variation in the technical competence of their designers across the United States should not be so high to make their websites incomparable to each other.

2.4 Agency Outreach

Different state agencies are created with different objectives and functions in mind. There are agencies that have as a main focus the outreach to the public. For example, we expect an agency focused on tax collection to have a large website, including all information for all fiscal categories about where, when, how and how much to pay taxes. Other agencies might instead have just a single webpage, barely notifying citizens of their existence. For this reason the number of pages crawled and the website traffic are not a good indicator of agency size. However, we expect these tendencies to be relatively consistent within the same function, enabling interstate comparisons between the websites of tax collection agencies.

2.5 Alexa

To collect traffic information we rely on the information provided by Alexa. Alexa rarely collects information about subdomains: if we request `http://www.mass.gov/courts/court-info/sjc/`, Alexa will return the traffic information for `mass.gov`. Likewise, `http://elmwood.millburyschools.org` is resolved by Alexa into `millburyschools.org`.

For this reason, we do not use Alexa’s information directly. We instead split the traffic information equally among all the domains sharing the same results. This introduces the bias that overcounts page views for most of the subdomains and undercounts page views for the main domain, which is likely the one generating most of the traffic. To account for this bias, we only use Alexa traffic data to understand the traffic of general (not specific) functions, which are less likely to be located within the same main domain.

2.6 Intranet Activities

A state might have a whole range of activities that could be reported beyond a login screen, as well as activities for which the government believes that secrecy is necessary and thus does not reflect online. Our crawler cannot access internet resources that are not openly accessible. Thus our results reflect only the structure of activities that is "visible" online to citizens without special credentials.

This bias may be significant. Yet there are reasons to think that it is less of a problem at the level of aggregation we employ in our analyses. Many functions whose specifics are concealed may still be explained in a publicly accessible page; in this case we do not miss the existence of the function or agency, only some of its density (for which our measure employs the page count, and will therefore underestimate the density of sites with substantial pages behind logon screens). We also expect not to find beyond a login screen external links to other state agencies that cannot be found in any other part of the website. Finally, we expect the concealment of activities to be much more of an issue at the federal level than for state governments.

Chapter 3

Government Activity Classification

To determine the functional structure of state governments reflected in their web footprints, we need to classify the main activity of each government agency: to define a list of government activities and to connect each agency to one of these activities. In this section, we present the challenges we face in classifying agency functions, the procedure we define, and our validation strategy, which relies on crowdsourcing.

3.1 The Procedure

In principle, an agency should be classified as performing a given function if the content of its website reflects activities related to that function. This task is similar to document classification, a task in natural language processing. Given the website content, we can cluster together websites if their content is sufficiently similar.

We can do so by using a variety of metadata about the website: the URL and the agency name. One standard approach is to perform latent semantic analysis [1], which is fully automatic. However, a fully automatic strategy failed for a number of reasons.

First and foremost, the content of many pages—and their titles—employs non-standard language or is otherwise unreliable. Automated techniques will only perform as well as the language they examine is consistent. Second, the language-based differences identified by automated approaches frequently do not align with distinctions of substantive social science interest. For example, there are very specific and standard vocabularies around primary and secondary school calendars and athletics; while there are wide-ranging and inconsistent vocabularies around many health and social services. Automated approaches consistently resulted in a “school calendar” topic and a “school athletics” topics (which are not of distinct substantive interest), while frequently failing to

distinguish between social service functions that are both of interest to, and discernible by, social scientists.

Third and relatedly, the differences between functions that are meaningfully distinct to humans can be too nuanced for machines to consistently differentiate. For example, labor and industrial relations boards that are independent, quasi-judicial agencies and adjudicate collective bargaining disagreements should be functionally distinct from generically-titled Departments of Labor and Industrial Relations. Similarly, a public service commission may refer either to an independent advocacy and regulatory agency for utilities provision, or it may refer to volunteer work; the Office of Public Counsel, or the People’s Counsel, may refer to utilities regulation, or to legal public and indigent defense services. Automated approaches cannot consistently distinguish between such functions.

Finally, many agencies perform multiple functions (e.g., a university may have its own library, a corrections facility may have its own police force). We aim to classify agencies according to their primary function, similar to the method employed by the U.S. Census (Higher Education: Universities and Colleges and Public Safety: Corrections and Correctional Facilities, respectively). Automated approaches cannot reliably identify primary versus secondary functions from webpage text.

In the face of these hurdles, even semi-manual text cleaning and semi-supervised approaches did not produce a satisfying classification of government functions. As a result, we employ an inductive, supervised approach to classify node function.

First, we generate the most common bigrams and unigram keywords from the node titles. Second, we use these keywords as a guide to develop an initial nested set of 169 categories that substantively extend the existing Census categories of government functions. These 169 categories are nested within 29 broader category groups (e.g. the broader category group “Education” contains both “Higher Education: Universities and Colleges” and “Primary and Secondary Education: Schools and Districts”). Third, we use the major keywords associated with each functional classification to preliminarily categorize websites.

Fourth, and finally, we review each category and manually reassign all agencies that should be differently categorized—a process that required adding and merging several categories. In particular, we add categories when more than 10 websites appeared that, when reviewed, were substantively different from any existing category (typically the result of high variance in the language used for website titles). Among the categories added through this process were: additional financing authorities (Educational Facility Financing Authorities, Health Facility Financing Authorities, and Student Loan and Tuition Financing Authorities), Administrative Law Hearings and Appeals, and Utilities Regulation (as distinct from Utilities Provision). Categories were merged when a substantial number of websites used both categories, such that the number of agencies that appeared in just one category fell below 15. Through this process, for example, we combine the Probation and Parole categories with the Pardons

category and the Gaming Regulation category was combined with the Racing category.

Overall, this process resulted in 166 functional classifications, with 28 broad category groups. Table 3.1 reports the resulting classification. The Census classification is included in the Appendix of this document.

Administrative Law General Appeals and Hearings; Civil Service Appeals and Hearings; Tax Appeals and Hearings; Workers' Compensation Appeals and Hearings.
Audits, Inspectors General & Accountability
Boards & Professional Regulation General Licensing; Architects, Landscape Architects & Interior Designers; Athletic Trainers, Occupational Therapy & Physical Therapy; Audiology, Speech & Language Pathology; Barbers & Cosmetology; Chiropractic; Construction & Contractors; Counselors, Therapists, Psychologists & Social Workers; Dental; Engineers & Land Surveyors; Geologists and Soil Scientists; Massage Therapy; Medical Examiners; Medicine; Nursing; Nursing Home & Assisted Living Administrators; Optometry; Pharmacy; Alcohol, Drug, & Gambling Counselors; Private Investigators; Real Estate Appraisers; Veterinary.
Business Regulation General; Alcohol & Tobacco Regulation; Athletic & Boxing Commissions; Cemeteries; Land Use & Building Codes; Lotteries; Gaming & Racing Commissions; Real Estate.
Commerce and Economic Development General; Consumer Protection and Education; Film & Television Promotion; Housing and Community Development; Improvement Districts; Industrial Development; Minority & Women's Businesses; Small Businesses.
Criminal Justice General; Juvenile; Pardons, Parole & Probation; Sentencing; Sex Offenders Registry; Victims Advocate.
Data & Information General; Geographical Information Systems; Health & Vital Statistics; Information Technology & Communications Technology.
Education General; Adult Education; Higher: General; Higher: Universities and Colleges; Libraries; Primary and Secondary: General; Primary and Secondary: Career & Technical Centers; Primary and Secondary: Early Childhood Centers; Primary and Secondary: Schools and Districts.
Environment Agriculture, General; Agriculture, Product Promotion and Marketing Boards; General; Conservation & Protection, General; Coastal Protection; Fish, Game & Wildlife; Flood Control & Levees; Forestry; Land, Soil & Water; Parks and Recreation; Rivers; Energy; Minerals and Mining; Oil and Gas; Power Agencies & Authorities.
Ethics, Lobbying & Campaign Finance Commissions
Financial Administration

General & General Regulation; Banking; Budget Offices; Credit Unions; Fiscal Security; Insurance & Risk Management; Purchasing; Securities.
Financing Authorities General; Educational Facilities; Health Facilities; Housing and Community Development; Student Loans & Tuition.
General Administration Arts & Humanities; Arts & Humanities: Public Broadcasting; General; Public Buildings & Facilities; Secretary of State; Volunteering & Community Service; Elections; Municipal and County Administration; Governors & Lieutenant Governors.
Health & Human Services General; Public Health, General; Children and Families: General; Children and Families: Child Services, General; Children and Families: Family Services, General; Children and Families: Child Abuse Prevention; Children and Families: Child Advocate; Children and Families: Child Support; Deaf, Hard of Hearing & Blind; Developmental Disabilities; Mental & Behavioral Health, Substance Abuse; Rehabilitation, Independent Living & Disability Services; Elder Care, Aging & Long Term Care; Hospitals & Medical Centers.
Historic Preservation General; Archives & Records; Museums.
Homeland Security
International Trade
Judicial and Legal General; Attorney General, Solicitor & District Attorneys; Civil & Human Rights; Courts, General; Public & Indigent Defense; Supreme Court.
Labor and Human Resources General; Labor & Industrial Relations Boards; Occupational Safety & Industrial Accidents; Personnel, Human Resources & Jobs; Workforce Development & Improvement.
Legislative
Military State Guards; National.
Public Safety General; Corrections & Correctional Facilities; Emergency Management & 911; Fire Protection; Office of the Medical Examiner; Police, Sheriffs & Patrol Offices.
Research & Science General; Geological Surveys; National & Federal Laboratories.
Social Services General; Health Care & Insurance Exchanges; Indian & Native American Affairs; Medicaid; Minority Affairs; Retirement; Social Security, Social Insurance Administration, Medicare; Unemployment Insurance & Employment Security; Veterans Affairs; Women's Affairs.
Tourism & Travel General; Conventions & Fairs.
Transportation

General; Air Transportation, Aviation & Airports; Highways & Roads; Mass Transit; Motor Vehicles; Port Authorities.
Treasury & Revenue General; Investments.
Utilities Oversight and Adjudication; General; Sanitation; Waste Management.

Table 3.1: The two level functional classification of government activities. In each cell, the first line in bold represents the broad category group (first level of the classification). The following list includes the functional classifications (second level) separated by semicolons.

3.2 Crowdsourced Validation

Because our classification is inductive, we verify that it was not idiosyncratic to our understanding of the functions agencies reported undertaking in their websites. Our validation strategy is to prove that an untrained observer, when faced with the task of classifying the function of a government agency, is likely to agree with our classification. In this section we describe the setup of the validation procedure and its results.

3.2.1 Setup

We relied on the services provided by the Amazon Mechanical Turk¹. In the Mechanical Turk, one can assign Human Intelligence Tasks (HITs) to the pool of current workers. Our setup follows closely the best practices used for behavioral research [2]. In our case, a HIT consisted in classifying five government agencies. The worker has to provide her best guess based on the agency name and a screenshot of the agency website, which we collected during the web crawling.

For each agency, the observer was given seven options. Four options were randomly selected functions, one option was the function we assigned to the agency. These five options were scrambled randomly. Two extra options were provided if the observer felt the agency was not classifiable under any of the five options: “None of the above” and “Not sure”.

One of the five agency included in each HIT was a control agency. We selected as control agency schools, because they are the most easily classifiable. For each observer, we record their accuracy in classifying schools. We discard answers from observers whose school classification accuracy is lower than 90%. The position of the control question in the list of five questions per HIT is random.

Each HIT is performed by three independent observers. This implies that every agency-function connection is validated by three observers. Given pricing constraints, we are able to classify 9,061 agencies out of our entire set.

¹<https://www.mturk.com/>

This allows us to validate fully all functions aside from schools, libraries and county/city administrations, which should be the most easily classified, and for which we validate only a sample.

3.2.2 Results

We evaluate the results of the validation process by calculating raw agreement measures. The first measure is the worker quality distribution. We want to know how each observer performed on the control agencies. There was not a observer worker who failed to pass the 90% test. The minimum observer accuracy on schools was 90.19%, while the best observer had an accuracy of 97.35%.

The second evaluation criterion is a inter-worker agreement measure. With this measure we aim at understanding how hard the classification task was for the observer. In practice, we report the share of agencies for which at least two out of three observers agree on the answer. This share equals to 98.71%. The high value means that the classification task was not difficult for the observers.

The third evaluation is the observer agreement with our agency classification. We count the number of agencies for which at least two observers agree with our classification. This is necessarily lower than the inter-observer agreement, because we can only discard agencies. In fact, this measure equals to 94.98%. Interestingly, if we discard all HITs in which the observer mis-classified the control agency, we obtain an agreement of 95.25%. Note that for many of the $\sim 3.5\%$ agencies where the observers agree with each other and disagree with our classification, the agreement between the observers is on the “not classifiable” options.

From these evaluations we can conclude that observers performed well on average, that the classification task was not hard, and that the observers overwhelmingly agreed with our classification. However, the reported statistics do not allow us to know if the average observer agrees with our classification more than the observers agree with each other. To do so, we use the Fleiss’ Kappa, which is a standard test to evaluate for inter-rater reliability in categorical ratings [3, 4]. We discard all agencies for which at least two out of three workers responded “None of the above” or “Not sure”, or any agencies for which all three observers disagree.

When scoring the agreement among observers, we obtained a Fleiss’ Kappa $\kappa = 0.89$, which confirms the high degree alignment of the observers. When scoring the agreement of the average observer (when at least two out of three observers agree) with our classification, we obtain a higher score, $\kappa = 0.945$. We can conclude that our classification indeed captures the agreement between observers better than any single observer’s.

3.3 Ubiquity of Government Functions

Not all functions are performed by all states. Figure 3.1 displays a histogram showing, for each of the 28 functions composing the top level of our classifica-




Figure 3.1: For each of the 28 functions, the figure reports the number of states with at least one agency performing the function.

tion system, the number of states in which at least one agency is performing the function. Only seven out of the 28 functions that are performed by all states—primarily core function such as the judiciary or the legislature. But most remaining functions are implemented by most states: for example, 24 functions are performed by at least 40 states. Only a few functions are rare: research and science; homeland security; ethics, lobbying, and campaign finance regulation; and international trade.

Chapter 4

Building the Government Map

We model the structure of government as a network. In this network, the nodes are the government agencies. We connect two government agencies with an edge if they directly link to each other through at least one hyperlink in their websites.

The edge creation process in the structural map is ruled by the decision of a government agency to link up to another government agency. If such link is present, this means that the agency thinks that the readers of its website might want or need to read information contained in the other agency's website. Each link can be considered as a vote cast locally by an agency, making its relations with another agency explicit. The collection of all these local decisions made by individual government agencies reveals the overall network structure of the government.

Figure 4.1 (top) depicts the structural map of the government of Ohio. The structural map of a government is a directed network, given the asymmetry of web hyperlinks – one agency can connect to another even if that latter agency does not connect to it. In the network, the node size is proportional to the in-degree of the nodes: we think that an agency's size and importance is proportional to the number of agencies that link to it. The color of the node is proportional to its rank, calculated using betweenness centrality, from orange (very central node) to blue (peripheral node). Betweenness centrality aligns the degree of a node with its structural importance in the network: a node is more important if the nodes it links to are otherwise unreachable following the edges of the network. The edge size is proportional to the number of links between the two websites. The edge color is proportional to its relative strength. An edge is strong if the source website links to the target website in many of its pages. Hence the color is proportional to the ratio between the number of links over the total number of pages of the website. Strong links are reported in orange and weak links in blue.




Figure 4.1: The structural network for Ohio including all links (top) and only the links significant at $p < .1$ according to [5] (bottom).

Figure 4.1 (top) reports every single observed hyperlinks. The underlying structure is made clearer by selecting only those links which are unusually strong. We do so by including only links that pass a test of statistical significance, such as that described in [5]. With this test, we select only the edges that are significantly stronger with $p < .1$ than the average tendency of agencies to connect. The result is depicted in Figure 4.1 (bottom). Given the lower number of nodes and edges, we label the most important nodes of the network with the agency name.

Figure 4.2 is a full page expansion of Figure 2 in the main paper, showing the aggregated online network including all agencies from all states. It then compares the specific state views of three states: Georgia, South Carolina, and Massachusetts. Finally, it shows the aggregated functional network, where nodes are government functions rather than agencies, and links are the sum of all links from all 50 states connecting the various functions (more details in the paper).




Figure 4.2: (A) The full web network. Each node is an agency and connections are established if two agencies linked on the web. (B) The extract from the full network including only Massachusetts agencies. (C) The aggregated functional network: each node is a government function and they are connected if agencies performing the two functions connected on the web. Aggregation of all 50 US states. (D) Comparison of the functional network using only links from Georgia, South Carolina, and Massachusetts. 23

Chapter 5

Hierarchies in Government Networks

In the paper, we claim that the network of hyperlinks among government functions is strongly hierarchical. In this chapter we provide additional evidence for this point.

How to determine if a directed network exhibits a clear hierarchy is an open problem in the network literature. There are different definitions of “hierarchy”: for example, order, nested, and flow [6, 7]. Here we focus on flow hierarchy, because order hierarchy is a form of simple node ranking – each node is associated with a score – while nested hierarchy seeks higher-order structures that fully contain lower order structures – e.g. communities. In a flow hierarchy, nodes in a higher level connect to nodes at the level directly beneath, and can be seen as managers spreading information or messages to the lower levels. This is the most commonly understood definition of hierarchy.

The first test to determine whether a directed network is hierarchical is to test for reciprocity. High reciprocity means that, if a connects to b , then b often links back to a . If that happens often, it is impossible to tell whether a is on a hierarchically higher level than b . As reported in the main paper, the level of reciprocity in the functional government network is low ($\sim 24\%$) and, more interestingly, this is higher when looking at connections inside the functions ($\sim 35\%$) – agencies performing the same functions connecting to each other – than across functions ($\sim 16\%$). This is an initial proof of functional hierarchy: agencies reciprocate connections when they perform the same functions – i.e. the intra-function connections are horizontal, non-hierarchical. However, when two agencies performing different functions connect, it is less likely that the connection will be reciprocated, as expected if connections flow from functions on one level of the hierarchy to another.

Different functions show different reciprocity levels. In the previous paragraph we make a distinction between two reciprocity measures: between the websites of agencies performing the same function and those performing differ-

ent functions.

Table 5.1 reports the value distributions for both versions of the reciprocity measure. As expected from the higher average (35% vs 16%), reciprocity is higher within the function than with other functions for all functions except the last two. As discussed in the paper, Boards and Professional Regulatory agencies are an outlier with unusually high reciprocity with other such agencies (71%). Education connections to other functions are the most reciprocated, but the proportion (25%) is lower than the average reciprocity within education – and, indeed, lower than reciprocity within almost all functions.

Second, we calculate the fraction of edges not participating in cycles in a directed graph: the higher this fraction, the more a network is hierarchical [8]. This fraction is $\sim 56\%$ in the functional network, meaning that almost half of the edges are not part of cycles. The expected fraction of edges outside a cycle in a corresponding null network is only 5%.

Our third method, “global reaching centrality” (GRC), is quantified by identifying the node which can reach the largest fraction of nodes in the network via outgoing edges [6]. The more distinct this root node is from the rest of the network, the more the structure is hierarchical, because there is one node “commanding” it. We compare the GRC of the functional network with the result one would get from a randomized version of the network, such as one depicted in Figure 3B in the main paper. We find that the randomized networks have significantly lower hierarchical values ($p < 0.05$).

The fourth method was developed by the authors of this paper [9]. It is quantified by reducing the directed network to a perfect hierarchy (an arborescence) by collapsing to a single node all strongly connected components and deleting edges when a node has more than one parent. The fraction of edges preserved at the end of this reduction over the original number of edges in the network is the hierarchy score. We find that $\sim 2\%$ of the edges in the functional network are preserved. This is significantly higher than what one would expect in the null model ($p < 0.001$).

None of these four tests by itself is perfect by itself, but their complete agreement allows us to conclude that the network of hyperlinks between government functions is very likely hierarchical. Figure 5.1 shows the hierarchy of the functional network. It is a reproduction of Figure 3 in the main paper, adding the function labels on the x-axis.

Function	Within The Function	With Other Functions
boards & professional regulation	0.7186	0.2189
transportation	0.4895	0.1831
military	0.4762	0.1426
labor and human resources	0.4664	0.1549
data & information	0.4500	0.1682
legislative	0.4455	0.1096
financial administration	0.4390	0.1500
health & human services	0.4266	0.1599
environment	0.4205	0.2092
audits, inspector general & accountability	0.4000	0.1002
treasury & revenue	0.4000	0.1154
utilities	0.3902	0.1796
tourism & travel	0.3750	0.1764
judicial and legal	0.3727	0.1746
commerce and economic development	0.3543	0.1876
social services	0.3478	0.1220
financing authorities	0.3478	0.1403
general administration	0.3440	0.2031
historic preservation	0.3421	0.1394
education	0.3404	0.2566
criminal justice	0.3396	0.1232
public safety	0.3340	0.1744
administrative law	0.3256	0.1779
research & science	0.2857	0.1726
business regulation	0.1905	0.1355
ethics, lobbying & campaign finance commissions	0.0000	0.0610
homeland security	N/A	0.1485
international trade	N/A	0.0692

Table 5.1: Directly reciprocated connections within and across functions. The table is sorted by direct reciprocity within functions. The last two functions have undefined value, because there is no edge in the network connecting them to other agencies performing the same function.




Figure 5.1: (A) The hierarchical layout of the functional network. Nodes are functions connected if two agencies performing them connect to each other. Centrality of the Level 1 function determines the position on the x-axis (from the middle out, categorical), centrality of the Level 2 determines the y-axis position (from top to bottom, logarithmic). (B) The same layout applied to a null model shows lack of hierarchical organization.

Chapter 6

Significance

This work is based on a fundamental assumption: the structure of the government as extracted from the online interactions of its agencies is a faithful representation of the actual structure of the government, one detailed enough to provide insights about its *de facto* interactions that are not captured in other ways - for example by examining its *de jure* relations. In this section we prove this assumption.

6.1 In-Depth: Massachusetts

First, we make an in-depth analysis of the structure of the Massachusetts state government. We first create the mandated (*de jure*) network of Massachusetts by manually parsing the structure as described in the Massachusetts General Laws, noting agencies established by statute and connecting agencies when they are mandated to interact by statute. We note four possible ways in which agencies are mandated to interact: 1) if one has direct oversight over the other, 2) if one is supposed to give advice to the other, 3) if one has to ask permission of the other, and 4) if officials of one agency serve on a committee with officials from the other.

To compare the mandated network of Massachusetts with the Web network, we first match the nodes. For 24% of the agencies in the mandated network we find no website. Most agencies without a website are small advisory committees and groups¹. For the remaining 76% of agencies we either find a perfect match, meaning that the agency has its own website², or a partial match, meaning that the agency's website is actually included in a parent website³. In 4.5% of the cases we have neither type of match and we drop the agency.

¹Examples are “Energy advisory board” and “Watershed lands preservation committee”

²For instance, at the time of writing this article the office of the governor has its own website at <http://www.mass.gov/governor/>

³E.g., agency “Boston municipal court department” has <http://www.mass.gov/courts/court-info/trial-court/bmc/> as its website, which we match to <http://www.mass.gov/courts/>




Figure 6.1: The overlap between the constitutional and the crawled network. (Left) The distribution of matching and non-matching weight ratios for the null models. For each ratio value we count the number of null models scoring that particular value. The black line reports the value of the observed ratio. (Right) The number of hyperlinks per node pair at different path distances in the constitutional network.

We now turn to matching the edges. The expectation is that agencies that are connected in the mandated network should be more strongly connected on the web. To prove this, we implement the following simple algorithm. We first consider all Web connections. A Web connection is, for instance, between the office of the governor and the office of Energy and Environmental Affairs. The latter has 30 hyperlinks linking to the former, thus the edge has weight 30. If the edge is also in the mandated network we label the edge as “matching”, otherwise we label it as “non-matching”. Around 35% of the mandated edges are matching. Once we label all edges, we calculate the average edge weight for both matching and non-matching edges. The ratio between the two is around eight, meaning that between connected agencies in the mandated network there are 8 times more hyperlinks than agencies that are not directly connected by statutory mandate. If as edge weight we use not the total number of links but the average link per webpage, the ratio is around six.

These results indicate that mandated connections are very strongly reflected in the web network. To estimate the significance, we compare the observed values with one obtained from randomly generated networks. In our null models, the mandated network’s edges are reshuffled randomly, maintaining both the in- and the out-degree distributions (all agencies are still connected and connects to the same number of agencies). We then recompute the matching and non-matching edges. We create 10,000 null models; the distributions of the matching non-matching ratios is reported in Figure 6.1 (left). Out of the 10,000 null models, only one outscored the observation, giving us a confidence associated with $p = 0.0001$.

We can distinguish between the different types of mandated edges as presented before. Table 6.1 reports the ratios calculated for the four mandated edges types. The strongest overlap between mandated and Web network is ob-

Type	Ratio	% Edges	p-value
All	5.9698	100%	0.0001
Direct oversight	1.8640	40.8%	0.0313
Advice	8.6565	36.9%	0
Permission	0.2942	3%	0.844
Committee	1.2407	19.3%	0.0625

Table 6.1: The ratio between matching and non-matching edge weights, split across the different constitutional edge categories. In the last column, we report the relative size of each edge category in the constitutional network.




Figure 6.2: The relationship between the betweenness centrality of Massachusetts agencies in the Web (x-axis) and mandate (y-axis) networks.

served for the advice relationship. When two agencies are connected with a “permission” type of link, the edge weight ratio is below one, meaning that there are fewer hyperlinks between agencies connected this way. This fact does not hinder the overall result, as the number of permission edges is very low, around 3% of the mandated network.

We further hypothesize that the closer two agencies are in the mandated network, the more they should link to each other. In other words, agencies directly connected in the mandated network should have strong connections in the web network, while those at three or more degrees of separation should not be very related. We test this hypothesis by counting the number of hyperlinks established between websites of agencies at a given path distance in the mandated network. The result is depicted in Figure 6.1 (right). From the Figure, we see that agencies at distance one (directly connected) have an average of a hundred hyperlinks between them, and the number of hyperlinks decreases exponentially as the distance in the mandated network increases. For agencies at three degrees of separation in the mandated network, the average number of hyperlinks is lower than one. At distance five, the diameter of the network, there is only one hyperlink per hundred node pairs.

Finally, we hypothesize that each agency that is central in the mandated network is also central in the reconstructed Web network. We test this hypothesis by calculating the betweenness centrality of the nodes both in the mandated and

in the Web network. We then calculate the correlation between the centrality vectors. There is a strong linear relationship between the two, with $R^2 \sim 55\%$. Figure 6.2 depicts this relationship as a scatter plot. Most nodes cluster on low values: these are the peripheral and less important nodes we allow to move around. When looking at the non-zero centrality nodes, we see a significant relationship.

6.2 High Level: North Carolina, Arkansas, and Idaho

As a second analysis, we further investigate the high-level structure of the state agencies of three states: North Carolina, Arkansas, and Idaho. In this case, we retrieve each state’s organizational chart⁴⁵⁶. The organizational charts report only the main departments of a state government. For each of these departments, we look if it is present in our node list.

For all three states we count the number of total departments and the number of the departments with a matching node in our network. We obtain the following 1-to-1 match levels:

- North Carolina: 33 of the 38 official governing bodies have a 1-to-1 match.
- Arkansas: 36 of the 39 official governing bodies have a 1-to-1 match.
- Idaho: 32 of the 34 official governing bodies have a 1-to-1 match.

In some cases, the lack of a 1-to-1 match is due to the fact that the web data is more granular than the organizational charts. For instance:

- North Carolina, of the matches that are not 1-to-1, two are 1-to-many: the Department of Natural and Cultural Resources is split into two nodes (Natural Resources, node ID 44418, and Cultural Resources, node ID 43967); the Department of Environmental Quality is split into two nodes (Air Quality, node ID 44199, and Water Quality, node ID 43975).
- Arkansas, of the three not 1-to-1 matches two are 1-to-many: Boards and Commissions, and “Separate agencies”.

The remaining cases are either many-to-1 match in which the web data is less granular than the organizational chart (e.g. the Court of Appeals of North Carolina shares its node with the Administrative Office of the Courts), have websites but they were not part of our seed list because they were not linked by the central government (e.g. the Legislative Audit of Arkansas), or do not have a website (e.g the Idaho Department of Self-Governing Agencies).

⁴<http://digital.ncdcr.gov/cdm/ref/collection/p249901coll122/id/704114>

⁵https://ballotpedia.org/File:Arkansas_exec_org_chart.png

⁶<https://legislature.idaho.gov/wp-content/uploads/budget/publications/Legislative-Budget-Book/2016/0.Front-End/-32.Organizational%20Chart.pdf>

Chapter 7

Variation Across States

State governments vary both quantitatively and qualitatively. In this chapter we explore both dimensions. First, we explore the quantitative variation, e.g. scaling: how does the government of a state change when the state increases in size? We use population as a measure of a state's size. Then, we explore the qualitative differences: how much the implementation of each government function differs across states.

7.1 Quantitative Variation: Scaling

To explain why states' governments differ, one possibility is that they reflect differences in the numbers of citizens state governments are responsible for governing. Some state functions, such as the governor or the state legislature, are typically performed by a small number of state bodies in both large and small states, but others such as education, health care, social services, court systems, public safety, and many other government functions involve highly localized implementation and frequently face-to-face contact with citizens, necessitating implementation by agencies that are more numerous, complex, or interconnected in more populous states. Thus to some degree government, like most other complex social and biological systems [10, 11, 12, 13], should scale.

To examine scaling of government, in Figure 7.1 we plot each state's population against the scale (panel A), density (panel B), and the number of hyperlinks (panel C) of its network, across all functions. We find that state governments as a whole scale on all three metrics: more populous states have more state agencies and those agencies appear to be more complex and more interlinked. But the scaling is mostly sublinear, implying that the governments of larger states realize some economies of scale: state population explains 54% of cross-state variance in the total number of state agencies, 45% of the variance in their density, and only 15% of the variance in links.

We find that only two of the 28 general functions are responsible for most of this scaling: education and municipal and county administrations, the functions




Figure 7.1: Scaling with population in the structural network, by state (Panels A-D) and by function (Panels E-H). Panels A-D show scaling of state governments as a whole. Each point is a state. The x-axis plots state population; the y-axis plots: (A) nodes (state agencies, representing the scale of each states government), (B) pages (revealed scope, or density, of the activities state agencies report engaging in), (C) links (connectedness of a states agencies); and (D) traffic (visits to state government webpages, a proxy for the online engagement of citizens with the government). Panels E-H show the scaling law of each function by state population. Each point is a function. The x-axis plots the population scaling exponent, and the y-axis plots the amount of variation state population explains in each functions (E) nodes (agencies engaged in the function), (F) pages (revealed scope of activities of agencies engaged in the function), (G) links (connectedness of agencies engaged in the function with agencies in other functional clusters); and (H) traffic (visits to webpages of agencies engaged in the function).

that also show the greatest tendency toward structural similarity across states. For the other 26 general functions, state population explains only 4% of cross-state variance in agencies, 12% of their density, and 4% of their links (Figure 7.1, panels E-G). Many regulatory functions show little evidence of scaling – including utilities, audits, financial administration, professional boards and regulation, environment, and business regulation – as, more surprisingly, do health and human services, social services, and public safety – all service-provision functions that involve significant interaction with citizens and thus might be expected to be larger, more complex, or more connected when there are more citizens to serve.

We do find, however, that agency websites are more visited in larger states across a much wider range of functions. Figure 7.1D and 7.1H plot website usage against state population, overall (panel D) and by function (panel H), using web traffic data from Alexa. By far the most visited function is education, which accounts for 67% of all page-visits to government websites, followed by general administration (8.7%), legislative (4.8%), environment (4%), and business regulation (2%). Unlike the scale, connectedness, and density of the state governments that the websites reflect, we find that page visits to the websites themselves tend to scale close to linearly with population for a number of functions including education, municipal and county administration, and public safety, and superlinearly for several others, including business regulation and administrative law, indicating that citizens interaction with these functions is exponentially larger in more populous states.

In sum, state government scales with population. Citizens’ online interactions with agencies show linear or superlinear scaling, similar to a number of other urban phenomena such as personal income, education, and innovation [12, 13]. The bureaucratic networks implementing two functions, schooling and local administration, also scale, but sublinearly – similar to biological phenomena [10] and some characteristics of urban communities, including infrastructure, community size, and interactions among residents [12, 14]. But the extensive differences in the structures different states use to implement their remaining functions do not appear to be related to differences in the number of citizens they are responsible for serving.

7.2 Qualitative Variation: Function Implementation

In this section we discuss the measure we use in the paper to estimate functional variation across states, which equals 0 if the function is applied equally in all states, and increases for larger differences. We also compare this measure with some alternatives and discuss the results of its application to our dataset.

7.2.1 Diversity Measure

Our diversity measure has to match the intuition that functions implemented differently across states have a high value. We have to first quantify how a function can be implemented differently. We focus on three intuitive aspects.

1. A function can be either implemented by a single agency or multiple agencies;
2. A function can be simple or complicated;
3. The agencies performing the function can be connected to other agencies or can act in isolation.

We capture each aspect with a different measure:

1. Scale: The number of nodes (organizations) performing the function;
2. Density: The average number of pages in the websites of the organizations performing the function (in logs);
3. Connectedness: The amount of dispersion of agencies performing the function in different connected components.

These three types of organizational difference – scale, density, and connectedness – are distinct in the government’s Internet footprint: they are correlated at less than .1

A note on Connectedness: while all other functions are a simple count (of nodes and pages), Connectedness requires further understanding. First, we create an extract of the state network containing only the nodes performing the given function. From this extract we then identify the distribution in size of the connected components. For two agencies to be able to communicate, it is not necessary that they directly connect, as long as there is a path of links from an agency to another. Second, we calculate the normalized information entropy of the size distribution of the components. Meaning that we calculate:

$$C(x) = 1 - \left(\frac{-\sum_{i=1}^n P(x_i) \log P(x_i)}{\log n} \right),$$

where n is the number of nodes implementing the function (i.e. the sum of the component sizes), and $P(x_i)$ is the fraction of nodes that are part of the i th component. The advantage of using the information entropy over, say, the fraction of nodes in the largest connected component, comes from the fact that information entropy weights all components in the estimation of Connectedness. The scenario in which all nodes are split 50/50 in two component is different from the one in which 50% of the nodes are in the largest component and the rest is distributed in isolated 1-node components. Information entropy can discern between the two cases. Figure 7.2 depicts two examples of high and low Connectedness.




Figure 7.2: Two examples of high and low Connectedness. (Left) The network extract for Environment organizations in Indiana. Indiana has 12 organizations classified under Environment and 11 of them link to each other in a single connected component, making its Connectedness equal to $C(x) = .885$. (Right) The network extract for Business Regulation organizations in Ohio. Out of the nine organizations dealing with Business Regulation in Ohio, the largest connected component only hosts three, while most of the nodes are isolated in their own component making Ohio's Connectedness in this function equal to $C(x) = .237$.

To obtain a single measure of diversity we collapse these three dimensions into a single number for each state. To do so we first normalize Scale and Density to take values between 0 and 1 and then average the three measures. For each function we now have a vector of 50 numbers: one per state. The standard error of the mean of this vector is our measure of diversity. The higher the standard error, the more states tend to differ from the average implementation.

7.2.2 Alternatives

We tested several alternatives to our standard error measure of diversity:

- **State Distance:** We consider each state as a point in a three dimensional space. The value of each measure represents the position of the state in one of these dimensions. We then calculate the generalized euclidean distance among all possible pairs of states in this space. The higher the average distance, the more the states are dispersed, and therefore the more diverse their implementation of the function is.
- **Cluster Distance:** This measure is similar to State Distance, but calculated as the average distance between clusters of states rather than individual states. States sufficiently similar are grouped together and the centroid of the cluster represents all the states in the cluster.
- **Silhouette:** Silhouette is a quality measure used in data clustering. It is defined as the mean intra-cluster distance and the mean nearest-cluster




Figure 7.3: Variation in scale, density, and connectedness of six functions, arranged in decreasing diversity from top-left to bottom-right. Each data point is a state. Scale and density determine the x and y coordinates, while connectedness determines the size of the point.

distance for each sample [15]. The higher the Silhouette of a function, the more its clusters are interleaved with each other.

- Entropy: We use the Shannon information entropy formulation [16] applied to the clusters of the function. If all states belong to the same cluster, entropy is zero. If they are equally distributed among the clusters, entropy is at its maximum.

Each of these alternative measures has a downside relative to our standard error measure. State Distance is inversely proportional to the number of states performing the function, which might force lower diversity for functions implemented by many states. The other three measures are based on the clustering: the downside is that all intra-cluster variation is ignored.

Figure 7.3 shows six functions at different levels of diversity across states.

Chapter 8

State Difference Regression

In this chapter we explain in detail the regression results reported in Table 1 of the main paper. We first explain the interpretation of the dependent variables, then describe the independent variables. Finally, we report additional information and robustness checks.

8.1 Dependent Variables

Table 1 in the main paper reports the coefficients and their significance of various state-state distance measures for predicting the similarity between two state governments. We generate several measures of similarity in state governments, all using the same procedure.

Let us start by considering the simplest similarity measure: Scale. This measure is based on the number of agencies (nodes) in a given function. A state i is represented by a vector s_i , which has one entry per function. The value of each entry is the base 10 logarithm of its number of agencies. Then, the similarity σ_{ij} of state i with state j is given by:

$$\sigma_{ij} = 1 - \rho(s_i, s_j),$$

where ρ is the function calculating the correlation distance between two vectors.

The σ similarities for Density and Connectedness are calculated using the same procedure, with the exception that Connectedness is the raw number, unnormalized. For the overall diversity, the three vectors per state (Scale, Density, and Connectedness) are stacked together in the same vector.

Figure 8.1 depicts the distributions of each of the four measures (scale, density, connectedness, and overall). Each measure distributes normally or quasi-normally across all observed state pairs.




Figure 8.1: The distribution of the pairwise state similarities according to our four measures: the bar reports how many state pairs have a similarity value included in the interval covered by the bar itself.

8.2 Independent Variables

Table 1 in the main paper proposes four different possible hypotheses and one control for the similarity of state governments. In each of the following subsections we describe the measure associated with each hypothesis and control, including the source of the data and the preprocessing performed to construct the variable. Once calculated, every independent variable is then normalized by subtracting its average and dividing by its standard deviation. In this way, the normalized coefficients from the regression are easily comparable.

8.2.1 Ideology

The first hypothesis is that states whose citizens have similar ideologies on the typical left-right dimension will have similar governments. Political scientists have developed numerous measures to place voters on this dimension; we use the NOMINATE ideology score [17]. To account for the likely lag in ideological opinion to be reflected in government structure, we use average ideology in each state from 2005 to 2010. In NOMINATE, a score of zero means that the state is dominated by Republican ideology, while 100 indicates Democrat dominance. We transform this into a state-state distance measure by taking the absolute value of the difference of the scores of the two states.

8.2.2 Economic Structure

The second hypothesis is that states with similar economic structures should have more similar governments. To measure economic structure, we examine the structure of employment across industries, using state-industry employment data from the County Business Patterns¹.

Each state is represented by a vector; each entry in the vector corresponds to an industry. The value of the entry is the Revealed Comparative Advantage (RCA) of the state in the industry. To calculate the RCA, we start by considering p_{si} as the total payroll in industry i and state s . RCA is then calculated as:

$$RCA(p_{si}) = \frac{p_{si}}{\sum_s p_{si}} \bigg/ \frac{\sum_i p_{si}}{\sum_{si} p_{si}},$$

which is the relative importance of i in s on the relative importance of i in the aggregation of all states. We transform this into a state-state distance measure by calculating the Euclidean distance between two state vectors.

8.2.3 Income

The third hypothesis is that differences in government structure may reflect differences in state income. For each state, we collect from the US Census information about its GDP per capita. Similarly to ideology, we transform this into a pairwise distance measure by taking the absolute value of the difference of the two states.

8.2.4 Geographical Distance

The fourth hypothesis is that proximity between states leads to similarity in their governments, by facilitating interaction among their agencies, providing opportunities for learning and common understanding of how legitimate and effective governmental organizations are structured.

To measure geographic distance, we represent the United States as a graph, in which each node represents a state. States are connected by an edge if they share a border. The geographical distance between two states is the number of edges (border crossings) necessary to cross from one state to another.

Non-continental states like Alaska and Hawaii do not share borders with any US state. For this reason, we consider as a node in the network also the Pacific Ocean and Canada. As a consequence, while for continental states the lowest geographical distance is one (from the states with which they share a border), for Alaska and Hawaii the minimum distance is two: they first have to hop to the Pacific ocean and then to the destination state.

¹<https://www.census.gov/programs-surveys/cbp.html>




Figure 8.2: The distribution of the pairwise normalized state distances: the bar reports how many state pairs have a normalized distance value included in the interval covered by the bar itself.

8.2.5 Network Position

A similar measure is used for our control measure: two states whose agencies are more closely connected online might tend to converge around similar ways of organizing, and thus have more similar governments. Here, we consider the full network of government agencies. Two agencies are connected if they link directly to each other on the Web. We calculate all pairwise shortest paths between pairs of agencies belonging to two states in the network. The average length of these shortest paths is used as our network distance.

Figure 8.2 depicts the distributions of each of the five distance measures (ideology, income, economic, geography and network), after the normalization. Each measure distributes normally or log-normally across all observed state pairs.

8.3 Robustness Checks

8.3.1 Statistical Tests

	indu dist	geo distance	ideo distance	income diff	fullnet dist	traffic diff
indu dist	1.000	0.295	-0.022	0.236	0.182	0.296
geo distance	0.295	1.000	0.091	0.191	0.261	0.111
ideo distance	-0.022	0.091	1.000	-0.049	-0.015	0.011
income diff	0.236	0.191	-0.049	1.000	0.156	-0.028
fullnet dist	0.182	0.261	-0.015	0.156	1.000	-0.012
traffic diff	0.296	0.111	0.011	-0.028	-0.012	1.000

Table 8.1: Pairwise Pearson correlations between the vectors of our five independent variables.

	indu dist	geo distance	ideo distance	income diff	fullnet dist	traffic diff
indu dist	1.000	0.327	-0.010	0.201	0.155	0.252
geo distance	0.327	1.000	0.095	0.202	0.258	0.044
ideo distance	-0.010	0.095	1.000	-0.054	-0.004	0.088
income diff	0.201	0.202	-0.054	1.000	0.150	-0.022
fullnet dist	0.155	0.258	-0.004	0.150	1.000	-0.128
traffic diff	0.252	0.044	0.088	-0.022	-0.128	1.000

Table 8.2: Pairwise Spearman rank correlations between the vectors of our five independent variables.

In the main paper, we note that the independent variables in the Table 2 regressions are not strongly correlated with each other. Tables 8.1 and 8.2 report their pairwise Pearson and Spearman correlation coefficients, respectively. We can see that all correlation coefficients are lower than .3, with the sole exception of the Spearman correlation between industrial composition and geographical distance.

In Table 2 in the main paper, we only test the relationship between government functional similarity and all the distance measures taken together. Here,

Distance	Scale	Density	Connectedness	Overall
Ideology	0.000575 (0.00137)	-0.0039 (0.00214)	-0.0043 (0.00278)	-0.0028 (0.0018)
Income	-0.000585 (0.00149)	-0.00588* (0.00232)	-0.00691* (0.00302)	-0.00527** (0.00195)
Economic	-0.0244** (0.00169)	-0.0186** (0.00283)	-0.0177** (0.00371)	-0.0168** (0.00237)
Location	-0.0053** (0.00149)	-0.0119** (0.00232)	-0.0139** (0.00303)	-0.00969** (0.00195)
Network	-0.00928** (0.00232)	-0.0144** (0.00364)	-0.0184** (0.00474)	-0.0112** (0.00306)
State F.E.	Y	Y	Y	Y
Obs.	1225	1225	1225	1225

Table 8.3: Results of bivariate version of the models reported in Table 1 in the main paper. Each distance measure (Ideology, Income, Economic, Geographic, and Network) is tested independently against each measure of government functional similarity (Scale, Density, Connectedness, and Overall).

in Table 8.3, we test the relationship between each distance measure with each government similarity measure independently. We can see that many distance measures have a negative correlation with government similarity, meaning that states with similar levels of income, economic structure, geographical location and network connections tend to have similar governments. The only exception to the rule is ideology.

	Scale	Density	Connectedness	Overall
Ideology	0.7969	0.4721	0.3702	0.4632
Income	0.7969	0.4735	0.3717	0.4659
Economic	0.811	0.4913	0.3814	0.4861
Geographic	0.7977	0.4826	0.3803	0.4739
Network	0.7989	0.4778	0.3772	0.4687

Table 8.4: Bivariate R^2 for the regressions in section A of Table 1 in the main paper.

For space constraints, we split out from Table 8.3 the R^2 of those regressions, which we report in Table 8.4. We can see that the Scale similarity is the variable that is the most easily predicted, while Connectedness is harder to explain with our covariates. Economic structure is the best bivariate predictor of government similarity across all measures, confirming the results of the multivariate regression in Table 1 of the main paper

In the main paper we show a version of the regressions in which Scale equals to the logarithm of the number of agencies performing a function in a state. Here we check that the results are robust to measuring scale relative to population, by replacing the Scale measure in the paper with the number of agencies in the function per million inhabitants in the state. Table 8.5 show the regression results with this population-adjusted Scale measure, both alone and as a part of the overall diversity measure. The results are consistent with the main table in the paper.

The results reported in the table include fixed effects for state, to account for unobserved state characteristics (such as population, history, culture, and any political, economic, or geographic characteristics not reflected in our other measures). Table 8.6 shows regression results without these fixed effects, which

Distance	A: Bivariate				B: Multivariate			
	Scale	Density	Connectedness	Overall	Scale	Density	Connectedness	Overall
Ideology	0.00115 (0.00199)	-0.0039 (0.00214)	-0.0043 (0.00278)	-0.00267 (0.00184)	0.00168 (0.00183)	-0.00236 (0.00211)	-0.00242 (0.00277)	-0.0014 (0.0018)
Income	4.63e-05 (0.00216)	-0.00588* (0.00232)	-0.00691* (0.00302)	-0.00557** (0.002)	0.00422* (0.00202)	-0.0026 (0.00234)	-0.00364 (0.00307)	-0.00233 (0.002)
Economic	-0.03** (0.00252)	-0.0186** (0.00283)	-0.0177** (0.00371)	-0.0192** (0.00242)	-0.0226** (0.00281)	-0.0123** (0.00324)	-0.0108* (0.00425)	-0.0135** (0.00277)
Location	-0.00459* (0.00217)	-0.0119** (0.00232)	-0.0139** (0.00303)	-0.01** (0.002)	0.000271 (0.00218)	-0.00726** (0.00252)	-0.00862** (0.0033)	-0.00574** (0.00215)
Network	-0.0114** (0.00338)	-0.0144** (0.00364)	-0.0184** (0.00474)	-0.0117** (0.00314)	-0.00442 (0.00329)	-0.0067 (0.0038)	-0.0105* (0.00499)	-0.00446 (0.00325)
State F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Obs.	1225	1225	1225	1225	1225	1225	1225	1225
R^2					0.8456	0.5394	0.44	0.5393
$Adj R^2$					0.8315	0.4975	0.3891	0.4974

Table 8.5: Robustness check for Table 2 in the the main paper: changing the Scale component of the state similarity.

Distance	A: Bivariate				B: Multivariate			
	Scale	Density	Connectedness	Overall	Scale	Density	Connectedness	Overall
Ideology	0.00813** (0.003)	-0.00181 (0.00243)	0.00327 (0.0029)	0.00107 (0.00201)	0.00777** (0.00245)	-0.000624 (0.00242)	0.00265 (0.00285)	0.00174 (0.00195)
Income	-0.00494 (0.00301)	-0.00482* (0.00243)	-0.0108** (0.00288)	-0.00677** (0.00201)	0.0128** (0.00254)	-0.00254 (0.0025)	-0.00653* (0.00295)	-0.00163 (0.00201)
Economic	-0.0587** (0.0025)	-0.00249 (0.00243)	-0.0123** (0.00288)	-0.0131** (0.00198)	-0.053** (0.00276)	0.00341 (0.00272)	-0.00325 (0.0032)	-0.00514* (0.00219)
Location	-0.029** (0.00289)	-0.0141** (0.0024)	-0.00833** (0.00289)	-0.0155** (0.00197)	-0.0142** (0.00267)	-0.0133** (0.00263)	-0.00151 (0.00309)	-0.0125** (0.00212)
Network	-0.0295** (0.00289)	-0.00686** (0.00242)	-0.0218** (0.00283)	-0.0133** (0.00198)	-0.00907** (0.00264)	-0.00422 (0.0026)	-0.0192** (0.00307)	-0.00783** (0.0021)
State F.E.	N	N	N	N	N	N	N	N
Obs.	1225	1225	1225	1225	1225	1225	1225	1225
R^2					0.3498	0.03184	0.05386	0.08472
$Adj R^2$					0.3466	0.02707	0.0492	0.08022

Table 8.6: Robustness check for Table 2 in the the main paper: removing state fixed effects controls.

Distance	A: Bivariate				B: Multivariate			
	Scale	Density	Connectedness	Overall	Scale	Density	Connectedness	Overall
Ideology	0.00813** (0.003)	-0.00181 (0.00243)	0.00327 (0.0029)	0.00107 (0.00201)	0.00776** (0.00246)	-0.000597 (0.00242)	0.00275 (0.00285)	0.00177 (0.00195)
Income	-0.00494 (0.00301)	-0.00482* (0.00243)	-0.0108** (0.00288)	-0.00677** (0.00201)	0.0127** (0.00254)	-0.00253 (0.0025)	-0.00649* (0.00294)	-0.00162 (0.00202)
Economic	-0.0587** (0.0025)	-0.00249 (0.00243)	-0.0123** (0.00288)	-0.0131** (0.00198)	-0.0531** (0.00278)	0.00354 (0.00274)	-0.00279 (0.00322)	-0.00499* (0.0022)
Location	-0.029** (0.00289)	-0.0141** (0.0024)	-0.00833** (0.00289)	-0.0155** (0.00197)	-0.0142** (0.00267)	-0.0134** (0.00263)	-0.00168 (0.0031)	-0.0126** (0.00212)
Network	-0.0295** (0.00289)	-0.00686** (0.00242)	-0.0218** (0.00283)	-0.0133** (0.00198)	-0.00906** (0.00264)	-0.00424 (0.0026)	-0.0192** (0.00307)	-0.00785** (0.0021)
Ideo-Income Int.	0.00455 (0.00305)	0.000711 (0.00246)	0.00461 (0.00293)	0.00187 (0.00204)	-0.000459 (0.00248)	0.00111 (0.00244)	0.00398 (0.00288)	0.00122 (0.00197)
State F.E.	N	N	N	N	N	N	N	N
Obs.	1225	1225	1225	1225	1225	1225	1225	1225
R^2					0.3498	0.032	0.05535	0.08501
$Adj R^2$					0.3461	0.02643	0.04991	0.07975

Table 8.7: Robustness check for Table 2 in the the main paper: removing state fixed effects controls and adding ideology-income interaction term.

allows us to see the raw predictive power of each variable.

Although ideology and income are not strongly predictive of state government similarity, they may be interactively predictive: for example, wealthy states may have more freedom to fashion government according to their citizens' ideological preferences. Table 8.7 shows the regression results with such an interaction, without fixed effects for states. This interaction is not significant in any of the regressions.

Finally, we substitute state fixed effects with random effects. Table 8.8 shows the regression results. The results are substantively similar; notably, geography and economic structure are still the strongest predictors of state-government similarity.

8.3.2 Alternative Ideology Measures

One of the most important results in the main paper is the absence of an association between ideology and structural similarity in state governments. Our main measure of ideology is NOMINATE [18] scores, which are based on roll call votes of each state's congressional delegation. Specifically we averaged 2005 to 2010 NOMINATE scores per state and then defined the ideology distance of two states as the absolute value of the difference between the two states.

Distance	A: Bivariate				B: Multivariate			
	Scale	Density	Connectedness	Overall	Scale	Density	Connectedness	Overall
Ideology	0.000794 (0.00136)	-0.00354* (0.00209)	-0.00322 (0.00271)	-0.00224 (0.00175)	0.00138 (0.00118)	-0.00223 (0.00206)	-0.00174 (0.00269)	-0.00108 (0.00172)
Income	-0.000729 (0.00148)	-0.00584** (0.00225)	-0.0077** (0.00291)	-0.00549** (0.00189)	0.00295 (0.00131)	-0.00271 (0.00227)	-0.00429 (0.00295)	-0.00222 (0.00189)
Economic	-0.0255** (0.00167)	-0.0163** (0.00268)	-0.0171** (0.00344)	-0.0167** (0.00221)	-0.0186** (0.00179)	-0.0107** (0.00292)	-0.011** (0.00373)	-0.0126** (0.0024)
Location	-0.00589** (0.00148)	-0.0118** (0.00224)	-0.0126** (0.00291)	-0.0105** (0.00188)	-0.00164 (0.00141)	-0.00756** (0.00241)	-0.00635* (0.00312)	-0.00583** (0.002)
Network	-0.0105** (0.00227)	-0.0123** (0.00323)	-0.0196** (0.00402)	-0.0124** (0.00264)	-0.00346 (0.00211)	-0.00544 (0.00339)	-0.0132** (0.00428)	-0.00609* (0.00275)
State Rand.E.	Y	Y	Y	Y	Y	Y	Y	Y
Obs.	1225	1225	1225	1225	1225	1225	1225	1225

Table 8.8: Robustness check for Table 2 in the the main paper: substituting state fixed effects with state random effects controls.

Distance	A: Bivariate				B: Multivariate			
	Scale	Density	Connectedness	Overall	Scale	Density	Connectedness	Overall
Ideology	0.00077 (0.00129)	-0.00363 (0.00203)	-0.00375 (0.00265)	-0.0022 (0.0017)	0.00233* (0.00115)	-0.000598 (0.00207)	-0.000288 (0.00272)	0.000514 (0.00173)
Income	-0.00058 (0.00148)	-0.00585* (0.00232)	-0.0069* (0.00303)	-0.00518** (0.00195)	0.00221 (0.00133)	-0.00253 (0.00238)	-0.00367 (0.00314)	-0.00247 (0.002)
Economic	-0.0245** (0.00168)	-0.0186** (0.00282)	-0.0176** (0.00371)	-0.0167** (0.00236)	-0.0171** (0.00182)	-0.0121** (0.00326)	-0.0106* (0.00429)	-0.0117** (0.00273)
Location	-0.00515** (0.00149)	-0.0117** (0.00232)	-0.0138** (0.00303)	-0.00948** (0.00195)	-0.00172 (0.0014)	-0.00723** (0.0025)	-0.00887** (0.00329)	-0.00576** (0.0021)
Network	-0.00929** (0.00232)	-0.0146** (0.00363)	-0.0184** (0.00474)	-0.0113** (0.00305)	-0.00287 (0.00212)	-0.00721 (0.00379)	-0.0107* (0.005)	-0.00499 (0.00318)
State F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Obs.	1225	1225	1225	1225	1225	1225	1225	1225
R^2					0.9065	0.5503	0.4512	0.5328
$Adj R^2$					0.898	0.5095	0.4013	0.4903

Table 8.9: Robustness check for Table 2 in the the main paper: substituting NOMINATE ideology with ideology score from [19].

NOMINATE scores are widely used for their availability and comparability but may be imperfect proxies of state ideology: either of the ideological views of voters or of their state representatives, whose decisions are more directly relevant to the functional structures of their state governments [19]. We therefore verify the absence of a robust association between structural similarity and two alternative measures: (i) NOMINATE scores calculated for each state legislature [20, 21], averaged over the same 2005-2010 period as the results we report in the main paper, which are strongly correlated with self-reported ideology questions from the 2000-2008 Annenberg National Election Surveys; and (ii) an alternative measure of the ideological preferences of voters in each state as measured in state surveys conducted by the Pew Research Center in 2014, the year of our crawl².

Tables 8.9 and 8.10 show the results, which confirm the lack of a robust association between ideology and state functional structures that we report in the main paper. In bivariate regressions NOMINATE scores based on votes of state legislators are unassociated with any of our measures of state functional structures, and are only weakly associated with scale in the full model controlling for income, economic structure, and location, and in the wrong direction: implying, implausibly, that when states are more similar ideologically, the functional structures of their governments are more likely to be different. The Pew surveys

²<http://www.pewforum.org/religious-landscape-study/compare/political-ideology/by/state/>

Distance	A: Bivariate				B: Multivariate			
	Scale	Density	Connectedness	Overall	Scale	Density	Connectedness	Overall
Ideology	-0.00134 (0.00148)	-0.00848** (0.00232)	-0.00551 (0.00303)	-0.00567** (0.00195)	0.00243 (0.0014)	-0.00344 (0.0025)	0.000935 (0.0033)	-0.000924 (0.0021)
Income	-0.00058 (0.00148)	-0.00585* (0.00232)	-0.0069* (0.00303)	-0.00518** (0.00195)	0.00204 (0.00137)	-0.00166 (0.00244)	-0.00402 (0.00322)	-0.00208 (0.00205)
Economic	-0.0245** (0.00168)	-0.0186** (0.00282)	-0.0176** (0.00371)	-0.0167** (0.00236)	-0.0172** (0.00184)	-0.0114** (0.00328)	-0.0109* (0.00433)	-0.0114** (0.00275)
Location	-0.00515** (0.00149)	-0.0117** (0.00232)	-0.0138** (0.00303)	-0.00948** (0.00195)	-0.00181 (0.0014)	-0.00689** (0.00251)	-0.00899** (0.00331)	-0.00562** (0.00211)
Network	-0.00929** (0.00232)	-0.0146** (0.00363)	-0.0184** (0.00474)	-0.0113** (0.00305)	-0.00291 (0.00212)	-0.00701 (0.00379)	-0.0108* (0.005)	-0.0049 (0.00318)
State F.E.	Y	Y	Y	Y	Y	Y	Y	Y
Obs.	1225	1225	1225	1225	1225	1225	1225	1225
R^2					0.9064	0.5511	0.4512	0.5328
$Adj R^2$					0.8979	0.5103	0.4013	0.4903

Table 8.10: Robustness check for Table 2 in the the main paper: substituting NOMINATE ideology with Pew ideology.

from the same year as our crawl are strongly correlated with similarity in the density of state functional structures (Table 8.10), but not once we add income, economic structure, and location to the model. In these specifications, economic structure and location remain strongly correlated with similarity in density of state functional structures, suggesting that these other factors are more strongly associated with similarity in density than ideology. (Indeed, more recent work by Tausanovitch and Warshaw [22] explicitly includes economic and geographic variables in multi-level modeling of public opinion in geographic sub-units.)

Bibliography

- [1] Susan T Dumais. Latent semantic analysis. *Annual review of information science and technology*, 38(1):188–230, 2004.
- [2] Winter Mason and Siddharth Suri. Conducting behavioral research on amazons mechanical turk. *Behavior research methods*, 44(1):1–23, 2012.
- [3] Joseph L Fleiss. Measuring nominal scale agreement among many raters. *Psychological bulletin*, 76(5):378, 1971.
- [4] Kilem L Gwet. *Handbook of inter-rater reliability: The definitive guide to measuring the extent of agreement among raters*. Advanced Analytics, LLC, 2014.
- [5] M Ángeles Serrano, Marián Boguñá, and Alessandro Vespignani. Extracting the multiscale backbone of complex weighted networks. *Proceedings of the national academy of sciences*, 106(16):6483–6488, 2009.
- [6] Enys Mones, Lilla Vicsek, and Tamás Vicsek. Hierarchy measure for complex networks. *PloS one*, 7(3):e33799, 2012.
- [7] Enys Mones. Hierarchy in directed random networks. *Physical Review E*, 87(2):022817, 2013.
- [8] Jianxi Luo and Christopher L Magee. Detecting evolving patterns of self-organizing networks by flow hierarchy measurement. *Complexity*, 16(6):53–61, 2011.
- [9] Michele Coscia. Using arborescences to estimate hierarchicalness in directed complex networks. *PloS one*, 13(1):e0190825, 2018.
- [10] Geoffrey B West, James H Brown, and Brian J Enquist. A general model for the origin of allometric scaling laws in biology. *Science*, 276(5309):122–126, 1997.
- [11] Alberto Alesina and Enrico Spolaore. *The size of nations*. Mit Press, 2005.
- [12] Luís MA Bettencourt. The origins of scaling in cities. *science*, 340(6139):1438–1441, 2013.

- [13] Andres Gomez-Lievano, Oscar Patterson-Lomba, and Ricardo Hausmann. Explaining the prevalence, scaling and variance of urban phenomena. *Nature Human Behaviour*, 1(1):0012, 2017.
- [14] Luís MA Bettencourt, José Lobo, Dirk Helbing, Christian Kühnert, and Geoffrey B West. Growth, innovation, scaling, and the pace of life in cities. *Proceedings of the national academy of sciences*, 104(17):7301–7306, 2007.
- [15] Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.
- [16] Claude Elwood Shannon. A mathematical theory of communication. *Bell System Technical Journal*, 27(3):379–423, 1948.
- [17] Keith T Poole and Howard Rosenthal. A spatial model for legislative roll call analysis. *American Journal of Political Science*, pages 357–384, 1985.
- [18] Keith T Poole and Howard Rosenthal. *Congress: A political-economic history of roll call voting*. Oxford University Press on Demand, 2000.
- [19] Boris Shor and Nolan McCarty. The ideological mapping of american legislatures. *American Political Science Review*, 105(3):530–551, 2011.
- [20] William D Berry, Evan J Ringquist, Richard C Fording, and Russell L Hanson. Measuring citizen and government ideology in the american states, 1960-93. *American Journal of Political Science*, pages 327–348, 1998.
- [21] William D Berry, Richard C Fording, Evan J Ringquist, Russell L Hanson, and Carl E Klarner. Measuring citizen and government ideology in the us states: A re-appraisal. *State Politics & Policy Quarterly*, 10(2):117–135, 2010.
- [22] Chris Tausanovitch and Christopher Warshaw. Measuring constituent policy preferences in congress, state legislatures, and cities. *The Journal of Politics*, 75(2):330–342, 2013.

Appendix A

Government Function Classifications

A.1 US Census

Code	Label
01	Air Transportation (Airports)
02	Cemeteries
09	Education (School Building Authorities)
24	Fire Protection
32	Health
40	Hospitals
41	Industrial Development
42	Mortgage Credit
44	Highways (Regular or Toll)
50	Housing and Community Development
51	Drainage
52	Libraries
59	Other Natural Resources
60	Parking Facilities
61	Parks and Recreation
63	Flood Control
64	Irrigation
79	Public Welfare Institutions
80	Sewerage
81	Solid Waste Management
86	Reclamation
87	Water Transport and Terminals
88	Soil and Water Conservation
89	Other Single-Function Districts
91	Water Supply Utility

92	Electric Power Utility
93	Gas Supply Utility
94	Mass Transit System Utility
96	Fire Protection and Water Supply
97	Natural Resources and Water Supply
98	Sewerage and Water Supply
99	Other Multi-Function Districts

A.2 Web Classification

Administrative Law General Appeals and Hearings; Civil Service Appeals and Hearings; Tax Appeals and Hearings; Workers' Compensation Appeals and Hearings.
Audits, Inspectors General & Accountability
Boards & Professional Regulation General Licensing; Architects, Landscape Architects & Interior Designers; Athletic Trainers, Occupational Therapy & Physical Therapy; Audiology, Speech & Language Pathology; Barbers & Cosmetology; Chiropractic; Construction & Contractors; Counselors, Therapists, Psychologists & Social Workers; Dental; Engineers & Land Surveyors; Geologists and Soil Scientists; Massage Therapy; Medical Examiners; Medicine; Nursing; Nursing Home & Assisted Living Administrators; Optometry; Pharmacy; Alcohol, Drug, & Gambling Counselors; Private Investigators; Real Estate Appraisers; Veterinary.
Business Regulation General; Alcohol & Tobacco Regulation; Athletic & Boxing Commissions; Cemeteries; Land Use & Building Codes; Lotteries; Gaming & Racing Commissions; Real Estate.
Commerce and Economic Development General; Consumer Protection and Education; Film & Television Promotion; Housing and Community Development; Improvement Districts; Industrial Development; Minority & Women's Businesses; Small Businesses.
Criminal Justice General; Juvenile; Pardons, Parole & Probation; Sentencing; Sex Offenders Registry; Victims Advocate.
Data & Information General; Geographical Information Systems; Health & Vital Statistics; Information Technology & Communications Technology.
Education General; Adult Education; Higher: General; Higher: Universities and Colleges; Libraries; Primary and Secondary: General; Primary and Secondary: Career & Technical Centers; Primary and Secondary: Early Childhood Centers; Primary and Secondary: Schools and Districts.

Environment Agriculture, General; Agriculture, Product Promotion and Marketing Boards; General; Conservation & Protection, General; Coastal Protection; Fish, Game & Wildlife; Flood Control & Levees; Forestry; Land, Soil & Water; Parks and Recreation; Rivers; Energy; Minerals and Mining; Oil and Gas; Power Agencies & Authorities.
Ethics, Lobbying & Campaign Finance Commissions
Financial Administration General & General Regulation; Banking; Budget Offices; Credit Unions; Fiscal Security; Insurance & Risk Management; Purchasing; Securities.
Financing Authorities General; Educational Facilities; Health Facilities; Housing and Community Development; Student Loans & Tuition.
General Administration Arts & Humanities; Arts & Humanities: Public Broadcasting; General; Public Buildings & Facilities; Secretary of State; Volunteering & Community Service; Elections; Municipal and County Administration; Governors & Lieutenant Governors
Health & Human Services General; Public Health, General; Children and Families: General; Children and Families: Child Services, General; Children and Families: Family Services, General; Children and Families: Child Abuse Prevention; Children and Families: Child Advocate; Children and Families: Child Support; Deaf, Hard of Hearing & Blind; Developmental Disabilities; Mental & Behavioral Health, Substance Abuse; Rehabilitation, Independent Living & Disability Services; Elder Care, Aging & Long Term Care; Hospitals & Medical Centers.
Historic Preservation General; Archives & Records; Museums.
Homeland Security
International Trade
Judicial and Legal General; Attorney General, Solicitor & District Attorneys; Civil & Human Rights; Courts, General; Public & Indigent Defense; Supreme Court.
Labor and Human Resources General; Labor & Industrial Relations Boards; Occupational Safety & Industrial Accidents; Personnel, Human Resources & Jobs; Workforce Development & Improvement.
Legislative
Military State Guards; National.
Public Safety General; Corrections & Correctional Facilities; Emergency Management & 911; Fire Protection; Office of the Medical Examiner; Police, Sheriffs & Patrol Offices.

Research & Science General; Geological Surveys; National & Federal Laboratories.
Social Services General; Health Care & Insurance Exchanges; Indian & Native American Affairs; Medicaid; Minority Affairs; Retirement; Social Security, Social Insurance Administration, Medicare; Unemployment Insurance & Employment Security; Veterans Affairs; Women's Affairs.
Tourism & Travel General; Conventions & Fairs.
Transportation General; Air Transportation, Aviation & Airports; Highways & Roads; Mass Transit; Motor Vehicles; Port Authorities.
Treasury & Revenue General; Investments.
Utilities Oversight and Adjudication; General; Sanitation; Waste Management.