The Cultural Evolution of National Constitutions

Daniel N. Rockmore[1,2,3,4], Chen Fang[1], Nicholas J. Foti[5], Tom Ginsburg[6], and David C. Krakauer[3,7,8]

[1]Department of Computer Science, Dartmouth College, Hanover, NH, USA 03755

[2]Department of Mathematics, Dartmouth College, Hanover, NH, USA 03755

[3]The Santa Fe Institute, Santa Fe, NM, USA 87501

[4]The Neukom Institute for Computational Science, Dartmouth College, Hanover, NH, USA 03755

[5]Department of Statistics, University of Washington, Seattle, WA USA 98195-4322

[6]University of Chicago Law School, The University of Chicago, Chicago, IL, USA 60637

[7]Wisconsin Institute for Discovery & Department of Genetics, University of Wisconsin, Madison, WI, USA 53715

[8]Department of Genetics, University of Wisconsin, Madison, WI, USA 53715

February 29, 2016

Abstract

We introduce a hybrid of approaches, inspired by biology and genetics, to analyze patterns of cultural inheritance and innovation through the study of the diffusion of ideas through a corpus of 591 national constitutions spanning 1789–2008. We extract information from a topic modeling of the complete corpus and construct cultural diffusion trees of topics (in the topic modeling sense) in order to characterize constitutions as cultural recombinants borrowing from ancestral constitutions back to the Last Universal Common Ancestor of Constitutions (LUCAC), the US Constitution of 1789. We discover constitutions cluster into distinct epochs within which legal concepts are frequently shared. We find constitutions vary systematically in their patterns of borrowing from ancestral texts – from asexual copying through to polysexual borrowing. Most constitutions are very similar and have only a short term influence on descendant constitutions but a few are surprisingly innovative with very many offspring with a long lasting influence. These highly influential constitutions tend to be the oldest. We find that constitutions behave “biologically” in that their patterns of inheritance follow a characteristic negative-binomial distribution of “offspring” arising through a preferential-attachment process. These findings support a principled definition of memes in which the particulate inheritance of topics reproduces regularities in both constitutional statistics and dynamics.

Background

Cultural inheritance involves the diffusion of innovations, a process of interest to both biologists [1] and social scientists [2]. In biology inheritance is governed by mechanisms of genetic transmission, which have been quantified [3].

Cultural inheritance takes a variety of forms which can resemble variants of biological inheritance [4, 5], including cultural selection [6, 7]. In cultural domains, complex forms of knowledge are encoded in social norms, legal principles and scientific theories [8, 9] and follow complex forms of transmission that involve the coordinated borrowing and learning of constellations of ideas.

Now that a large body of the cultural record has been digitized (including books [10], music [11], art [12], etc.) new techniques of machine learning are making the quantitative analysis of high-dimensional cultural artifacts possible. In analogy with the biological sciences, and genetics in particular, this data mining approach to the analysis of culture is sometimes referred to as "culturomics" [13], a term born of the consideration of the frequency distribution of an *n*-gram over time [14] as proxy for how memes move in and out of the cultural record. Literature (and text generally) remains a primary focus of such work (see e.g., [15, 16, 17]). A fascinating challenge is to supplement these correlation-based approaches to culture with principled causal mechanisms directed at discovering fundamental, extra-biological evolutionary processes.

We consider the notion of *diffusion patterns* in the study of cultural inheritance as a means of tracking the diffusion of *topics* through the documents in a legal text corpus of five hundred and ninety-one national constitutions (the full list is given at the end of the paper in Table 2a–e). "Topics" has a technical meaning here (and throughout this paper that is the sense in which the word is used) as probability distributions over words (positive weights that sum to one) that are the output of *topic modeling* is a computational and statistical methodology for text analysis that has made great inroads throughout the humanities (see e.g., [18, 19]), to the point of reaching an almost "plug-and-play" form [20, 21] for easy deployment. A set of topic is "learned" (i.e., automatically derived) from the corpus. The various topic distributions highlight (i.e., attach high weight to) different sets of words. In the best cases those words usually suggest a particular theme and associated labeling of the topic. Pieces of a given text in the corpus are thus represented as varying mixtures of topics. In this way topics provide a low-dimensional representation of the corpus in terms of higher level ideas and could provide a rigorous operational basis for a meme, to be tested against a suitable dynamics of inheritance. Although we focus on its use in the analysis of text, the topic modeling framework is more general and has been used in a number of areas [22].

Given a topic of some significance in a work, embodied in a set of semantically correlated legal concepts, we track its appearance and prevalence in subsequent works within the corpus, as well as its extinction. While dynamical considerations have been incorporated previously into topic models [23, 24] this analysis differs in that we account for the diffusion of topics from document to document, and in this way reveal more clearly the patterns of genealogy and the essentially recombinant nature of textual artifacts. It is our contention that while culture is clearly an active *in situ* feature of human brains [25], it is also present in material artifacts which afford rich forms of combinatorial manipulation and transmission *ex situ*.

The corpus of national constitutions is particularly well-suited to a framing and analysis as a document corpus composed of units of correlated meaning evolving according to idea diffusion and borrowing. Indeed, scholars have demonstrated that many provisions in constitutions are copied from those of other countries. For example, Law and Versteeg [26] have shown that rights provisions have spread around the globe. Through *n*-gram analysis Ginsburg et al. show that constitutional preambles, which are conceptualized as the most nationally localized part of constitutions, also speak in a universal idiom and include a good deal of borrowing [27]. Elkins et al. [28, 29] show that some rights, such as freedom of expression, have become nearly universal, while others have not. Some even argue that there is a kind of global script at work, whereby nation-states seek to use constitutions to participate in global discourses [30, 31, 32]. This evolutionary framing of the creation of national constitutions draws on broader biological analogies for legal development across time and space [33]. Our use of diffusion trees for this problem can be seen as a new quantification of this biological analogy.

Results

A *topic* is a probability distribution over a fixed *vocabulary* derived from a text *corpus*. As such it represents a correlated set of words[1] encoding something like a "meme" or stochastic set of associations. The text corpus is partitioned into *documents*, sets of roughly contiguous groupings of 500 words[2] from the vocabulary (e.g., in the best case each constitution would be partitioned into contiguous word-blocks, but processing may remove the odd abbreviation, title, etc. besides respecting natural boundaries, such as the end of one constitution and the beginning of another). In the case of our corpus of constitutions, each constitution generally comprises a subset of such documents. The model does not care about word order, just which words occur and in what frequencies. This is the so-called "bag-of-words" model or representation, which is thus encoded as a probability distribution over the vocabulary. *Topic modeling* is a methodology for learning topics such that each document (as a bag of words) is represented as a weighted sum (mixture) of topics. In its generative form, each document is created by first choosing a topic according to the mixture it comprises and then choosing a word according to the distribution of that particular topic. In this respect a constitution can be thought of as a "meme cloud". We use the latent Dirichlet allocation (LDA) topic model (see [34] for a discussion of the various parameters that define the model). We tested several choices for the number of topics and chose 100 which we then validated.

The topic model forms the basis for our results. They include (1) the discovery of the topics that make up the corpus of constitutions, (2) the determination of their flow through time ("information cascades"), the (3) reconstruction of cultural diffusion trees; (4) network analysis of diffusion trees; and (5) discovery of a very biological pattern of inheritance with a highly skewed pattern of cultural fertility.

Topics.

We successfully labeled (i.e., recognized a clear theme in) 95 of the 100 topics (complete labeling is generally not achieved). Since generally each constitution comprises a set of corpus documents we assign an overall constitutional weight for a topic as the average topic weight over the documents that it comprises. In Table 1 we list the ten topics with largest average topic weight, along with the ten most probable words (in decreasing order) for each topic. A full list of the topics, in order of average weight, with the weights of the top 20 words can be found at https://www.math.dartmouth.edu/~rockmore/topics_weight_order.txt.

Influence and clustering.

The identification of the topics now gives a natural way to represent a constitution as a mixture of probability distributions. With that, we can compare quantitatively constitutions and get at a quantitative notion of influence, completely driven by the data of the words. A first coarse pass at this is to create a constitutional "family tree", where the (unique) immediate ancestor of any given constitution is simply the constitution closest to it among all earlier constitutions.

[1] Technically, the pre-processing of the texts may result in some elements of the vocabulary set that are not words per se, but instead word stems, often called "tokens". We will use the more colloquial term "word" in this paper.

[2] This is a standard topic modeling document length, short enough to reflect local context and long enough to make sensible the statistical model.

Given that our constitutions are now represented as probability distributions (over topics), a natural measure of distance is the Kullback-Liebler (KL) divergence.[3] KL is inherently non-symmetric. A standard interpretation (see e.g., the Wikipedia entry[4]) is the degree to which a distribution Q approximates another distribution P. So thinking of an earlier constitution as a potential model for a newly written constitution, the KL divergence of their underlying topic probability distributions is a natural measure of similarity.

The "KL Constitution Tree" is shown in Figure 1. Note that the figure is not scaled horizontally for time. The size and form of the representation presents some difficulty for reproducing legibly herein, so a separate pdf document, readily magnifiable, can be found at https://www.math.dartmouth.edu/~rockmore/kl_tree.pdf. We also provide a detail in Figure 2.

The KL-tree is a coarse and aggregate articulation of the notion that constitutional ideas flow in time. The "flows" of the individual topics gets more at flow of individual (if overlapping) ideas, with the topics acting as the fundamental organizational unit of heredity. This local flow of topics is articulated in terms of the notion of an *information cascade*. We follow standard conventions [35] and define an *information cascade* (at a set threshold) as a collection of constitutions and their timestamps where a particular topic makes up a proportion greater than the threshold. When two constitutions (nodes) both express a topic above some threshold then we see this as the information "cascading" from the earlier to the later.

The topic cascades allow us to infer how ideas represented by topics are likely to have propagated through the corpus over time. As stated previously, we view the observation of a topic (above some threshold) in two constitutions as a quantitative measure indicating correlation across time. Given the content of the topics and the fact that they are ordered chronologically and typically clustered spatially, shared topics are likely to have spread from the earlier to the latter, but we do not know by which path this diffusion occurred. In order to learn the most likely propagation structure of the topics we estimate an underlying *diffusion network* for the corpus [36]. A *diffusion network* is a directed graph with nodes corresponding to constitutions and where the edges satisfy the condition that the source constitution predates the destination constitution. This imposes a weak causal structure on the correlations. Importantly, we do not observe the diffusion network, but only the cascades that are assumed to diffuse over it. In brief, we define a probabilistic model describing the consistency of the observed cascades with respect to a fixed diffusion network. The diffusion network is that which (approximately) maximizes this probability [36].

The presentation of the full diffusion tree on our corpus presents some visualization challenges. To give a sense of what it looks like, Figure 3 shows the entire learned diffusion tree on a restricted set of ninety-nine constitutions. Even this is too dense to be inspected visually for information, but the figure gives a good sense of the way in which the methodology reifies the phenomena of the idea diffusion. Each of the edges (directed and extending downward) indicate particular topics diffusing forward in time to be taken up by subsequent constitutions. Issues of readability make it impossible to put labels on the various edges. The optimization algorithm that produces the diffusion network only collects a subset of the topics that appear in a constitution. Some diffuse forward, others do not. The "offspring" of a given constitution thus borrow certain "ideas" of the parents, but others are created afresh, presumably depending on legally appropriate contextual factors.

[3] Given two probability distributions P and Q, $KL(P||Q) = \sum_i P(i) \log \frac{P(i)}{Q(i)}$.

[4] https://en.wikipedia.org/wiki/Kullback-Leibler_divergence

Network analysis

In order to discern patterns in the diffusion tree the network is subjected to a clustering analysis – picking out communities of constitutions by methods of community detection and optimal modularity in which groups of constitutions which share topics – and thereby a directed edge – in an amount above that expected by chance. Such a community constitutes a cluster [37]. Figure 4 displays the results of a network reconstruction of the full circuit along with two color codings of the network resulting from the application of two forms of clustering analysis to the network. The network is illustrated using spring embedding whereby densely connected nodes appear packed together. The network has the form of a ”constitutional caterpillar” with a temporal spine threaded through the network spanning 1789 to 2014. This temporal structure is very clear in the clustering coloring. Using community structure algorithms [38] we observe three clear constitutional communities, each of which describes a span of time: epoch 1: from 1789 to 1936; epoch 2: from 1937 to 1967; and epoch 3: from 1968 to 2014. Using a spectral technique for community detection we can further partition these network data into higher order communities [37]. This analysis maintains the chronological structure and illustrates the way in which clusters that are growing in absolute size (more consitutions in each) have evolved to encompass roughly decreasing ranges of time.

Each constitution in the diffusion tree can be described in terms of its *transmission motif – “t-motif”*, which visualizes the indegree and outdegree for each constitution. A selection of these motifs are shown in Figure 5 and Figure 6. The motifs demonstrate the variation to be found in balancing in-bound and out-bound influence for each constitution. Early constitutions tends to have few parents (Canada only has one – the US constitution) whereas subsequent constitutions vary significantly in their ancestry. This variation can be explained thorough a combination of both time (earlier constitutions present more opportunities for imitation) and how representative, novel and applicable each constitutions is as a model for imitation.

Models for transmission

We can gain further insights into the patterns of inheritance by studying directly the distributions of indegree and outdegree across the entire dataset, whose underlying data is shown in Figure 5. The distributions are shown in Figure 6. Figures 6A and 6B represent the pdf (probability density function) and cdf (cumulative distribution function) for the indegree for all constitutions. Illustrated in blue is the data and in orange the maximum likelihood parameter estimates for the best fitting distribution. The indegree distribution is well-captured by a Gaussian distribution with a mean of 8.8 and a standard deviation of 2.9. The estimated distribution does tend to slightly underestimate the mean but captures the tails very accurately. Outdegree is quite different. Figures 6C and 6D show the best fitting Poisson distribution and the outdegree distribution. Whereas the mean is effectively recovered, the tails of the distribution are poorly fitted; the Poisson underestimates the number of constitutions with few offspring and overestimates the number of constitutions with many offspring. On the other hand consider Figures 6E and 6F where we show the best fitting negative binomial distribution to the data. This very accurately recovers the entire offspring distribution with maximum likelihood parameter estimates for the two shape parameters of the distribution as $r = 2.5$ and $p = .22)$ Recall that for a negative binomial r describes the number of offspring observed before no more offspring are generated and that the probability of producing an offspring is given by the value of p. We view this as a pure birth process as constitutions never die – in the sense that they are always available as inspiration for a newly written constitution.

Methods

As outlined in the previous sections, our results and methodology depend on the use of topic models (see e.g., [22]) and diffusion networks. Topic models represent texts (embodied as a "bag-of-words") as mixtures of "topics", defined as probability distributions of the vocabulary in the corpus. The bag of words representing the constitution can be viewed as a mixture of topics, which are hidden and reflected by the words in it. Topic models are statistical models to learn the underlying structure of a corpus of documents. The underlying topics are represented as latent variables in a hierarchical Bayesian model. A generative model is assumed to be responsible for the observed documents and the word distributions of each topic. The topic proportions of each document can be learned via estimation of the posterior distribution of latent variables conditioned on the observed documents.

There are many flavors of topic model. We use the Latent Dirichlet Allocation (LDA) [34] probabilistic generative topic model. Details can be found in that reference. LDA depends on four parameters, the Dirichlet parameters α, η, the number of topics K and the document length l. The topics are learned via Markov chain Monte Carlo (MCMC) and in order to expedite the mixing of the related Markov chain, we fix α and η, and vary the value of K and l. The choice of appropriate number of topics and a document length l for a given corpus is a problem of *model selection*. We carried out 5-fold cross-validation to optimize for both parameters. Specifically, we split the corpus evenly into 5 parts, then run 5 iterations, in which each of the five parts will be held-out set once, and model estimation will be carried out on the remaining four parts. With the learned model, the likelihood of the held-out part is an measurement of the generalization ability of the model. Unfortunately, the computation of held-out likelihood $P(W|K,l)$ is intractable. We adopted the Chib-style estimation in [39] to approximate it. Finally, the best combination of K and l should have the highest overall likelihood over 5 folds. We fix l to be 500, which is the optimal choice affirmed by cross-validation. The number of topics that maximize the held-out likelihood is 100.

Inference for the diffusion network also depends on certain parameters. We follow [36] to find optimal choices for a fixed threshold. We fix a final threshold as that which simultaneously maximizes the mean indegree and outdegree in the diffusion network.

Discussion

We have searched for regular patterns of transmission in complex cultural artifacts. If there are cultural analogs to genotypes, and perhaps even phenotypes, we should be able to observe their signatures in a temporally resolved study of evolving documents. Much like organisms that adapt to local environments, constitutions must be adapted to local cultural and legal requirements. And as with organisms, a great deal of variability in constitutions has been documented or inferred as derived from ancestral documents.

The motifs (Figures 7 and 8) illustrate clearly how constitutions are "cultural recombinants" borrowing extensively from their ancestors. Constitutions vary in their hybridicity, with a few "asexual" constitutions borrowing from a single source, such as Australia (1901) and France (1791) both borrowing from the Last Universal Ancestor of Constitutions (LUCAC), the 1789 US Constitution, and several borrowing promiscuously from many ancestors, such as Hungary (1949) borrowing from seven ancestors, Greece (1975) borrowing from ten, and Egypt (1971) and Jordan (1952) borrowing from eleven. At the other end of the extreme are "infertile" constitutions with no descendants, such as Slovenia, Antigua and Barbados, and Myanmar. This diversity of heredity reflects the diversity of legal concepts each nation seeks to encode through its constitution. The motif variations suggest a constitution taxonomy, of *minor,* *major, idiosyncratic, and innovative* depending on where in the distribution the indegree and outdegree lie. Figure 9 represents these in a matrix. This diversity highlights an important difference from biology where species of organisms show far less variation in the basic mechanics of transmission.

That said, the distributions of the indegree and outdegree support different biological analogies. Consider again the striking result of the fit of the outdegree distribution to the negative binomial and the indegree to the Gaussian. A principled way to understand these distributions is to derive them from suitable stochastic processes. The Gaussian distribution arises naturally from the sum of independent random variables with a well defined mean and variance. Poisson distributions are attractors of the Galton-Watson process whereas negative binomial distributions are attractors of the Yule process (see e.g.[40]). Both Poisson and negative binomial offspring distributions are observed frequently in biological systems. The Galton-Watson process was derived to explain the extinction of family names. The idea is that at each generation a parent can transmit their name to either $0, 1, \ldots, n$ offspring. Each parent samples the number of offspring independently from the same distribution. Our data support a negative binomial distribution so we shall focus on the Yule process. The Yule process is also known as a preferential attachment process as it can be derived from an "urn process" in which balls of a given color are sampled in linear proportion to the number of balls already in each urn. The negative binomial distribution is derived by solving a simple recurrence equation describing the temporal evolution of a probability distribution of the form,

$$P_n'(t) = -n\lambda P_n(t) + (n-1)\lambda P_{n-1}(t).$$

Here $P_n(t)$ is the probability of finding n constitutions at time t. The rate of offspring production in some interval δt is parameterized by λ. Hence at a time t a number n of constitutions will decline through the addition of more offspring proportional to $n\lambda P_n(t)$ and increase through the production of offspring by the class $n-1$ at a rate $(n-1)\lambda P_{n-1}(t)$. If we establish an initial condition as the number of constitutions at the start of constitutional history as 1, $P_0(0) = 1$, we find that,

$$P_n(t) = (\frac{n-1}{n-n_0})e^{-\lambda n_0 t}(1-e^{-\lambda t})^{n-n_0}.$$

Which takes the form of the negative binomial distribution in which we observe exactly n_0 offspring in n trials with a success probability, $p = e^{-\lambda t}$.

We can test the assumptions of the Yule process by looking directly at the imitation dynamics of any given constitution. We simply plot the date on which the descendant of a given constitution was created against the order in which it was created. In Figure 10A we look at the evolution of the first 20 constitutions. By far the majority have fewer than 10 offspring and these offspring span a range of under 50 years. However a few of these constitutions are exceptional. The most remarkable is the 1813 constitution of Paraguay that has influenced 70 descendant constitutions and remained influential over 200 years. This is followed by the original constitution of the Unites States of America from 1789 that produces 20 descendant constitutions and remained influential for around 80 years. The Canadian constitution of 1791 produces 11 descendants but remained influential for over 150 years. Figure 10B includes the first 100 constitutions, 10C the first 200 , and 10D all 591 in the data set. A clear relationship between offspring number and longevity emerges consistent with preferential attachment in which a small number of constitutions are of dominant influence, these appeared early in constitutional history gaining a significant foothold, and with the vast majority of constitutions both short lived and producing less than 10 offspring.

The analysis of cultural recombination through a principled decomposition of textual artifacts suggests new domains of cultural inheritance. Unlike simple Mendelian systems, or simple learning models with homogeneous rules, we observe diverse patterns of variation in the way in which nations encode important moral and legal principles.

Moreover we can obtain a principled definition of a meme that goes beyond the "word" and captures highly linked sets of words expressing a functional, legal category – much the way a gene, composed of linked sets of nucleotides – contributes to a function. We would argue that the semantic interpretation of a given constitution and its practical legal impact is what we might mean by the phenotype. Nations differ in their debt to the past and their original contributions to the future. This allows us to speak in a rigorous fashion about phylogenetic concepts like analogy and homology when it comes to cultural artifacts.

Reconciling statistical patterns of influence with potential biases and patterns in thinking and writing will bring us closer to frameworks that connect methods of mathematical science with objects of humanistic interest in the service of new models and theories of cultural transmission and influence.

References

[1] Hart DL, Clark AG (1997) Principles of Population Genetics. Sinauer Associates, Inc Publishers, Mass.

[2] Rogers EM (1995) Diffusion of Innovations. The Free Press, NYC.

[3] Christiansen FB (2008) Theories of Population Variation in Genes and Genomes. Princeton University Press, Princeton, NJ.

[4] Sforza LLC, Fedlman MW (1981) Cultural Transmission and Evolution: A Quantitative Approach. Princeton University Press, Princeton, NJ.

[5] Richerson PJ, Boyd R (2006) Not by Genes Alone: How Culture Transformed Human Evolution. University Of Chicago Press.

[6] Rogers DS, Ehrlich PR (2007) Natural selection and cultural rates of change. Proceedings of the National Academy of Sciences 105: 3416–3420.

[7] Pagel MD (2013) Wired for Culture. W.W. Norton & Company.

[8] Wimsatt WC (1999) Genes, memes, and cultural heredity. Biology and Philosophy 14: 279–310.

[9] Kaiser DI (2009) Drawing theories apart: The dispersion of Feynman diagrams in postwar physics. University of Chicago Press.

[10] The Google books website. URL `http://books.google.com/`. Accessed April, 2013.

[11] International Music Score Library Project. URL `http://imslp.org/`. Accessed July, 2013.

[12] ARTstor. URL `http://www.artstor.org`. Accessed April, 2013.

[13] Michel JB, Shen YK, Aiden AP, Veres A, Gray MK, et al. (2010) Quantitative analysis of culture using millions of digitized books. Science 331: 176-182.

[14] The Google books N-Gram Viewer. URL `http://books.google.com/ngrams`. Accessed April, 2013.

[15] Moretti F (2005) Graphs, Maps, Trees. Verso Books.

[16] Jockers M (2013) Macroanalysis: Digital Methods and Literary History. University of Illinois Press.

[17] Hughes JM, Foti NJ, Krakauer DC, Rockmore DN (2012) Quantitative patterns of stylistic influence in the evolution of literature. Proceedings of the National Academy of Sciences 109: 7682–7686.

[18] Riddell A (2014) How to Read 22,198 Journal Articles: Studying the History of German Studies with Topic Models. In: Erlin M, Tatlock L, editors, Distant Readings: Topologies of German Culture in the Long Nineteenth Century. Rochester, NY, USA: Camden House, pp. 91–114.

[19] Newman DJ, Block S (2006) Probabilistic topic decomposition of an eighteenth-century american newspaper. Journal of the American Society for Information Science and Technology 57: 753–767.

[20] Stanford Topic Modeling Toolbox. URL `http://nlp.stanford.edu/software/tmt/tmt-0.4/`. Accessed July, 2013.

[21] Mallet: MAchine Learning for LanguagE Toolkit. URL `http://mallet.cs.umass.edu/topics.php`. Accessed July, 2013.

[22] Blei DM (2012) Probabilistic topic models. Communications of the ACM 55: 77-84.

[23] Blei DM, Lafferty JD (2006) Dynamic topic models. In: Proceedings of the 23rd International Conference on Machine Learning. New York, NY, USA: ACM, ICML '06, pp. 113–120. doi:10.1145/1143844.1143859. URL `http://doi.acm.org/10.1145/1143844.1143859`.

[24] Wang X, McCallum A (2006) Topics over time: A non-Markov continuous-time model of topical trends. In: Proceedings of the 12th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining. New York, NY, USA: ACM, KDD '06, pp. 424–433. doi:10.1145/1150402.1150450. URL `http://doi.acm.org/10.1145/1150402.1150450`.

[25] Boyd R, Richerson PJ (1996) Why culture is common but cultural evolution is rare. Proceedings of the British Academy 88: 73-930.

[26] Law DS, Versteeg M (2011) The evolution and ideology of global constitutionalism. Cal Law Review 99: 1163.

[27] Foti N, Ginsburg T, Rockmore D (2014) 'We the Peoples': The global origins of constitutional preambles. George Washington International Law Review, 46: 101–134.

[28] Elkins Z, Ginsburg T, Melton J (2009) The Endurance of National Constitutions. Cambridge University Press.

[29] Elkins Z, Ginsburg T, Simmons B (2013) Getting to rights: Constitutions and international law. Harvard International Law Journal 51: 201–34.

[30] Go J (2003) A globalizing constitutionalism? Views from the postcolony 1945-2000. International Sociology 18: 71-95.

[31] Boli-Bennett J (1987) Human rights or state expansion? Cross-national definitions of constitutional rights, 1870-1970. In: Thomas G, Meyer J, Ramirez F, Boli J, editors, Institutional Structure, Sage. pp. 71–91.

[32] Law D (2005) Generic constitutional law. Minn L Rev 89: 652.

[33] Watson A (1974) Legal Transplants. Cambridge University Press.

[34] Blei DM, Ng AY, Jordan MI (2003) Latent Dirichlet allocation. Journal of Machine Learning Research 3: 993-1022.

[35] Leskovec J, McGlohon M, Faloutsos C, Glance NS, Hurst M (2007) Patterns of cascading behavior in large blog graphs. In: SDM. pp. 551-556.

[36] Gomez-Rodriguez M, Leskovec J, Krause A (2012) Inferring networks of diffusion and influence. TKDD 5: 21.

[37] Newman M (2006) Modularity and community structure in networks. Proceedings of National Academy of Sciences, USA 103: 8577-8582.

[38] Girvan M, Newman ME (2002) Community structure in social and biological networks. Proceedings of the National Academy of Sciences 99: 7821–7826.

[39] Wallach HM, Murray I, Salakhutdinov R, Mimno DM (2009) Evaluation methods for topic models. In: ICML. p. 139.

[40] Karlin S, Taylor HM (1975) A First Course in Stochastic Processes. Academic Press.

Topic name	Top 10 words in topic
General rights	right rights citizens freedom law public guaranteed citizen everyone religious
Sovereignty	national people sovereignty law rights state flag language international equal
public order	law public cases order one property laws authority liberty civil
allocation of powers	congress executive laws power ministers state secretaries order necessary public
misc v	law government president organization national organic public laws social functioning
socialism II	people socialist country revolution working popular citizens system society development
legislative sessions	session deputies sessions deputy members elected first vote majority extraordinary
bureaucracy	papers years state department necessary respective individuals departments body power
Socialism I	people organs state supreme work organ presidium elected decisions committees

Table 1: Most popular topics across entire corpus, and their corresponding top 10 words. A full list of the topics, in order of average weight, with the weights of the top 20 words can be found at https://www.math.dartmouth.edu/∼rockmore/topics_weight_order.txt.

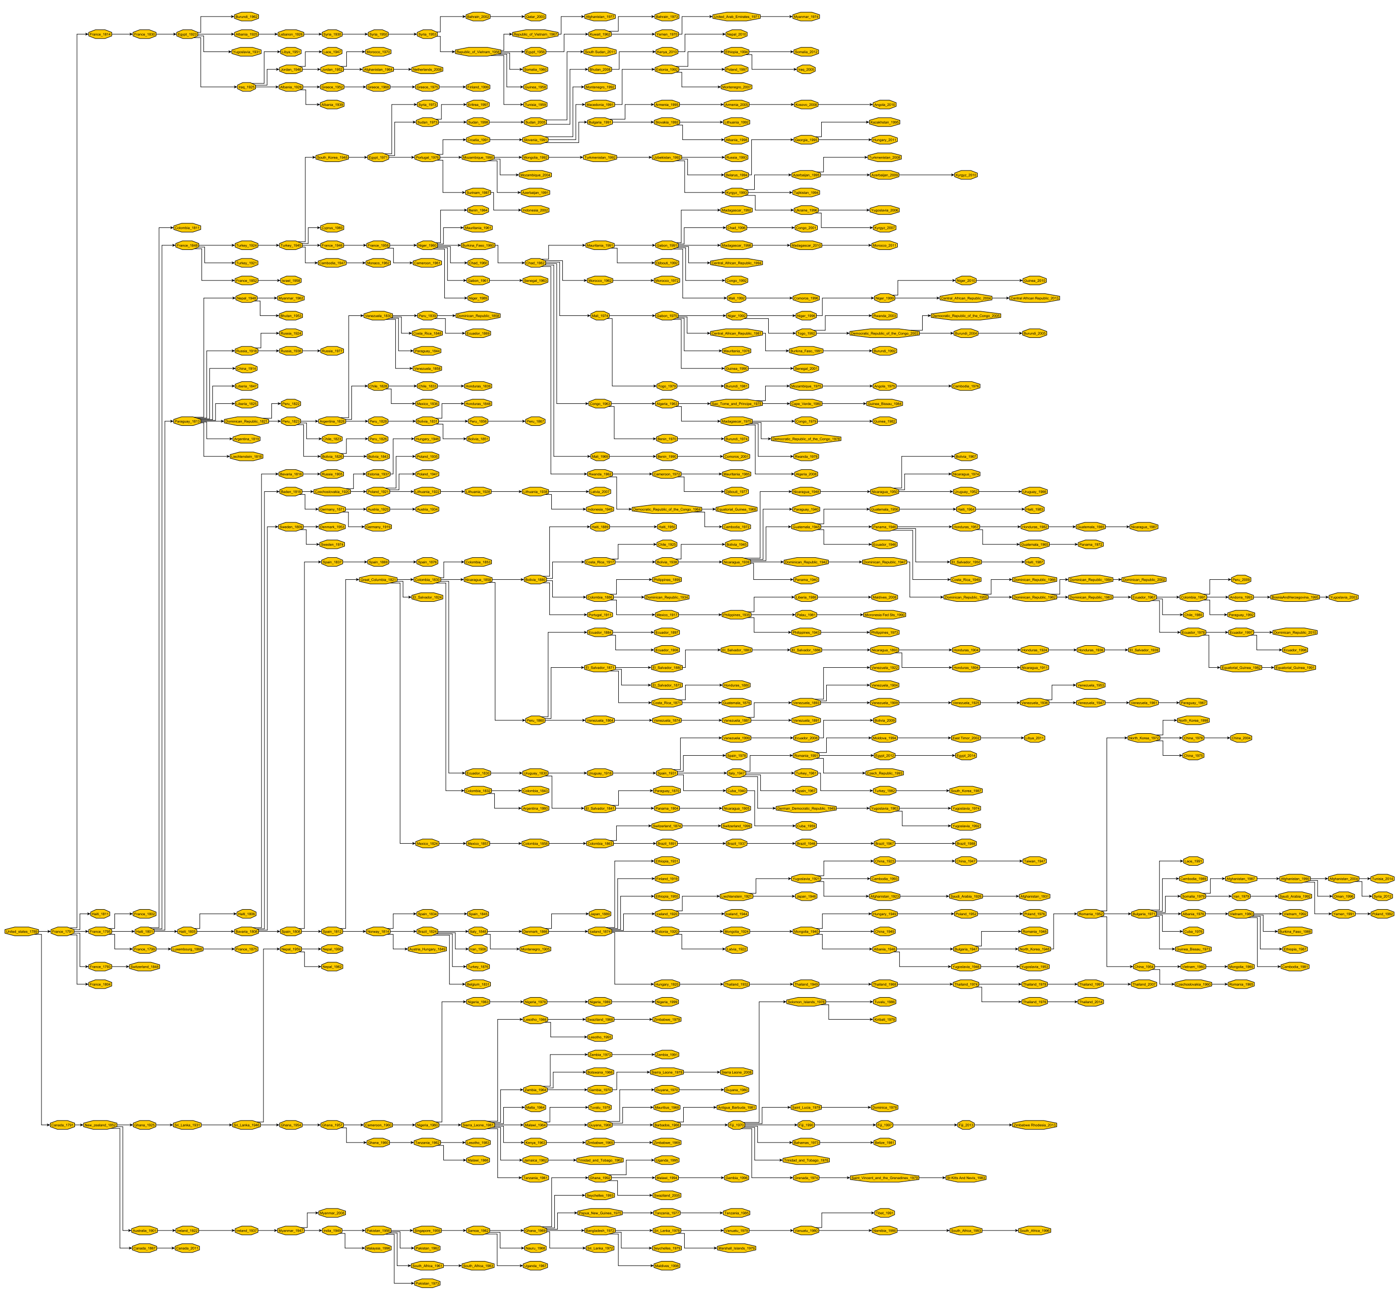

Figure 1: The "family tree" of constitutions. The United States constitution of 1789 is the root and thus the "Last Universal Common Ancestor Constitution". Any other constitutions is deemed as having as its most recent ancestor the closest earlier constitutions where distance is measured as the KL-divergence of the former to the latter. A pdf document of this tree, easily magnifiable, can be found at https://www.math.dartmouth.edu/∼rockmore/kl_tree.pdf.

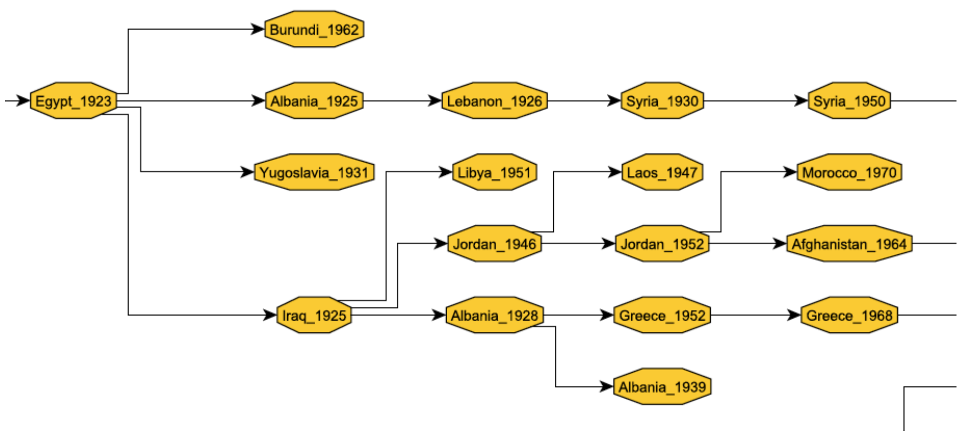

Figure 2: A detail of the tree around the Egyptian constitution of 1923. Note the fertility of that constitution, as well as the sterility of the constitutions of Burundi (1962), Morocco (1970), and Albania (1939). The last of these is particularly interesting as we see a line of descendants issuing forth from the Albania constitution of 1925. So, earlier versions of constitutions can have patterns of transmission that do not include all of their descendants. A pdf document of the full tree, easily magnifiable, can be found at https://www.math.dartmouth.edu/∼rockmore/kl_tree.pdf.

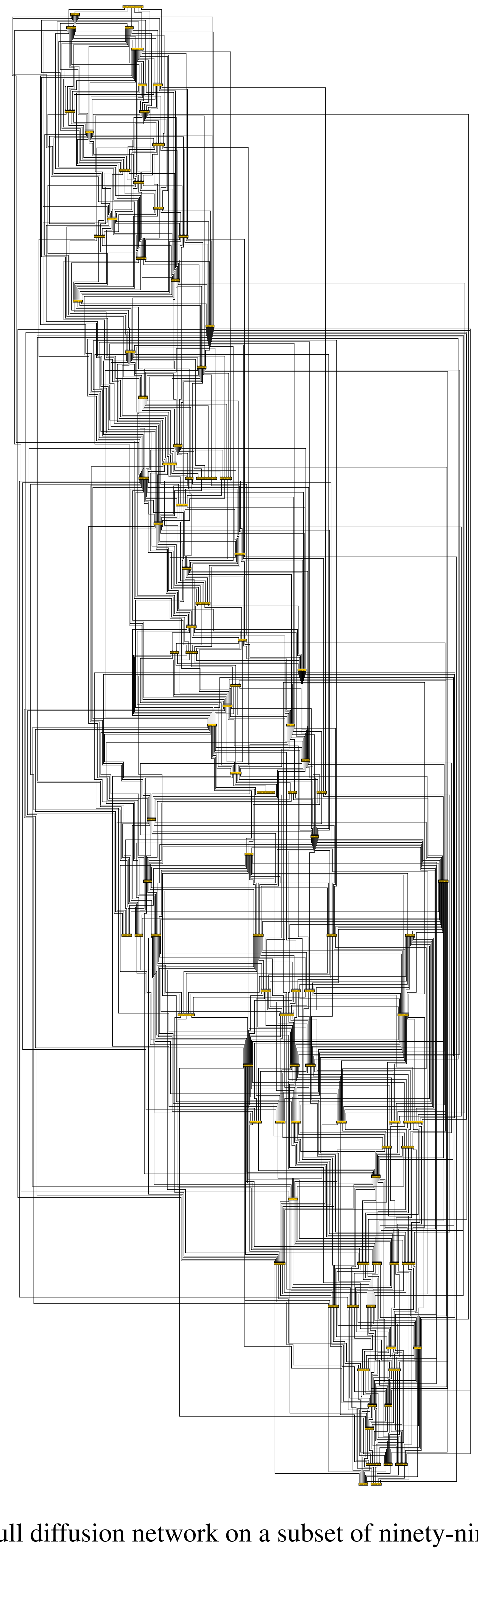

Figure 3: The full diffusion network on a subset of ninety-nine constitutions.

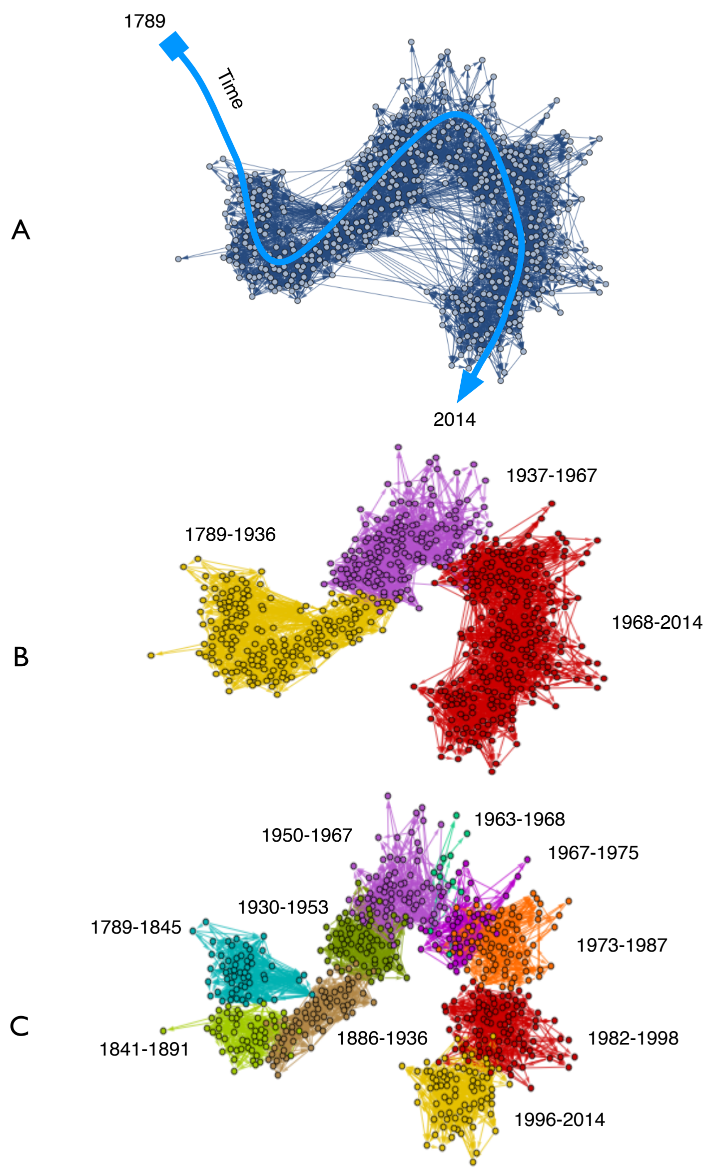

Figure 4: (A) Spring embedded reconstruction of constitutional diffusion network. Nodes correspond to constitutions and directed edges encode topic borrowing. The blue arrow traces time forward through the network starting in 1789 and ending in 2014. Time is the dominant factor in explaining the geometric form of the network. (B) Application of a community detection algorithm [38] to the thresholded diffusion tree reveals three clear epochs of constitutional inheritance. The oldest epoch spans 147 years and contains 175 constitutions generating an average of 1.2 constitutions per year. The second epoch spans 30 years and contains 148 constitutions and an average of 4.9 constitutions per year. The third epoch spans 46 years and contains 267 constitutions and an average of 5.8 constitutions per year. The rate at which constitutions are being written has increased through time whereas the temporal influence of constitutions into the future has contracted. (C) Use of more sensitive optimal modularity methods [37] provides a decomposition of each of these epochs into a further three epochs. Each induced cluster preserves the largely temporally contiguous ordering demonstrating that time remains a dominant dimension of variation at the microscopic level.

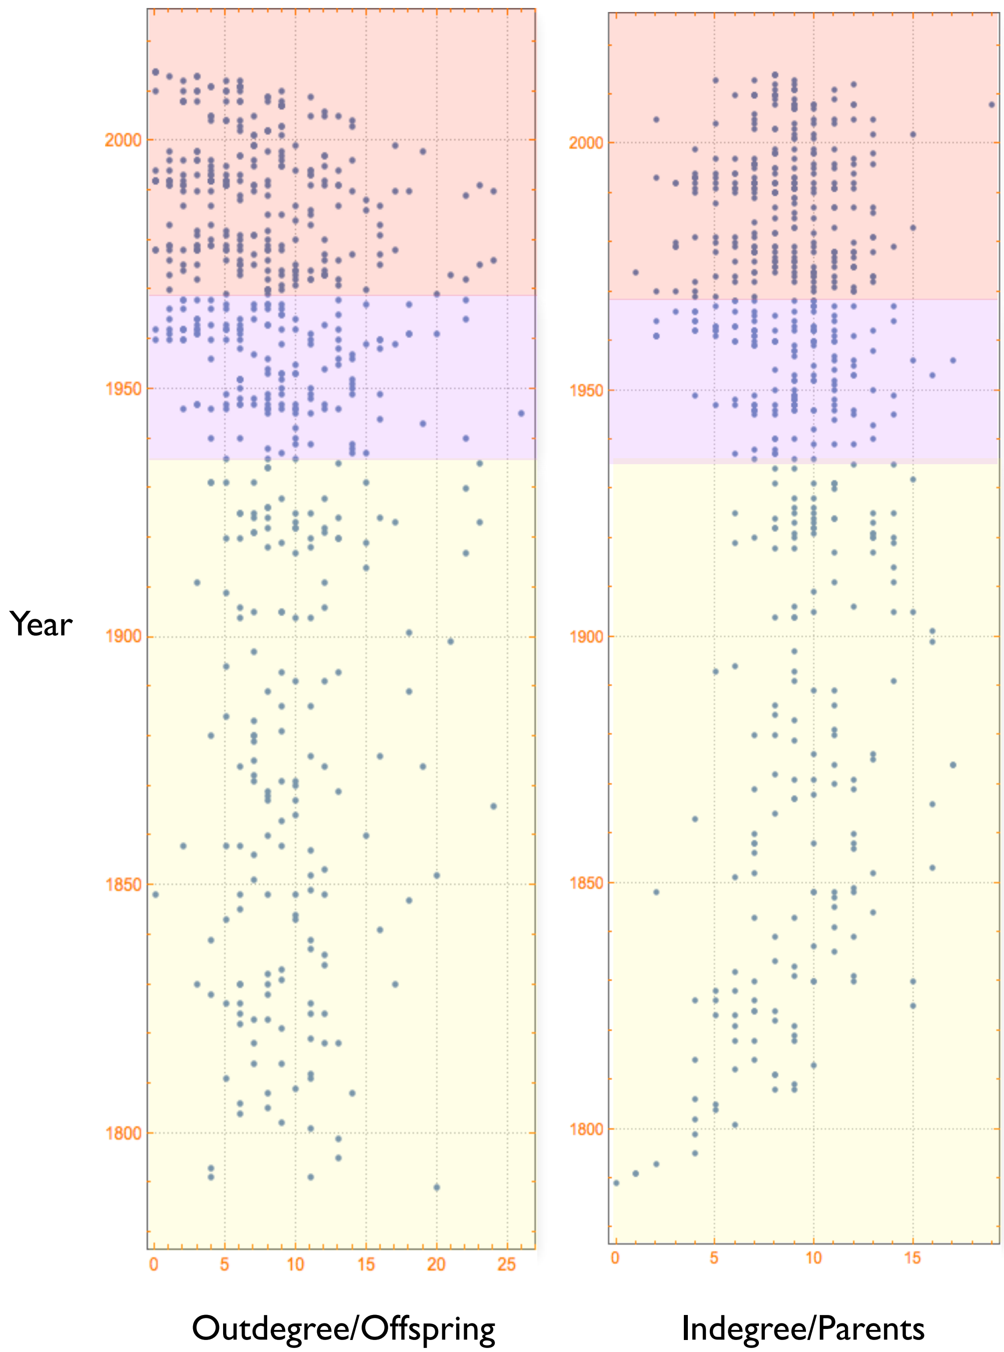

Figure 5: The relationship of outdegree and indegree against time of writing color-coded into three constitutional epochs. We see how the longest epoch is the sparsest and the most recent epoch the densest. The outdegree of the most recent constitutions tend to show a significantly greater concentration for 5 or less offspring. The average indegree increased over the course of the first epoch but was stable by the second epoch.

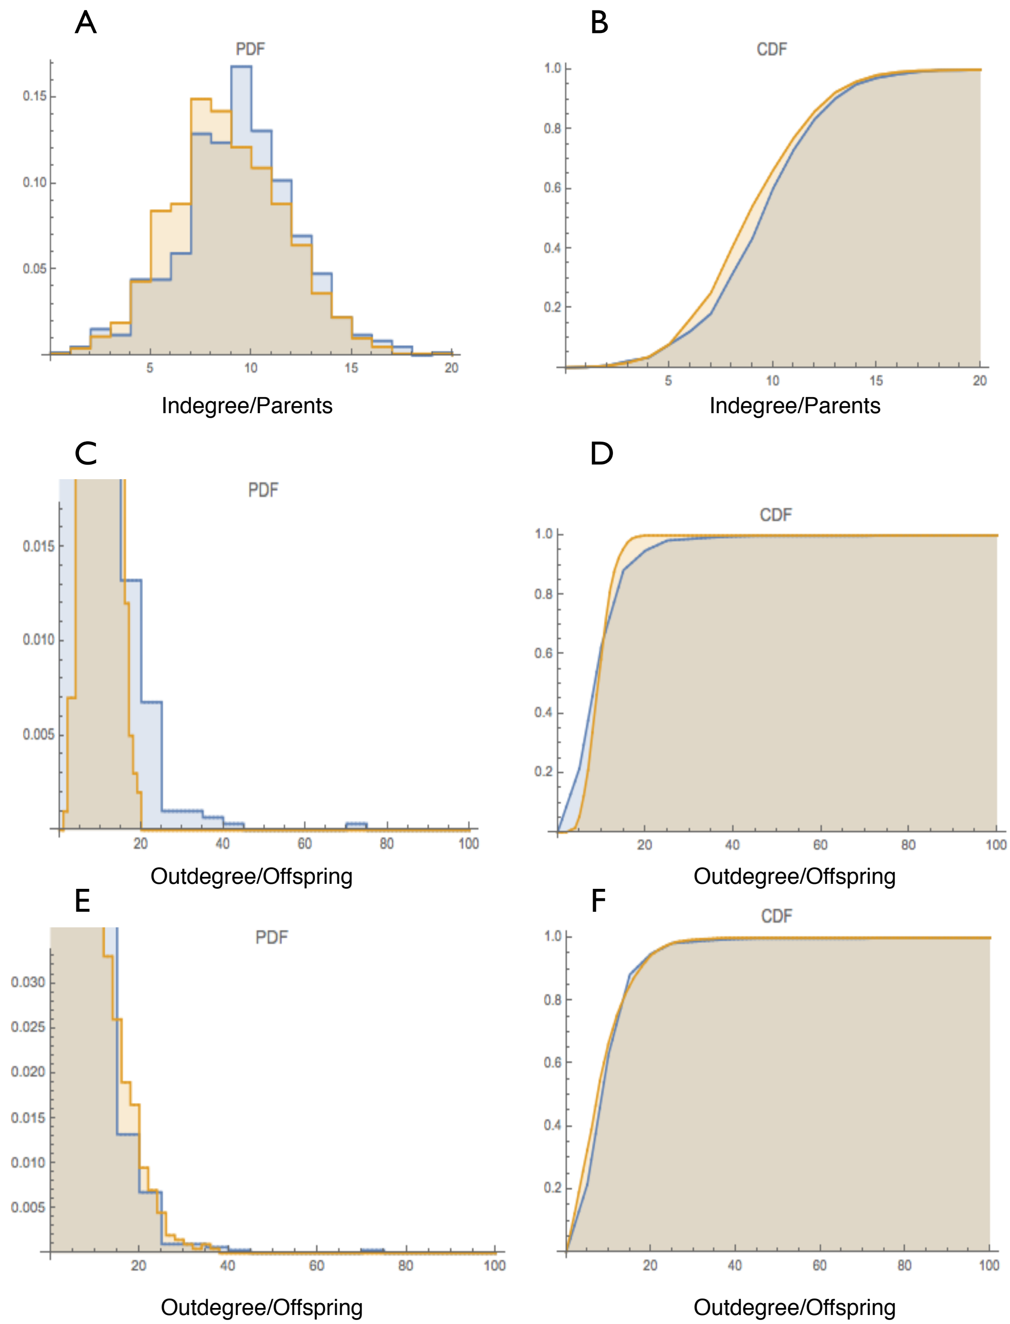

Figure 6: Illustrated in blue are the inferred connectivity data and in orange the maximum likelihood parameter estimates for the best fitting distributions for constitutional indegree (A,B) and outdegree (C,D,E,F). The indegree is best approximated by a Gaussian distribution with a mean of 8.8 and a standard deviation of 2.9. Figures 4C and 4D plot the outdegree distributions and the best fitting Poisson distribution. Whereas the mean is effectively recovered, the tails of the distribution are poorly fitted. The Poisson underestimates the number of constitutions with few offspring and overestimates the number of constitutions with many offspring. In Figure 4E and 4F we show the best fitting negative binomial distribution to the data. This very accurately recovers the entire offspring distribution with maximum likelihood parameter estimates for the two shape parameters of the distribution as $r = 2.5$ and $p = .22$.

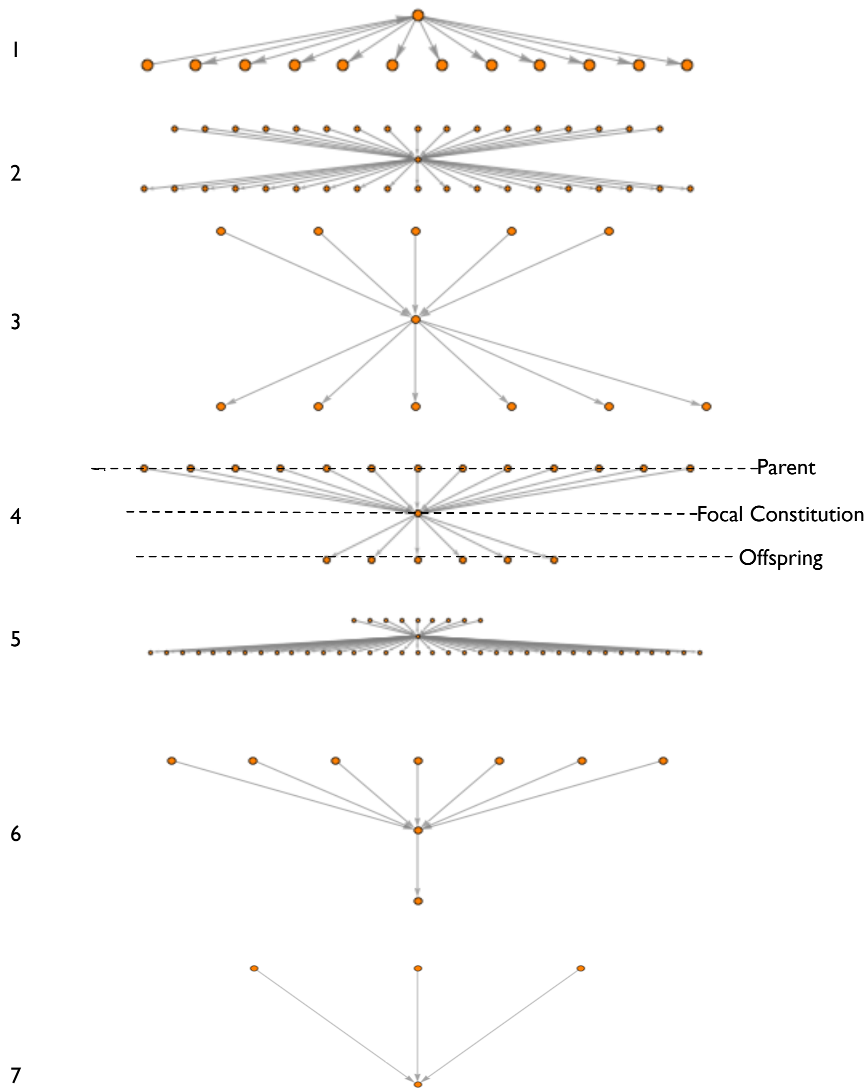

Figure 7: Each constitution in the diffusion tree can be described in terms of a *transmission motif*, which visualizes the indegree and outdegree for each target constitution. The motifs demonstrate the balance between in-bound and out-bound influence for each constitution in terms of a threshold number of topics that are borrowed. (1) Early constitutions tends to have few parents, e.g., Canada (1791) only has one (the U.S.(1789) constitution, the leftmost node). Subsequent constitutions vary significantly in their ancestry: (2) Iceland (1874)'s constitution has many parents and many offspring; (3) Bolivia (1826) constitution has fewer parents and few offspring (4) Venezuela (1830) exhibits many parents and few offspring; (5) South Korea (1948) has few parents and many offspring; (6) Albania (1976) has several parents and only one offspring; (7) Montenegro (1992) has no offspring. This variation in parentage and fertility can be explained thorough a combination of both the time at which they were written and the tendency to preferentially attach to a small number of highly favored models for imitation.

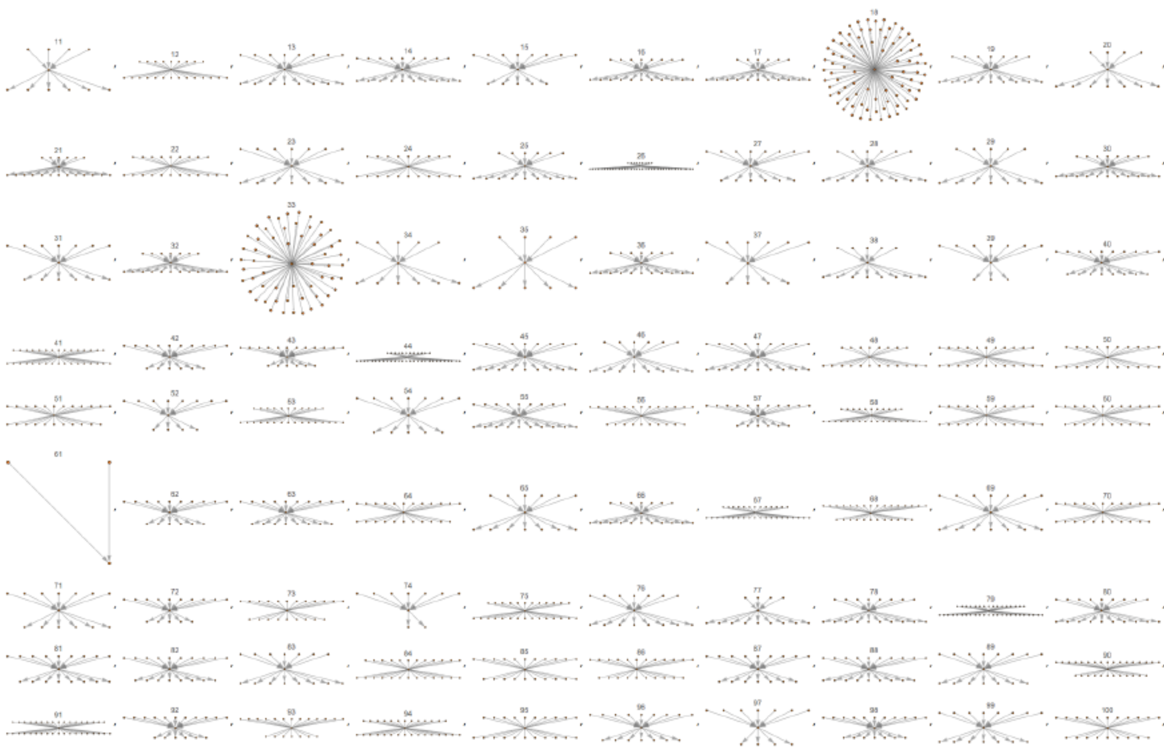

Figure 8: A larger sample of motifs.

Constitutional Influence	Below Mean Parents	Above Mean Parents
Below Mean Offspring	Minor	Idiosyncratic
Above Mean Offspring	Innovative	Major

Figure 9: A taxonomy of constitutions derived from the co-consideration of the place in the distributions of the indegree and outdegree of a given constitution.

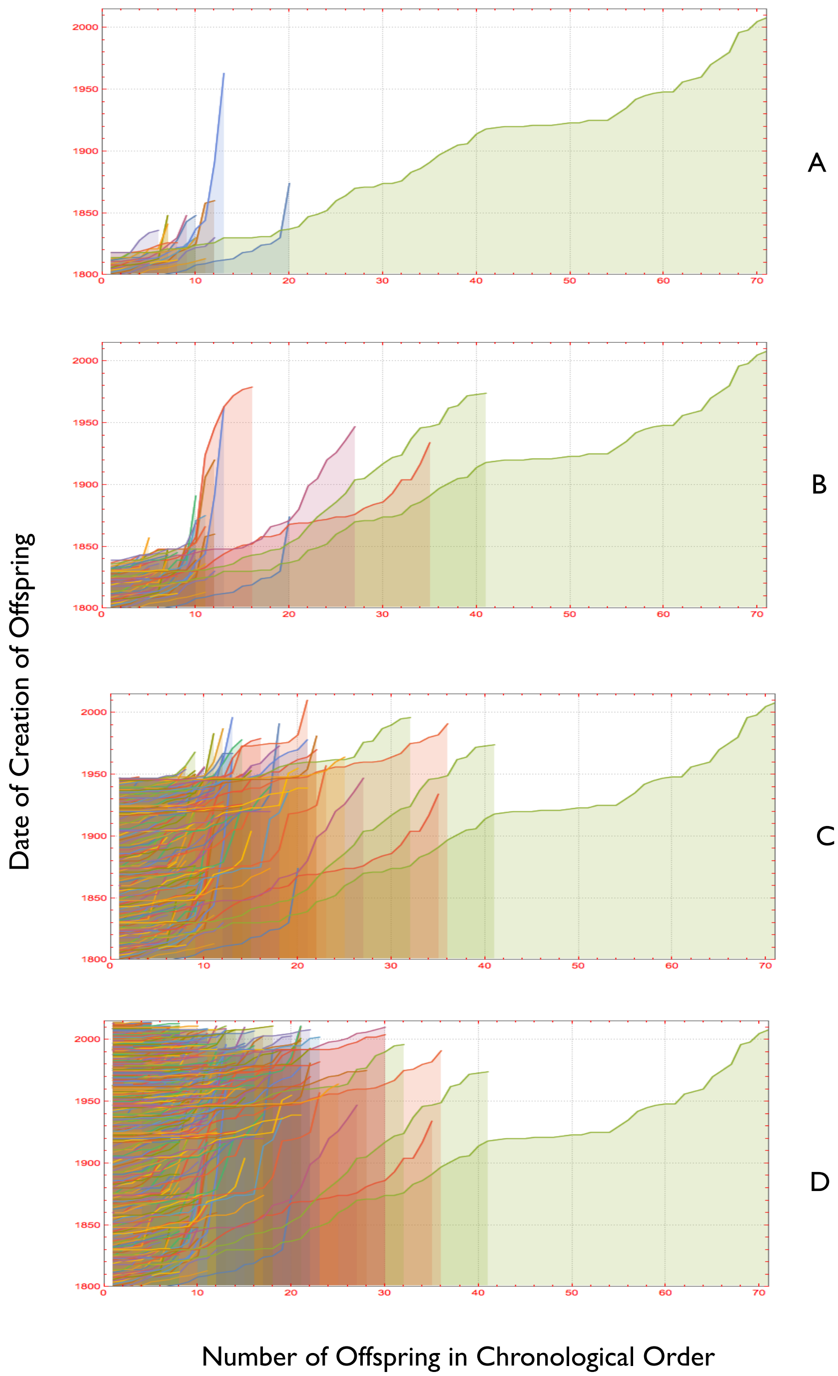

Figure 10: Fecundity and longevity of constitutions. On the x-axis are the number of descendant constitutions arranged in chronological order and on the y-axis the date of their appearance. In Panel A we plot the first 10 constitutions. In Panel B the first 50. Panel C the first 200. Panel D all 591. Most constitutions have few descendants and these appear over a relatively short span of time. Constitutions with many descendants tend to span longer periods of time. Most of the longest-lived constitutions in terms of influence/borrowing were written in the first of the three epochs of constitutional history (as in Figure 4A).

Afghanistan_1923	Bangladesh_1972	Burundi_1962
Afghanistan_1931	Barbados_1966	Burundi_1974
Afghanistan_1964	Bavaria_1808	Burundi_1981
Afghanistan_1977	Bavaria_1818	Burundi_1992
Afghanistan_1987	Belarus_1994	Burundi_2004
Afghanistan_1990	Belgium_1831	Burundi_2005
Afghanistan_2004	Belize_1981	Cambodia_1947
Albania_1925	Benin_1964	Cambodia_1972
Albania_1928	Benin_1970	Cambodia_1976
Albania_1939	Benin_1990	Cambodia_1981
Albania_1946	Bhutan_1953	Cambodia_1989
Albania_1976	Bhutan_2008	Cambodia_1993
Albania_1998	Bolivia_1826	Cameroon_1960
Algeria_1963	Bolivia_1831	Cameroon_1961
Algeria_2008	Bolivia_1843	Cameroon_1972
Andorra_1993	Bolivia_1851	Canada_1791
Angola_1975	Bolivia_1880	Canada_1867
Angola_2010	Bolivia_1938	Canada_2011
Antigua_Barbuda_1981	Bolivia_1945	Cape_Verde_1980
Argentina_1819	Bolivia_1967	Central African Republic_2013
Argentina_1826	Bolivia_2009	Central_African_Republic_1981
Argentina_1860	BosniaAndHerzegovina_1995	Central_African_Republic_1994
Armenia_1995	Botswana_1966	Central_African_Republic_2004
Armenia_2005	Brazil_1824	Chad_1960
Australia_1901	Brazil_1891	Chad_1962
Austria_1920	Brazil_1937	Chad_1996
Austria_1934	Brazil_1946	Chile_1823
Austria_Hungary_1849	Brazil_1967	Chile_1828
Azerbaijan_1991	Brazil_1988	Chile_1833
Azerbaijan_1995	Bulgaria_1947	Chile_1925
Azerbaijan_2009	Bulgaria_1971	Chile_1980
Baden_1818	Bulgaria_1991	China_1914
Bahamas_1973	Burkina_Faso_1960	China_1923
Bahrain_1973	Burkina_Faso_1988	China_1947
Bahrain_2002	Burkina_Faso_1991	China_1949

Table 2a. Constitution list part 1.

China_1954	Dominican_Republic_1821	Equatorial_Guinea_1991
China_1975	Dominican_Republic_1858	Eritrea_1997
China_1978	Dominican_Republic_1934	Estonia_1920
China_2004	Dominican_Republic_1942	Estonia_1937
Colombia_1811	Dominican_Republic_1947	Estonia_1992
Colombia_1830	Dominican_Republic_1955	Ethiopia_1931
Colombia_1832	Dominican_Republic_1962	Ethiopia_1955
Colombia_1843	Dominican_Republic_1963	Ethiopia_1987
Colombia_1853	Dominican_Republic_1966	Ethiopia_1994
Colombia_1858	Dominican_Republic_1994	Fiji_1970
Colombia_1863	Dominican_Republic_2002	Fiji_1990
Colombia_1886	Dominican_Republic_2010	Fiji_1997
Colombia_1991	East Timor_2002	Fiji_2013
Comoros_1996	Ecuador_1830	Finland_1919
Comoros_2001	Ecuador_1869	Finland_1999
Congo_1963	Ecuador_1884	France_1791
Congo_1979	Ecuador_1897	France_1793
Congo_1992	Ecuador_1906	France_1795
Congo_2001	Ecuador_1946	France_1799
Costa_Rica_1848	Ecuador_1967	France_1802
Costa_Rica_1871	Ecuador_1978	France_1804
Costa_Rica_1917	Ecuador_1997	France_1814
Costa_Rica_1949	Ecuador_1998	France_1830
Croatia_1991	Ecuador_2008	France_1848
Cuba_1940	Egypt_1923	France_1852
Cuba_1959	Egypt_1956	France_1875
Cuba_1976	Egypt_1971	France_1946
Cyprus_1960	Egypt_2012	France_1958
Czech_Republic_1993	Egypt_2014	Gabon_1961
Czechoslovakia_1920	El_Salvador_1824	Gabon_1975
Czechoslovakia_1960	El_Salvador_1841	Gabon_1991
Dem_Rep_of_the_Congo_1964	El_Salvador_1871	Gambia_1970
Dem_Rep_of_the_Congo_1978	El_Salvador_1872	Gambia_1996
Dem_Rep_of_the_Congo_2003	El_Salvador_1880	Georgia_1995
Dem_Rep_of_the_Congo_2005	El_Salvador_1883	GDR_1949
Denmark_1866	El_Salvador_1886	Germany_1871
Denmark_1953	El_Salvador_1939	Germany_1919
Djibouti_1977	El_Salvador_1950	Ghana_1925
Djibouti_1992	Equatorial_Guinea_1968	Ghana_1954
Dominica_1978	Equatorial_Guinea_1982	Ghana_1957

Table 2b. Constitution list part 2.

Ghana_1960	Hungary_1920	Lesotho_1993
Ghana_1969	Hungary_1946	Liberia_1825
Ghana_1992	Hungary_1949	Liberia_1847
Great_Columbia_1821	Hungary_2011	Liberia_1986
Greece_1952	Iceland_1874	Libya_1951
Greece_1968	Iceland_1920	Libya_2011
Greece_1975	Iceland_1944	Liechtenstein_1818
Grenada_1974	India_1949	Liechtenstein_1921
Guatemala_1879	Indonesia_1945	Lithuania_1922
Guatemala_1945	Indonesia_2002	Lithuania_1928
Guatemala_1956	Iran_1906	Lithuania_1938
Guatemala_1965	Iran_1979	Lithuania_1992
Guatemala_1985	Iraq_1925	Luxembourg_1868
Guinea_1958	Iraq_2005	Macedonia_1991
Guinea_1982	Ireland_1922	Madagascar_1975
Guinea_1990	Ireland_1937	Madagascar_1992
Guinea_2010	Israel_1958	Madagascar_1998
Guinea_Bissau_1973	Italy_1848	Madagascar_2010
Guinea_Bissau_1984	Italy_1947	Malawi_1964
Guyana_1966	Jamaica_1962	Malawi_1966
Guyana_1970	Japan_1889	Malawi_1994
Guyana_1980	Japan_1946	Malaysia_1996
Haiti_1801	Jordan_1946	Maldives_1998
Haiti_1805	Jordan_1952	Maldives_2008
Haiti_1806	Kazakhstan_1995	Mali_1969
Haiti_1811	Kenya_1963	Mali_1974
Haiti_1889	Kenya_2010	Mali_1992
Haiti_1950	Kiribati_1979	Malta_1964
Haiti_1964	Kosovo_2008	Marshall_Islands_1979
Haiti_1983	Kuwait_1962	Mauritania_1961
Haiti_1987	Kyrgyz_1993	Mauritania_1978
Honduras_1839	Kyrgyz_2007	Mauritania_1985
Honduras_1848	Kyrgyz_2010	Mauritania_1991
Honduras_1880	Laos_1947	Mauritius_1968
Honduras_1894	Laos_1991	Mexico_1824
Honduras_1904	Latvia_1922	Mexico_1836
Honduras_1924	Latvia_2007	Mexico_1857
Honduras_1936	Lebanon_1926	Mexico_1917
Honduras_1957	Lesotho_1966	Micronesia Fed Sts_1990
Honduras_1982	Lesotho_1983	Moldova_1994

Table 2c. Constitution list part 3.

Monaco_1962	Niger_1996	Philippines_1973
Mongolia_1924	Niger_1999	Poland_1921
Mongolia_1940	Niger_2010	Poland_1935
Mongolia_1960	Nigeria_1960	Poland_1947
Mongolia_1992	Nigeria_1963	Poland_1952
Montenegro_1905	Nigeria_1978	Poland_1976
Montenegro_1992	Nigeria_1989	Poland_1992
Montenegro_2007	Nigeria_1999	Poland_1997
Morocco_1962	North_Korea_1948	Portugal_1911
Morocco_1970	North_Korea_1972	Portugal_1976
Morocco_1972	North_Korea_1998	Qatar_2003
Morocco_2011	Norway_1814	Republic_of_Vietnam_1956
Mozambique_1975	Oman_1996	Republic_of_Vietnam_1967
Mozambique_1990	Pakistan_1956	Romania_1948
Mozambique_2004	Pakistan_1962	Romania_1952
Myanmar_1947	Pakistan_1973	Romania_1965
Myanmar_1962	Palau_1981	Romania_1991
Myanmar_1974	Panama_1904	Russia_1905
Myanmar_2008	Panama_1940	Russia_1918
Namibia_1990	Panama_1946	Russia_1924
Nauru_1968	Panama_1972	Russia_1936
Nepal_1948	Papua_New_Guinea_1975	Russia_1977
Nepal_1959	Paraguay_1813	Russia_1993
Nepal_1962	Paraguay_1844	Rwanda_1962
Nepal_1990	Paraguay_1870	Rwanda_1978
Nepal_2010	Paraguay_1940	Rwanda_2003
Netherlands_2008	Paraguay_1967	Saint_Lucia_1978
New_zealand_1852	Paraguay_1992	Saint_Vincent_and_the_Grenadines_1979
Nicaragua_1858	Peru_1822	Samoa_1962
Nicaragua_1893	Peru_1823	Sao_Tome_and_Principe_1975
Nicaragua_1905	Peru_1826	Saudi_Arabia_1926
Nicaragua_1911	Peru_1828	Saudi_Arabia_1992
Nicaragua_1939	Peru_1839	Senegal_1963
Nicaragua_1948	Peru_1856	Senegal_2001
Nicaragua_1950	Peru_1860	Seychelles_1979
Nicaragua_1974	Peru_1867	Seychelles_1993
Nicaragua_1987	Peru_2009	Sierra Leone_2008
Niger_1960	Philippines_1899	Sierra_Leone_1961
Niger_1989	Philippines_1935	Sierra_Leone_1978
Niger_1992	Philippines_1943	Singapore_1959

Table 2d. Constitution list part 4.

Slovakia_1992	Syria_1973	Uruguay_1952
Slovenia_1991	Syria_2012	Uruguay_1966
Solomon_Islands_1978	Taiwan_1947	Uzbekistan_1992
Somalia_1960	Tajikistan_1994	Vanuatu_1979
Somalia_1979	Tanzania_1961	Vanuatu_1980
Somalia_2012	Tanzania_1962	Venezuela_1830
South Sudan_2011	Tanzania_1977	Venezuela_1858
South_Africa_1961	Tanzania_1985	Venezuela_1864
South_Africa_1983	Thailand_1932	Venezuela_1874
South_Africa_1993	Thailand_1949	Venezuela_1881
South_Africa_1996	Thailand_1968	Venezuela_1891
South_Korea_1948	Thailand_1974	Venezuela_1893
South_Korea_1987	Thailand_1976	Venezuela_1904
Spain_1808	Thailand_1978	Venezuela_1909
Spain_1812	Thailand_1997	Venezuela_1922
Spain_1834	Thailand_2007	Venezuela_1925
Spain_1837	Thailand_2014	Venezuela_1936
Spain_1845	Tibet_1991	Venezuela_1947
Spain_1869	Togo_1979	Venezuela_1953
Spain_1876	Togo_1992	Venezuela_1961
Spain_1931	Trinidad_and_Tobago_1962	Venezuela_1999
Spain_1967	Trinidad_and_Tobago_1976	Vietnam_1960
Spain_1978	Tunisia_1959	Vietnam_1980
Sri_Lanka_1931	Tunisia_2014	Vietnam_1992
Sri_Lanka_1946	Turkey_1876	Yemen_1970
Sri_Lanka_1972	Turkey_1921	Yemen_1991
Sri_Lanka_1978	Turkey_1924	Yugoslavia_1921
St Kitts And Nevis_1983	Turkey_1945	Yugoslavia_1931
Sudan_1973	Turkey_1961	Yugoslavia_1946
Sudan_1998	Turkey_1982	Yugoslavia_1953
Sudan_2005	Turkmenistan_1992	Yugoslavia_1963
Surinam_1987	Turkmenistan_2008	Yugoslavia_1974
Swaziland_1968	Tuvalu_1978	Yugoslavia_1992
Swaziland_2005	Tuvalu_1986	Yugoslavia_2003
Sweden_1809	Uganda_1967	Yugoslavia_2006
Sweden_1974	Uganda_1995	Zambia_1964
Switzerland_1848	Ukraine_1978	Zambia_1973
Switzerland_1874	Ukraine_1996	Zambia_1991
Switzerland_1999	United_Arab_Emirates_1971	Zimbabwe Rhodesia_2013
Syria_1930	United_states_1789	Zimbabwe_1965
Syria_1950	Uruguay_1830	Zimbabwe_1969
Syria_1953	Uruguay_1918	Zimbabwe_1979

Table 2e. Constitution list part 5.