Advance Access publication December 23, 2010

Political Analysis (2011) 19:103–117
doi:10.1093/pan/mpq038

Stochastic Process Methods with an Application to Budgetary Data

Christian Breunig
Department of Political Science, University of Toronto, Sidney Smith Hall, Room 3018, 100 St. George Street, Toronto, Ontario M5S 3G3 Canada
e-mail: c.breunig@utoronto.ca (corresponding author)

Bryan D. Jones
Department of Government, The University of Texas at Austin, 1 University Station A1800, Austin, TX 78712-0119
e-mail: bdjones@austin.utexas.edu

Political scientists have increasingly focused on causal processes that operate not solely on mean differences but on other stochastic characteristics of the distribution of a dependent variable. This paper surveys important statistical tools used to assess data in situations where the entire distribution of values is of interest. We first outline three broad conditions under which stochastic process methods are applicable and show that these conditions cover many domains of social inquiry. We discuss a variety of visual and analytical techniques, including distributional analysis, direct parameter estimates of probability density functions, and quantile regression. We illustrate the utility of these statistical tools with an application to budgetary data because strong theoretical expectations at the micro- and macrolevel exist about the distributional characteristics for such data. The expository analysis concentrates on three budget series (total, domestic, and defense outlays) of the U.S. government for 1800–2004.

Standard statistical practice in political science involves examining typical or mean responses in a dependent variable to changes in independent variables and treats scatter around those predicted means as aggravations—chance aggravations, one generally hopes. Yet often theoretical reasoning or practical considerations lead directly not to mean differences among groups or across time but to the higher moments of distributions (i.e., variance, skew, and kurtosis) or to the particular shape of the full distribution of observations. For example, recent studies of inequality and the rise of the super-rich (e.g., Piketty and Saz 2003; Hacker and Pierson 2010) focus not on mean changes in income over time but on changes in the dispersion of income and wealth in America. Students of inequality rely on a variety of measures, such as the 90/10 percentile ratio and the Gini coefficient, in order to examine the widening gap in incomes since the 1980s. Similarly, quantitative works in comparative politics and international relations regularly employ log transformations when constructing variables to adjust for skewness in such as population size, number of battle deaths, and gross domestic product (GDP) per capita.

Further, a critical shift in thinking about distributional components of variables is occurring. Much statistical modeling requires distributional assumptions—usually Normality—for tests of statistical significance to be valid. As a consequence, a standard strategy is to perform transformations until the empirical distribution approximates the Gaussian. The researcher treats deviations from Normality as aggravations that require adjustments, not as quantities worthy of study. But what if the interest of the investigator centers not on mean differences in which deviations from Normality interfere with

Authors' Note: The authors would like to thank Michael D. Ward, John Ahlquist, and Samuel Workman, our reviewers, and the editors of *Political Analysis* for helpful comments. Supplementary materials for this article as well as R and Stata code for implementing the presented methods are available on the authors' Web site.

© The Author 2010. Published by Oxford University Press on behalf of the Society for Political Methodology.
All rights reserved. For Permissions, please email: journals.permissions@oup.com

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

deploying standard statistical tests but directly in the deviations themselves? What if the theories guiding empirical inquiry direct us not to the mean but to the higher moments of the distribution? Political scientists increasingly employ statistical tools, such as heteroskedastic probit, gamma and beta regressions (e.g., Alvarez and Brehm 1995; Paolino 2001), which allow modeling the variance term; this additional parameter can be used to generate important empirical insights. Bayesian methods and multilevel models are employed in part to model variance and other higher moments as dependent on exogenous variables (see *Political Analysis* 13(4)). Moreover, stochastic process approaches can be invaluable in a variety of practical and theoretical situations in which the study of outliers and abnormal heterogeneity is the point—including in detecting the possibility of voting fraud in elections from a study of aggregate returns (e.g., Alvarez 2009).

This article follows this burgeoning interest in higher moments by reviewing and illustrating how stochastic process methods have been employed in public policy and public finance and explicating their larger relevance to political science research. Specifically, we survey important statistical tools that are available for assessing data in situations where the entire distribution of values on a variable is of interest. We review in some detail recent progress in the study of public budgets, showing how theoretical developments imply the necessity of studying distributions rather than just measures of central tendency. We assess U.S. budgetary data using a variety of visual and analytical techniques, including distribution functions, histograms, smooth density estimates, nonparametric regression, direct parameter estimates of probability density functions (p.d.f.), and relative distributions.[1]

We focus on government budgets in this exercise for several reasons. Government budgets offer political scientists the longest historical trace of governmental activity available. When used properly, they offer quantitative evidence of past decision-making activities and its consequences for the allocation of resources across governmental objectives. Developments in the field have centered on time series analyses as well stochastic process methods, but it has been the distributional studies that have provided the empirical underpinnings for major theoretical progress.

Although our empirical exploration concentrates on public budgeting, our aims are general. Stochastic process methods are appropriate in many instances in political science, yet they are not utilized as frequently as applicable. As a consequence, it is important to delineate circumstances when stochastic process methods (or distributional methods) can be applied to political inquiry and what potential payoffs are derived from them. Because methodologies relying on stochastic processes examine the full distribution of outcomes, extreme observations are not simply relegated as outliers. Instead, they serve as keys to the mapping of theoretical distributions on the distribution of data. In many cases, these extreme cases supply important insights for theoretical developments.

Early applications of this approach include the study of income and inequality (Engel 1857; Pareto 1897; Lorenz 1905). Echoing Pareto's work, both the natural and social sciences have rejuvenated in recent years this method of inquiry with the application of power laws to such diverse phenomena as earthquakes, financial markets, and the industrial revolution (Mandelbrot 1997; Sornette 2003; Mandelbrot and Hudson 2004; Beinhocker 2006). Indeed, suggesting that extreme values are outliers that should be ignored when the Gaussian distribution fits otherwise fairly well is akin to ignoring very big earthquakes because they are atypical events.

Although in many cases stochastic process approaches are most relevant for examining the empirical implications of theories, they are also extremely useful in exploratory analyses. They can offer insights when exact modeling of social phenomena and clear empirical model specifications are not possible. There are at least three conditions under which this might be the case. First, specification might be impossible due to complexity and uncertainty of the process. This condition is typical for the social sciences, where open systems with many subsystems often interacting nonlinearly and with feedback are common. Important examples for economics include stock markets (Sornette 2003) and growth processes, such as firms, universities, cities, and countries (Gabaix 1999; Podobnik et al. 2009); political science examples include forms of cooperation and conflict (Richardson 1948; Cederman 2003; Biggs 2005). Generally, stochastic process methods are relevant analytical tools whenever we move from studying equilibrium systems to examining complex systems.

[1]All analysis is conducted in R (R Development Core Team 2008).

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

Second, stochastic process methods allow researchers to establish broad empirical generalizations in situations where the scientific process has not produced extensive theoretical knowledge. In these cases, stochastic process methods might aid researchers in crafting some initial typologies or developing rudimentary causal mechanisms. An exemplary work in this regard is the model of network structures by Barabási and Albert (1999). Barabási and Albert showed that two mechanisms—preferential attachment and population growth—generate scale-free networks that follow a certain type of distribution known as power law. Their initial insights rejuvenated the interest in network analysis across the natural and social sciences. Stochastic process methods can serve as a precondition to sound model building and testing.

Third, stochastic process methods provide a fruitful alternative to regression style analysis when researchers are unable to identify appropriate measures of key explanatory concepts. In these situations, distributional methods can establish broad empirical generalizations that can serve as a starting point for modeling and testing. The classical examples of this approach come from the physical sciences, including climate disruptions and earthquakes, in which empirical laws about the distribution of phenomena preceded modeling. Generalizations based on stochastic process methods thus deliver a starting point for modeling social phenomena that do not require model building based on observational and quasi-experimental studies whose aim it is to isolate individual effects.

Stochastic process methods have the potential of overcoming some of the limitations that political scientists frequently encounter in many observational studies. Stochastic process models try to ascertain what theoretical probability distribution could have accounted for an observed frequency distribution of outcomes. Once the researcher has a theoretical frequency distribution that describes the data, he or she is able to use known properties of the theoretical distribution to explore more precise hypotheses about the empirical phenomena under study. In the following sections, we show how in one major area of study in political science, that is, public budgeting, stochastic process methods were the keys to theoretical progress.

1 Theoretical Background on Budgeting and Implications for Stochastic Process Methods

For many years, the major theoretical tool for understanding government budgeting was the standard incremental model, in which only marginal adjustments are made to existing programs. The incremental model was essentially a behavioral model of choice, but it had implications for budgetary outcomes. If behavioral incrementalism held, then outcome incrementalism should similarly hold. Many scholars were highly dissatisfied with the approach, and several papers critiqued the approach without displacing it. Then in a seminal paper, Padgett (1980) related certain internal behavioral choice mechanisms across budget-alternatives to distinct distributional forms of resulting policy changes. Critically, he showed that the incremental model implied a Gaussian distribution for year-to-year first differences (or Student's t if policy changes vary across time and programs), allowing an easy and conclusive rejection of the incremental model with his data. Investigators have found no budgetary series that displays Gaussian distributions for first differences. Padgett offered a plausible substitute for incrementalism: a serial judgment model. The decision maker prioritizes possible choices in a series of steps (sieves) in which a decision maker compares the budget proposal with the set of constraints, one-by-one rather than comprehensively. The model implies a double exponential distribution of budget changes for a single program over time and a double Paretian for the mixed model aggregated across time and programs.

Padgett's model, like incrementalism, works from the inside out—it connects internal decision making to budgetary outcomes. Recent studies in public policy have worked from the outside in—that is, they ask what the implications of changes in information in the policymaking environment make for decision making within policymaking institutions. The theory of punctuated equilibrium claims that policy change occurs in an episodic and disjointed fashion (Baumgartner and Jones 1993). Decision makers react to changes in information through a combination of inattentiveness and rapid reaction. The latter results in a reprioritization of policy objectives, which in turn lead to budgetary punctuations. Subsequently, it was recognized that this pattern of policy dynamics implies that frequency distributions of outputs should display fat tails (indicative of internal reprioritizations and external policy punctuations), sharp central peaks (indicative of internal inattentiveness and external temporal stability), and ''weak shoulders'' (indicative of a relative lack of moderate change) (Jones et al. 2003). Frequency distributions of budget changes should be leptokurtic (Jones and Baumgartner 2005).

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

Ensuing research identified variation in the institutional structures of policy making (''friction'') as a source of the variation in the distribution of policy outcomes. Friction is simply the resistance to change built into policy-making institutions—such as the Senate filibuster in the United States. Since it is easier to raise an issue (put it on the agenda) than to solve it, friction increases as an issue proceeds through the lawmaking process. Jones, Sulkin, and Larsen (2003) hypothesized that these changes in friction or resistance would lead to more leptokurtosis in output distribution—which was confirmed. In addition, the basic form of budget change distributions holds in a variety of nations and levels of government (Jones and Baumgartner 2005; Mortensen 2005; Breunig 2006; Breunig and Koski 2006; Baumgartner, Foucault, and François 2006). The basic distributional form seems to be double Paretian, whether single budgetary series or amalgams of programs are studied, leading to a claim of a general empirical law of budget distributions (Jones et al. 2009).

Cumulatively, Padgett's decision-making model and Baumgartner and Jones's punctuated equilibrium model imply certain kinds of distributions of budget data and these theoretical propositions can be directly compared to empirical frequency distributions. In particular, studying distributions of budget change data allows the decisive rejection of two important models of budget behavior: the standard incremental model and an updating model based on a weighted average of inputs data. It also points to a behavioral decision-making model based on bounded rationality and the serial processing of alternatives, and to a public policy model based on institutional and cognitive factors. In the following, we show that the serial judgment and friction models are compatible with empirical observations and among a potential class of acceptable models of budgeting. It should be clear that the theoretical insights by Padgett as well as Baumgartner and Jones are not contingent on budget data but based on broader decision-making and public policy models. Nevertheless, stochastic process approaches have been essential to the development of budget studies.

2 Budgetary Data

Because stochastic process approaches are central in budgetary studies, it is useful to detail these methods in a manner that may find receptivity in other areas. In this paper, we rely on three data series in order to explore our ideas: real annual outlays for the U.S. government, for defense and domestic programs, as well as the total outlays for 1800–2004.[2] We compute the growth rates (i.e. the change in the log of real outlays) for the series. The transformation essentially represents percentage changes in spending and takes care of the changing size of the budget over time.

The broadest prediction from the new budgetary theory is that a frequency distribution of change data will be leptokurtic. In order to determine the degree of leptokurtosis, the actual budget data are compared to a hypothetical Gaussian distribution with the same sample mean and variance. Figure 1 is a frequency distribution of the growth rates for total U.S. outlays between 1800 and 2004. The growth rates are based on logarithmic returns, that is, $gr_t = \ln(x_t/x_{t-1})$. Its shape is the classic frequency distribution for budget data and is quite obviously far from the Gaussian distribution predicted by incremental budget theory. Our problem here is how to assess and characterize this distribution so that we can compare it to other distributions from other political entities and pursue an explanation of its shape. The next section provides a methodology for this task.

3 A Moment on Moments

Most quantitative political studies are based on analyses of means; even ''explaining the variance'' really means ''accounting for the variability in a variable with means.'' In many fields of endeavor, including the study of public budgets in political science, it is necessary to move beyond analyses of means and even beyond studying the first two moments of a distribution.

Moments characterize certain aspects of a probability distribution through a series of quantities. The second central moment (around the mean) of a probability distribution is the variance, and the third is skew. The fourth moment about the mean assesses how flat or peaked a probability distribution is. Kurtosis, a standardized measure of the fourth moment, provides a scale-free summary measure of the shape of a distribution (DeCarlo 1997). It can therefore be used as an appropriate tool for measuring

[2]U.S. Government Outlays are from *Historical Statistics of the United States*, compiled by the U.S. Census Bureau, updated from the Office of Management and Budget Web site, Historical Statistics, Table 1.1. The Consumer Price Index was used to adjust for inflation due to the absence of GDP deflators for the early part of the series, with June 1984 = 100.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

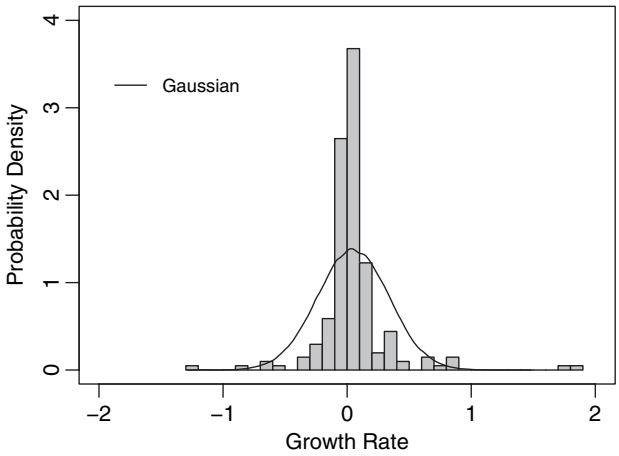


Fig. 1 Budget change frequency distribution. The histogram is based on the annual growth rates for total U.S. outlays, 1800–2004.

punctuations. Budget (and other policy) outcomes that are punctuated are described as leptokurtic (Jones, Sulkin, and Larsen 2003).

Unfortunately, kurtosis is somewhat problematic in practice (Groneveld 1998). Its sensitivity to extreme values in estimation (its quadratic form tends to overemphasize these) makes it a rickety empirical measure. Kurtosis measures based on L-moments are an appropriate alternative because they are less sensitive to extreme values and reliably computed for a relatively small number of cases (Hosking 1998). The kurtosis score can be computed as the fourth L-moments of a distribution (Hosking 1990). The resulting L-moment score ranges from zero to one, where an increasing number identifies a higher level of kurtosis. As a reference point, the standard Gaussian distribution has an approximate L-kurtosis score of 0.123.

Formally, L-moments are the expected values of linear combinations of order statistics $X_{i:n}$multiplied by scalar constants. If $F(x)$ is a distribution function of a random variable X and $X_{1:n} \leqslant X_{2:n} \leqslant \ldots \leqslant X_{n:n}$ are the order statistics associated with the distribution F (or the ordered values of a single sample of size n), then L-moments $L_r(F)$, $r = 1, 2, \ldots$ are defined as $L_r(F) = \frac{1}{r}\sum_{j=0}^{r-1}(-1)^j \binom{r-1}{j} E(X_{r-j:r})$. The first L-moment $L_1(F)$ is a measure of location. $L_2(F)$ is a scale measure for dispersion. The fourth moment (L-kurtosis) is obtained by normalizing $L_4(F)$ by $L_2(F)$. Thus, the L-moment ratio $\tau_4 = \frac{L_4(F)}{L_2(F)}$ measures the kurtosis of a distribution. Simply put, we can easily find summary statistics based on linear combinations of order statistics and thereby alleviate the imprecise and sensitive nature of the common moments (which suffer from their definition with powers).

Table 1 provides the most common descriptive statistics for the three U.S. budget series. Collectively, they provide a basic overview of the examined distributions. It is worth noting how little the first two moments really help in understanding the full distribution of outcomes. For all three series, one would guess (based on the mean and variance) that budgets incrementally increase over time. Specifically, the central moments suggest that these budgets grow between 2 and 5% annually. An inspection of the ordered statistics such as interquartile range (IQR) and L-kurtosis provide some initial insight in the full distribution of outcomes. IQR suggests that the middle 50% of the distributions is between 11 (domestic) and 18% (defense) wide. In other words, budgets do not slowly but steadily increase; instead their fortunes fluctuate considerably over time. The L-kurtosis scores of .40 (domestic) and above indicate that each distribution is highly leptokurtic. Simply put, the L-kurtosis suggests that budgets dramatically expand or shrink in some years (in fact, the minimum and maximum values suggests that they come close to total eradication but also more than double in size). The higher order moments are therefore necessary for understanding the complete decision-making process of budgeting. In other words, in order to develop a comprehensive model of budgeting, one needs to know why budgets sometimes more than double in size.

Table 1 also implies that political scientists should, at least in some areas of inquiry, be interested in applications that move beyond estimating some central characteristics (e.g., average vote return, average

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

Table 1 Descriptive statistics for growth rates in total, domestic, and defense outlays 1800–2004

	Outlays	*Domestic*	*Defense*
Mean	0.05	0.05	0.04
Median	0.02	0.03	0.02
Variance	0.08	0.05	0.15
IQR	0.12	0.11	0.18
Skewness	1.92	1.85	1.46
L-kurtosis	0.46	0.40	0.45
Minimum	−1.22	−0.77	−1.69
Maximum	1.85	1.43	2.41
N	204	204	204

demonstration size, mean inflation rate, average number of soldiers killed) based on a random sample taken from a population under study. Instead, one could be interested in examining the maximum or the minimum as key parameters for many political phenomena.[3] Some examples might include why some elections become landslides, why some demonstrations become million-man marches, why hyperinflation and world wars occur. In short, it is a promising endeavor to generate hypotheses that explain massive transformation or even distributional characteristics of political phenomena (see Cederman 2003 as an example in international relations).

4 Beyond Moments

We noted above that Padgett (1980) has derived full probability distributions for various budget change models. In the case of the incremental and implicit index models, these distributions are Gaussian. Thus, all we need to do for rejecting the incremental model is to demonstrate systematic deviations from the Normal distribution in our sample. But for Padgett's serial judgment model, it is necessary to compare the observed distribution with Paretian and exponential distributions. As a consequence, we cannot simply rely on the simple summary measure of kurtosis. Instead, all subsequent analyses rely on the usage of the cumulative density function of the data.

Throughout the paper, it is assumed that budget change is a continuous random variable X that is independent and identically distributed from the distribution F. Figure 2 displays the growth rates in the total outlays as well as a hypothetical Gaussian distribution for the cumulative density function (c.d.f.). The c.d.f. gives the probability that a randomly chosen value is less than or equal to x; it is the integral of the probability density function (p.d.f.). In short,

$$F(x) = \Pr(X \leqslant x) \equiv \int_{x_{\min}}^{x} \Pr(\xi) d\xi.$$

Figure 2 indicates that the budget data have weak shoulders. Note that the dotted line, the cumulative Gaussian, is higher than the empirical budget function below zero, but lower above zero, indicating that the shoulders of the empirical budget function Gaussian are ''more slender'' than those of the empirical distribution. In substantive terms, this means that totally outlays have less moderate changes than we would normally expect. Second, the graph hints at the existence of ''fat'' tails because the empirical budget curve has more area under it and hence more cases at both extremes. Third, the steep nature of the empirical slope around the central tendency also suggests that a high amount of stability is present at the center of the distribution.

The quantile function, also depicted in Fig. 2, confirms this. In general, the quantile $Q(p)$ can be defined as the value of y below which a proportion of p of the value falls. The weakness of the shoulders is illustrated by the fact that the difference in the growth rates and a Gaussian distribution at the 75th percentile

[3] By extension, this would imply that one moves outside of equilibrium theories of politics.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

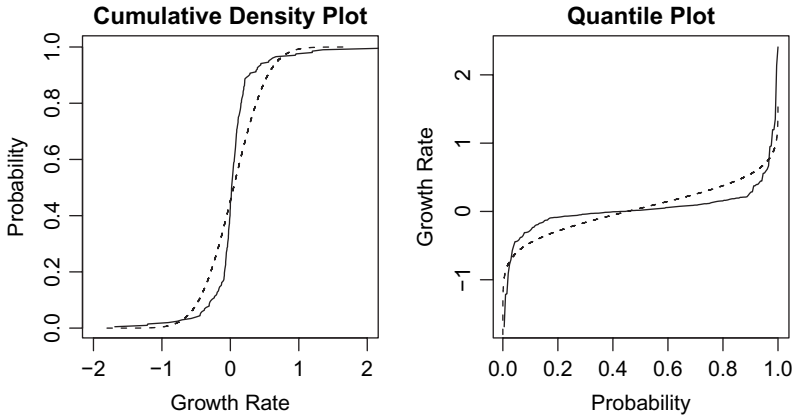


Fig. 2 Cumulative density function and quantile function for growth rate in total outlays, 1800–2004. The solid lines are total outlays and the dashed line the Gaussian distribution.

is about 15% points. In other words, at the 75th percentile, the total growth rate in outlays is roughly 11% points below the expectation derived from the Normal curve.

5 Constructing Histograms

The starting point of many studies of budget changes is a comparison of these changes Gaussian distribution using a histogram. The advantages of the histograms are its ease of interpretability and its convenient construction in most statistical packages. However, John and Margetts (2003, 423) suggest that the bin width, or size of the category used for accumulating cases, is more or less arbitrarily chosen and that the choice of bin width has considerable impact on evidence regarding leptokurtosis and the estimation of scaling parameter for Paretian distributions. How we can transform a continuous probability distribution into a step function that provides an appropriate histogram?

Sheather and Jones (1991) and Simonoff (1996, chap. 2) show that a trade-off between bias and variance exist when choosing the bin width and that it is possible to minimize this trade-off. Suppose that $F(y)$ has a finite range from A to B, how should we split up this range into K equisized intervals with width $h = (B - A)/K$? Simonoff (1996) shows that the minimization of the mean interval squared error (MISE) requires balancing bias and variance through the choice of the number of bins K. The minimizer of the asymptotic MISE is $h_0 = \left[\frac{6}{R(f')}\right]^{1/3} m^{-1/3}$, where $R(v) = \int_{-\infty}^{\infty} [v(x)]^2 \mathrm{d}x$ and m the sample size. To operationalize this rule, several candidates for specifying a particular estimate of f have been proposed. For example, a Gaussian distribution with mean and variance of the sample would lead to $h_0 = 3.491 s m^{-1/3}$, where s is the sample SD. Figure 1 displays the growth rate of total outlays using the optimal bandwidth of 8.37% points, which corresponds to about 37 bins for the data set.

6 Simple Comparisons with a Gaussian Distribution

Distributions can also be studied with simple statistical tests. A straightforward approach is to generate goodness-of-fit tests that assess whether the probability distribution underlying our observations reasonably could have come from a specified distribution—for example, the Normal. This is what the Kolmogorov-Smirnov (K-S) test does. The K-S test is based on the largest absolute difference between the observed and the expected cumulative distributions. The K-S test allows for a rejection of the hypothesis that a set of frequencies is Normally distributed.[4] Additionally, we can employ the Shapiro-Wilk (S-W) statistic as a more powerful test for assessing the Normality of a distribution. In contrast to the K-S test, the

[4]The K-S test can be expanded here in order to compare each tail of the three budget series with an Exponential, Pareto, or any other distribution. Using this method, we are able to reject that the six tails are Exponential.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

Table 2 Rejection of Normal distribution (K-S and S-W test for the growth rates in total, domestic, and defense outlays, 1800–2004)

	Outlays	*Domestic*	*Defense*
D	0.36	0.38	0.31
p-value	0.00	0.00	0.00
W	0.71	0.79	0.74
p-value	0.00	0.00	0.00

SW test does not require specifying the mean or variance of the hypothesized normal distribution in advance. The S-W test generates a *W* statistic, with small *W*-value indicating non-Normality.

Table 2 summarizes the test results for our budget series. The test statistics and the accompanying *p*-values for both the K-S test and the S-W test suggest that we can reject with a high confidence ($p > .001$) the hypothesis that the growth rates for all three outlay series are distributed Normally. It is important to stress that these goodness-of-fit tests only allow us to reject that the observed data are reasonably close to a specific distribution.

7 Comparing Distributions

Recent analyses of natural and social phenomena with a stochastic process view (Mantegna and Stanley 2000; Sornette 2003, 2006) are based on distribution fitting ideas and employ quantile-quantile plots (qq-plots), relative distributional methods, and direct parameter estimates using log-log and semi-log plots; qq-plots are a graphical technique used to determine whether differences in two distributions exist. A qq-plot graphs the quantiles of the first distribution against the second. If the two distributions were similar, its observations would fall on a 45-degree reference line. Greater divergence from that line suggests distinct distributions. A nice feature of qq-plots is that shifts in location or scale, and changes in symmetry are easily detectable. As indicated in Fig. 3, compared to a Normal distribution, U.S. budget outlays display more extreme features. Observed low growth rates (i.e., reductions) are well below the reference line for a Normal distribution and likewise, high growth rates are well above the reference line. These plots are quite protean, since one may calculate quantiles on any p.d.f. and compare the results with the observed frequency distribution.

A more fine-grained technique is the relative distribution. Relative distributional methods (Handcock and Morris 1999) combine previously considered methods, require only weak assumptions about the distributions, and are invariant to measurement scale. The analyst forms a new distribution out of the ratio of the empirical distribution to the reference distribution. Figure 4 shows the relative distribution of Gaussian distribution (the reference) to the three outlay series (the comparisons) using 100 bins. The interpretation of Fig. 4 is as follows: the relative density is the ratio of the frequency of the comparison distribution (budget growth rate) to the frequency of the reference distribution (Gaussian) at the *r*th quantile of the reference distribution level. The density estimates in Fig. 4 provide an intuitive assessment of the two distributions. If both distributions would be distributed as Gaussian, then the relative distributions would be uniform. These graphs make clear that the biggest difference between the two distributions is at the center of the distribution: at least more than twice as many outlays changes as expected by the Gaussian distribution are in the fourth decile. In substantive terms, this feature indicates that budget changes are overwhelmingly incremental. The second and third deciles have at times about half as many outlay changes as expected by the Gaussian distribution. Figure 4 shows that in the top (and bottom) one percentile empirical budget changes occurred more often than the normality assumption would expect.

Both qq-plots and the relative distribution method are quite general, in that the analyst can compare the empirical distribution at hand against any comparative theoretical distribution for which density calculations are available. The analyst can make direct comparisons with the Paretian and exponential p.d.f.s, both of which are implicated in budgetary theory. However, as we discuss next, a better method is available.

8 Direct Parameter Estimates of the Probability Distribution

P.d.f.s relate a range of the values of a variable of interest to the probability associated with that range. In producing a frequency distribution, we simply count the number of cases in a bin or obtain its

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

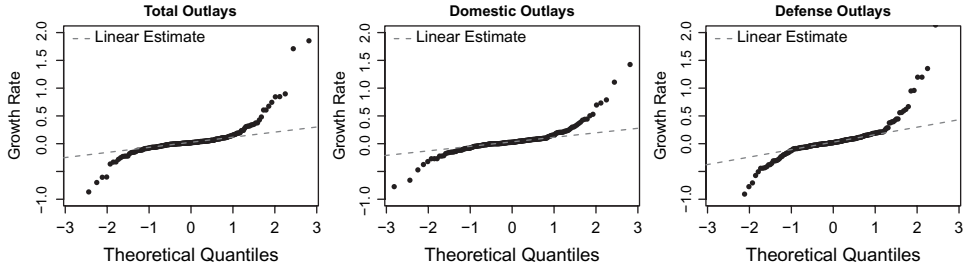


Fig. 3 Normal q-q plots for the growth rates in the three outlays series. The points represent the relationship between a theoretical Gaussian distribution and the actual growth rates at any given quantile. The dashed red line represents the normal quantile-quantile line that passes through the first and third quartiles.

probability density. The relative frequencies in a bin estimate the probability of occurrence associated with the bin.

For a theoretical continuous p.d.f., we define $\Pr(a<x<b) = \int_a^b f(x)\mathrm{d}x$. In some very important cases, we may transform the formula for a particular class of probability density distributions such that the probability is a linear function of the variable of interest. This allows for simple examination of a scatterplot of the empirical data and regression fits for judging how closely the empirical data follows the hypothesized probability law. The two most important p.d.f.s that can easily be transformed are the exponential and Paretian (or power) p.d.f.s. An exponential p.d.f. may be written as: $\Pr(x) = \alpha e^{\beta x}$ and is linear in logarithms. That is, taking logarithms of both sides transforms the exponential into a linear function yields $\ln(\Pr(x)) = \ln(\alpha) + \beta x$. Similarly, for the Paretian p.d.f.—$\Pr(x) = \alpha x^{\beta}$—is log-transformed into $\ln(\Pr(x)) = \ln(\alpha) + \beta \ln(x)$.

In practice, one finds considerable variability in the tails of the distribution, so that it is more convenient to rely on the cumulative relative frequencies of the empirical distribution rather than the relative frequencies themselves (Mandelbrot 1963). The cumulative density function of the power function is also linear in its logarithms.[5] These cumulative densities are then employed in order to assess visually and statistically whether the observed data might fall into these two important classes of distributions. If we plot the logarithm of the cumulative relative frequencies of an empirical distribution against the values of the bins or category midpoints, and the data fall along a straight line, then we may infer that the data approximates an exponential p.d.f.. If we plot the logarithm of the cumulative relative frequencies of the distribution against the logarithms of the category midpoints, and the data fall along a straight line, then we may hypothesize reasonably that the data approximates a Paretian p.d.f..

A major advantage of this approach is that one can obtain direct parameter estimates for the distributions by running the regression models associated with the graphs. Using this method requires considerable care because it is easy to confuse similar distributions, especially if too much attention is directed at the tails of the plotted distribution (Clauset, Shalizi, and Newman 2007). If one is examining a log-log plot, it should be common practice to plot the semi-log plot as well. Many distributions fall in between the two and can better be described as a ''stretched exponential'' (Laherrere and Sornette 1998). Log-normals may also be confused for exponentials or power functions. In any case, the best approach is to examine several competing models (certainly the power, exponential, and log-normal for leptokurtic frequency distributions), comparing goodness-of-fit measures among them.

What do the direct parameter estimates tell us about budgets? Figure 5 displays the log-log and semi-log plots for our three U.S. budget outlay series. In each of these plots, the negative slope (i.e., the values for budgetary declines) is reversed so that they may be easily compared with the positive slope. For the semi-log plots, we find that the direct parameter estimates ''misjudge'' the peaks as well as the tails of the distributions (the dots in each plot are above the estimated regression line). The ''moderate'' goodness-of-fit of these lines is indicated by R^2 ranging between 0.8 and 0.9. In comparison, the log-log plots nicely capture the relationship. The parameter estimates approximate straight lines with only slight deviations,

[5]Indeed, Clauset, Shalizi, and Newman (2007) show that OLS parameter estimates based on the p.d.f.s are often biased, whereas parameter estimates of the c.d.f.s are more robust. Alternatively, maximum likelihood estimates of the parameters might be obtained.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

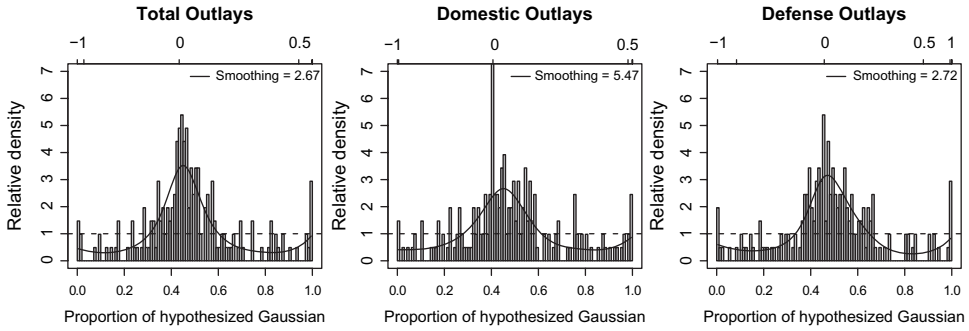


Fig. 4 Relative distribution plot for the three outlays series and a Normal distribution. The hypothesized Gaussian distribution is based on the mean and variance of each of the outlay distributions.

which is also illustrated by R^2 of .97 and higher. In short, budget growth is not Gaussian or exponential and might be distributed in Paretian fashion.

One of the major advantages of direct parameter estimation of a postulated p.d.f. is that it facilitates direct comparison of slopes: this in effect compares the magnitude of the exponents for the Paretian p.d.f.. A shallower slope—one that is smaller in absolute magnitude—indicates that more cases occur in the tail of the distribution. For change distributions, this means that the distribution is more punctuated. For example, the positive slope for changes in real domestic outlays is smaller in absolute magnitude than the negative slope, indicating more large leaps into new spending than is the case for budget cuts. For defense, the pattern reverses itself, with cutbacks being more severe than increases. We may also compare slopes for budgetary changes in different nations, or for different levels of government, asking two questions. Do all budgets follow a power function (the answer so far is ''yes''), and how do nations differ with regard to the exponent of the probability distribution?

9 Quantile Regression

Given the non-Normal nature of the growth rates in outlays, it is possible that different causal factors are involved in incremental growth and punctuations—the causes of extreme changes may be different from those that maintain a steady state. The semi-parametric technique called quantile regression makes it possible to examine whether the same causal processes influence large and small budget cuts, as well as small and large budget increases. Quantile regression, which was introduced by Koenker and Bassett (1978), is an extension of the classical ordinary least squares (OLS) estimation of conditional mean models to estimating models for conditional quantile functions.[6] The formulation of how this *location* model can be delineated as an optimization problem can be outlined as follows (Kroenker 2005).

Generally, we can say that a budget change at the τth quantile is higher than the proportion τ and worse than the proportion $(1 - \tau)$ of the reference group of all budget changes. More formally, any random variable $y \in \mathbb{R}$ may be characterized by its distribution function $F(y) = \Pr(Y \leqslant y)$ for any $0 < \tau < 1$, $Q(\tau) = \inf\{y : F(y) \geqslant \tau\}$ is called the τ quantile of X. The quantiles can be formulated as a solution to an optimization problem. For any $0 < \tau < 1$, describe the piecewise linear *check function*, $\rho_\tau(u) = u(\tau - I(u<0))$, where $I(\ )$ is the indicator function. Minimizing the expectation of $\rho_\tau(Y-\xi)$ with respect to ξ yields the solution, $\hat{\xi}(\tau)$, the smallest of which is $Q(\tau)$ defined above. For a random sample, $\{y_1, \ldots, y_n\}$, the τth sample quantile may be found by solving $\min_{\xi \in \mathbb{R}} \sum_{i=1}^{n} \rho_\tau(y_i - \xi)$. Finally, extending the idea of estimating the conditional mean function, the linear conditional quantile function for $y_i = x_i\beta_\tau + u_{\tau_i}$ is $Q_Y(\tau|X = x) = x_i\beta(\tau)$. This can be estimated by solving $\hat{\beta}(\tau) = \text{argmin}_{\beta \in \mathbb{R}^1} \sum \rho_\tau(y_i - x'\beta)$.

[6]Recent empirical quantile regression literature on labor data include, for example, Buchinsky (1994) and Handcock and Morris (1999).

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

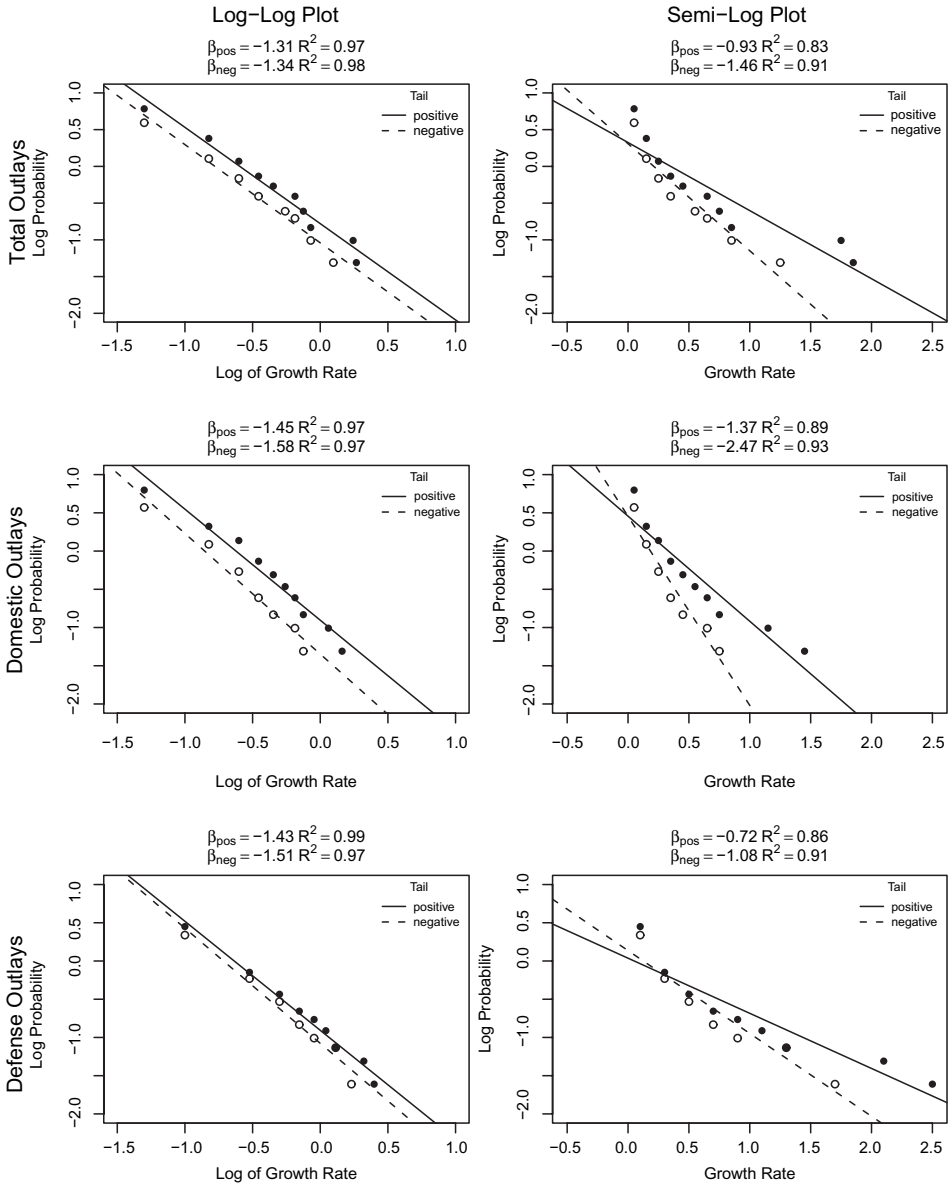


Fig. 5 Log-log (left) and semi-log plots (right) of the growth rates in the three outlay series (top to bottom). The two lines represent the regression estimates for the positive (solid) and negative (dashed) tails of each distribution.

To illustrate, we hypothesize that growth rates in total, domestic, and defense outlays respond to two major forces—the ideological composition of Congress and economic well-being of the nation. We postulate that the more conservative Congress, the less the probability that there will be high growth rates in domestic budgets, but that conservatism is associated with high growth rates in defense expenditures. For the second causal factor, there is a large economic literature (known as Wagner's Law) claiming that economic development is the primary determinant of the size of government. A more complete model might assess other aspects of the external environment, such as external threats and demographic pressures, and internal dynamics, such as the party of the president. But this simple approach using quantile regression offers substantial insights not presently reported in the literature. We operationalize the two covariates the following way. Congressional ideology is the mean of the first dimension DW-NOMINATE score (Poole and Rosenthal 1985). Economic growth is the growth in constant $ GDP (Historical Abstracts of the United States). The time series ranges from 1855 to 2004.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

Table 3 Quantile regression results for large cuts, median changes, and large increases plus OLS results

		Quantile			
Outlays	*Covariate*	*0.05*	*0.5*	*0.95*	*OLS*
Total	Intercept	−0.269*	0.018	0.393*	0.02
		−0.054	−0.016	−0.061	0.03
	House Ideology	−0.965*	−0.092	1.292	−0.026
		−0.356	−0.161	−0.887	0.306
	Econ.Growth	1.358*	0.033	2.269*	0.902*
		−0.606	−0.328	−0.676	0.457
Domestic	Intercept	−0.147*	0.038*	0.338	0.05*
		−0.023	−0.015	−0.112	0.021
	House.Ideology	−0.308	0.003	0.022	0.034
		−0.293	−0.138	−0.972	0.21
	Econ.Growth	−0.337	−0.179	−0.25	0.086
		−0.4	−0.275	−0.808	0.313
Defense	Intercept	−0.476*	0.021	0.454*	−0.008
		−0.069	−0.021	−0.052	0.041
	House.Ideology	−3.327*	0.019	1.846*	−0.158
		−0.731	−0.244	−0.809	0.418
	Econ.Growth	1.244	−0.057	3.253*	1.576*
		−0.913	−0.387	−0.698	0.623

Note. SEs are subset.
"*" Indicates $p < .05$.

We rely on both the regression results for three quantiles (Table 3) and a visual summary of all quantile regression results (Fig. 6) in order to illustrate the regression results. To restate, the regressions are estimates of linear conditional quantile functions just like OLS estimates conditional mean functions. Figure 6 provides a more complete description of how the conditional distribution of the budget changes depends on all the covariates. For each of these coefficients, the point estimates (solid line with dots) can be interpreted as the impact of a one-unit change of the covariate on budget change, holding other covariates fixed. Hence, each of the plots display the quantile τ on the horizontal axis and the vertical scale in percent change of budgets indicates the covariate effect.[7] The gray shaded area illustrates a 90% pointwise confidence band for the quantile regression.

The regression results suggest that economic and political forces have distinct estimated effects on cuts, marginal budget adjustments, and increases. The estimation indicates that for the total and the defense growth rates, when Congress becomes more conservative (House Ideology goes up), budget cuts will be more severe and budget increases become even larger. In short, the quantile regression suggests, most clearly for defense outlays, that a push in the conservative direction produces more extreme outcomes. Since the OLS regression is only concerned with the average growth rate, the described associations on the extremes cannot be detected.

Second, the quantile regression suggests that economic growth has a U-shaped relationship on total and defense growth rates. Better economic times alleviate budget cuts and amplify budget increases. All that the OLS regression suggests is that increases in economic growth lead to budget growth, missing the nonlinear effects associated with the extremes. The quantile regression estimates do not detect statistically significant differences of ideology and economic prosperity on growth in domestic budget outlays. Moreover, the existing literature has not highlighted this finding; as a consequence, it deserves further study. One suggestion, however, is that during times of moderate economic growth, budgets change

[7] Jones and Breunig (2007) extensively discuss memory effects in budget time series data and show that it is therefore tricky to identity a plausible stochastic process. Given this, the OLS estimates presented here are for illustrative reasons, and therefore, we assume stationarity.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

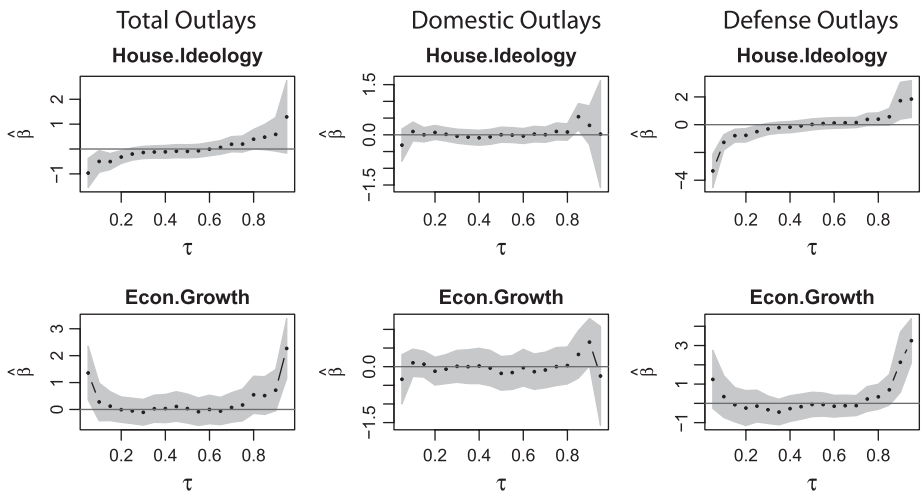


Fig. 6 Quantile regression of the outlays series. Each column interprets the estimation of each series labeled at the top. The *x* axis of each plot represents the quantiles of the budget changes (ranging from large cuts to large increases) and the *y* axis displays the estimated coefficients of each variable. The gray band is the 90% confidence band. The black lines are the point estimates.

proportionally to the economic growth rate, linked by simple decision rules as Wildavsky suggested. In times of economic stress, budgets cannot be cut as fast as the decline, and in modern times Keynesian theory dictates countercyclical spending. In periods of high growth, budget competition actually propels budget growth beyond the existing growth rate as program managers expect continued similar growth, which of course comes back to earth at some point. But whatever the explanation, the complex contingencies between economics and budgets defy simple models.

10 Conclusions

There are many situations in which the mean differences assessed by regression and related analyses do not model correctly the phenomena under study. As a consequence, political scientists are paying increasing attention to understanding and modeling higher moments of political phenomena. In this article, we have take a step back and considered situations in which (1) model specification might be impossible due to complexity and uncertainty of the process, (2) limited theoretical knowledge of a phenomena exists, or (3) indentification of appropriate measures of key explanatory concepts is impossible. We show that stochastic process methods have the potential to overcome some of these problems of political science research. In fact, our application of these analytical tools to budgetary data illustrates how distributional tools provide new insights in a well-established research field.

In general, we believe that social scientists can take advantage of the presented methods in situations when they are interested in theorizing about the full distribution of observed outcomes in order both to develop broad empirical generalizations for a class of social phenomena and to delineate some general causal mechanisms. Stochastic process methods therefore should be particularly rewarding for studies of dynamic (nonequilibrium) systems, which probably characterize many fields of political science research. We have discussed several different tools because it is challenging to distinguish carefully among different distributional forms; therefore, we urge researchers to employ a variety of methods in order to gain confidence in their analysis. Moreover, stochastic process methods can provide a starting point for more refined theorizing and modeling. In fact, for our application to budgets, we show how quantile regression can be utilized in order to identify distinct causal process across the whole outcome distribution. The lesson for the quantile regression is simple: once we move beyond conditioning our analysis on the mean, a richer and more complete set of empirical and theoretical insights emerges.

Theoretical developments in the study of public budgets imply particular distributions of year-to-year changes in budgetary allocations, but they do not always imply any particular time path of

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

budgets. As a consequence, it is often appropriate for students of budgets, and quantitative policy changes more generally, to rely on stochastic process methods rather than the more typical model estimation techniques relying on changes in means. Classic incremental theory implies a Gaussian distribution of budget changes (and hence a time path that is a random walk). However, the incremental model has been rejected decisively in our empirical investigations. Further analysis of government budgets requires distributional approaches since we have no theory that implies any particular time path of budgets but we do have theories that imply specific distributional patterns. Only by harnessing stochastic process methods based on full distributions of data can the full implications of the emerging theory of public expenditures come clear.

References

Alvarez, R. Michael, and John Brehm. 1995. American ambivalence towards abortion policy: Development of a heteroskedastic probit model of competing values. *American Journal of Political Science* 39:1055–82.

Alvarez, R. Michael. 2009. In *Election fraud: Detecting and deterring electoral manipulation*, eds. Thad E. Hall and Susan D. Hyde. Washington, DC: Brookings Institution Press.

Barabási, Albert-László, and Réka Albert. 1999. Emergence of scaling in random networks. *Science* 286:509–12.

Baumgartner, Frank R., and Bryan D. Jones. 1993. *Agendas and instability in American politics*. Chicago, IL: University of Chicago Press.

Baumgartner, Frank R., Martial Foucault, and Abel François. 2006. Punctuated equilibrium and French budgeting processes. *Journal of European Public Policy* 13:1086–103.

Beinhocker, Eric D. 2006. *Origin of wealth: Evolution, complexity, and the radical remaking of economics*. Cambridge: Harvard Business School Press.

Biggs, Michael. 2005. Strikes as forest fires: Chicago and Paris in the late nineteenth century. *American Journal of Sociology* 110:1684–714.

Breunig, Christian. 2006. The more things change, the more things stay the same: a comparative analysis of budget punctuations. *Journal of European Public Policy* 13(7):1065–81.

Breunig, Christian, and Chris Koski. 2006. Punctuated equilibria and budgets in the American States. *Policy Studies Journal* 34:363–79.

Buchinsky, Moshe. 1994. Changes in the U.S. wage structure 1963–1987: Application of quantile regression. *Econometrica* 62:405–58.

Cederman, Lars-Erik. 2003. Modeling the size of wars: From billiard balls to sandpiles. *American Political Science Review* 97:135–50.

Clauset, Aaron, Cosma Rohilla Shalizi, and M. E. J. Newman. 2007. Power-law distributions in empirical data. arXiv: 706.1062v1. *Physics.Data.an.* 7:1–26.

DeCarlo, Lawrence T. 1997. On the meaning and use of kurtosis. *Psychological Methods* 2:292–307.

Engel, Ernst. 1857. Die Produktions- und Konsumptionsverhältnisse des Königreichs Sachsen. *Zeitschrift des Statistischen Bureaus des Koniglich Sächsischen Ministeriums des Innern* 8:1–54.

Gabaix, Xavier. 1999. Zipf's law for cities: An explanation. *Quarterly Journal of Economics* 114:739–67.

Groneveld, Richard A. 1998. A class of quintile measures for kurtosis. *American Statistician* 51:325–9.

Hacker, Jacob S., and Paul Pierson. 2010. Winner-take-all politics: Public policy, political organization, and the precipitous rise of top incomes in the United States. *Politics & Society* 38:152–204.

Handcock, Mark and Martina Morris. 1999. *Relative distribution methods in the social sciences*. New York: Springer.

Hosking, Jonathan R. M. 1990. L-moments: Analysis and estimation of distributions using linear combinations of order statistics. *Journal of the Royal Statistical Society B* 52:105–24.

———. 1998. L-moments. In *Encyclopedia of statistical sciences*. Vol. 2, eds. S. Kotz, C. Read, and D. L. Banks, 357–62. Hoboken, NJ: Wiley.

John, Peter, and Helen Margetts. 2003. Policy punctuations in the UK. *Public Administration* 81(3):411–32.

Jones, Bryan D., Frank Baumgartner, Christian Breunig, Christopher Wlezien, Stuart Soroka, Martial Foucault, Able Francois, Christoffer Green-Pederson, Chris Koski, Peter John, Peter Mortensen, Frederic Varone, and Steffan Walgrave. 2009. A general empirical law of public budgets. *American Journal of Political Science* 53:855–73.

Jones, Bryan D., and Frank R. Baumgartner. 2005. *The politics of attention*. Chicago, IL: University of Chicago Press.

Jones, Bryan D., and Christian Breunig. 2007. Noah and Joseph effects in government budgets: Analyzing long-term memory. *Policy Studies Journal* 35:329–48.

Jones, Bryan D., Tracy Sulkin, and Heather Larsen. 2003. Policy punctuations in American political institutions. *American Political Science Review* 97:151–70.

Koenker, Roger. 2005. *Quantile regression*. Cambridge: Cambirdge University Press.

Koenker, Roger, and Gilbert Bassett Jr. 1978. Regression quantiles. *Econometrica* 46:33–50.

Laherrère, J., and D. Sornette. 1998. Stretched exponential distributions in nature and economy: ''fat tails'' with characteristic scales. *The European Physical Journal B* 2(4):525–39.

Lorenz, Max O. 1905. Methods of measuring the concentration of wealth. *Publications of the American Statistical Association* 9:209–19.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press

Mandelbrot, Benoit. 1963. The variation of certain speculative prices. *Journal of Business* 36:294–419.
———. 1997. *Fractals and scaling in finance*. New York: Springer.
Mandelbrot, Benoit, and Richard L. Hudson. 2004. *The (Mis)behavior of markets*. New York: Basic Books.
Mantegna, Rosario N., and H. Eugene Stanley. 2000. *An introduction to econophysics*. Cambridge: Cambridge University Press.
Mortensen, Peter. 2005. Policy punctuations in Danish local budgeting. *Public Administration* 83:931–50.
Padgett, John F. 1980. Bounded rationality in budgetary research. *American Political Science Review* 74:354–72.
Paolino, Philip. 2001. Maximum likelihood estimation of models with beta-distributed dependent variables. *Political Analysis* 9:325–46.
Pareto, Vinferdo. 1897. The new theories of economics. *Journal of Political Economy* 5:485–502.
Piketty, Thomas, and Emmanuel Saz. 2003. Income inequality in the United States, 1913–1998. *Quarterly Journal of Economics* 118:1–39.
Podobnik, B., D. Horvatic, A. M. Petersen, M. Njavro, and H. E. Stanley. 2009. Common scaling behavior in finance and macroeconomics. *The European Physical Journal B* 76(4):487–90.
Poole, Keith T., and Howard Rosenthal. 1985. A spatial model for legislative roll call analysis. *American Journal of Political Science* 29:357–84.
R Development Core Team. 2008. *R: A language and environment for statistical computing*. Vienna, Austria: R Foundation for Statistical Computing. http://www.R-project.org.
Richardson, Lewis F. 1948. Variation of the frequency of fatal quarrels with magnitude. *Journal of the American Statistical Association* 43:523–46.
Sheather, S. J., and M. C. Jones. 1991. A reliable data-based bandwidth selection method for kernel density estimation. *Journal of the Royal Statistical Society series B* 53:683–90.
Simonoff, Jeffrey. 1996. *Smoothing methods in statistics*. New York: Springer.
Sornette, Didier. 2003. *Why stock markets crash*. Princeton, NJ: Princeton University Press.
———. 2006. *Critical phenomena in natural sciences*. Berlin, Germany: Springer.

https://doi.org/10.1093/pan/mpq038 Published online by Cambridge University Press