

Policy design and local tax elections:
News from the search for a popular tax

Isaac William Martin
Heather Harper

University of California – San Diego
Paper prepared for the annual meetings of the ASA

August 2019

Abstract

What kinds of taxation are least objectionable to voters? The question is a classic preoccupation of public officials and social scientists alike. We exploit two advances that make it possible to shed new light on this question. The first is a large, new corpus of digitized documents describing municipal tax policies of heterogeneous design that have been directly subjected to popular referendum in the state of California. The second is a suite of new computational methods for the analysis of large corpora of texts and their associated quantitative metadata. By applying these new methods to those new data, we show that tax policies of different description vary systematically in their popularity at the ballot box. We find that official descriptions of tax policy differ along two interpretable dimensions that are associated with voters' willingness to approve the tax. We interpret these dimensions as *risk pooling* and *community orientation*, and we show that measuring these dimensions can modestly improve our ability to predict the popularity of a tax, relative to a conventional regression specification that omits information about qualitative policy design. We discuss implications of our findings for the research on tax policy design.

Almost 30 years ago, Alice Rivlin, the former director of the Congressional Budget Office, delivered a talk to the American Economics Association on “The Continuing Search for a Popular Tax” (Rivlin 1989). Because the tax debates of the future will take place under conditions of fiscal strain, she argued, these debates “are likely to be less concerned with the equity and efficiency of taxes than with their political acceptability”; and therefore, she predicted, “Economists may find themselves borrowing the tools of sociologists and psychologists to find out how people think about taxes and what kinds are least objectionable” (1989: 113).

Sociologists at the time had little substantive knowledge to offer. In the years since then, and especially in the last decade, sociologists have devoted increasing attention to questions of taxpayer consent (see the reviews in Campbell 1993; Martin, Mehrotra and Prasad 2009; Martin and Prasad 2014). Among the most important findings of the new literature are that many taxes *are* popular, but that there is also substantial variation in how voters respond to taxation (Pearson 2014; Williamson 2017; Martin, López and Olsen 2018). And while sociologists and political scientists have learned a lot about features of social context that can affect tax attitudes—such as employment insecurity (Barnes 2014, Dodson 2017), racial diversity (O’Brien 2017), and partisan polarization (Campbell 2014)—we still know relatively little about what features of a tax policy itself may make it more or less politically acceptable.

In this paper, we address Rivlin’s question: what kinds of tax are least objectionable to voters? We exploit two advances that make it possible to shed new light on this question. The first is a large, new corpus of digitized documents describing tax policies of heterogeneous design that have been directly subjected to popular referendum in the state of California (Martin, López and Olsen 2017). The second is a suite of new computational methods for the analysis of large corpora of texts and their associated quantitative metadata. By applying those new methods to these new data, we show that tax policies of different description vary systematically in their popularity. We proceed in two steps. First, we validate our approach to “text as data” (Grimmer and Stewart 2013) by showing that a simple clustering algorithm applied to textual policy descriptions yields a meaningful classification of tax policies, corresponding to types of taxation that other studies have found to be most politically salient. Second, we employ a supervised learning algorithm to identify dimensions that distinguish among policies in an abstract conceptual space. Our aim is to go beyond a mere catalogue of popular and unpopular types of taxation. By abstracting from particular policy types to dimensional concepts, we hope to contribute to greater theoretical understanding of *what it is* about particular types of policy that may conduce to more or less voter support.

We find that official descriptions of tax policy differ along two interpretable dimensions that are associated with taxpayer consent. The first dimension, which we call *risk pooling*, distinguishes between policies that socialize the risk of social exclusion and policies that individualize the costs of social inclusion. The second dimension, which we call *community orientation*, distinguishes between taxes on residency associated with the town, and taxes on commercial activity associated with the city. We show that measuring these dimensions can modestly improve our ability to predict the popularity of a tax at the ballot box, relative to a conventional regression specification that omits information about qualitative policy design. We conclude with some suggestions for future theoretical development.

Why social scientists do not know much about how to write a popular tax

The idea that some taxes enjoy greater consent than others was a commonplace of classical political economy (see e.g. Mill 2004 [1871]). In recent decades, it has received renewed attention from scholars of the welfare state, who have described a cross-national correlation between the choice of tax instruments and the political sustainability of social spending (see the review in Martin and Prasad 2014). Despite widespread agreement that some taxes may be more politically tolerable than others, there is little consensus on the precise characteristics of taxes that render them popular. A good example is the payroll tax authorized by Title VIII of the Social Security Act of 1935, which finances Old Age Insurance. Social scientists generally agree that this tax is comparatively popular, but there is no consensus about why. Scholars have suggested variously that the Social Security payroll tax is comparatively popular because it is called a “contribution” rather than a tax (Ikenberry and Skocpol 1987: 412); because, as tax on payroll, it symbolically associates the program with the dignity of work (Steensland 2008: 13-14); because it is earmarked for people regarded as deserving of benefits (Campbell 2003: 6); because it can be withheld from paychecks inconspicuously (Williamson 2017: 56-58); or because its incidence is obscured by the nominal splitting of the tax burden between employees and employers (McCaffery 1994: 1878-1880). These hypotheses are not mutually exclusive, but they are different. If we wish to draw any implications from this case for the design of *other* tax policies, it is important to understand precisely which of these dimensions of the payroll tax, perhaps among others, account for its comparative popularity.

Social scientists who examine questions like these often have attempted to isolate the effects of particular policy features through “framing” interventions that draw attention to some of those features, while taking the policy design itself to be fixed. Experimental studies have focused, for example, on whether the same levy will induce different behavior if it is described as a “tax” or as a “payment” (McCaffery and Baron 2006), or whether survey respondents react differently to the same policy when it is described as a “death tax” or an “estate tax” (Sides 2015). Research designs like these rely on a distinction between the effective policy text and the words used to frame it in a particular interaction. Such studies may reveal effects of “crafted talk” *about* tax policy in the course of a political campaign (Jacobs and Shapiro 2000, Graetz and Shapiro 2005), but may have little to say about the effects of “policy crafting” (Martin 2010) or policy design itself.

Other studies have attempted to isolate the relevant dimensions of policy design experimentally, by means of paired comparisons that hold many policy features constant (e.g. McCaffery and Baron 2006; Pampel, Andrighetto, and Steinmo 2019). The best of these studies provide convincing causal inferences, but it can be difficult to abstract theoretical conclusions from them, because they compare few policies at a time, and rely perforce on extremely abridged descriptions of those policies. Haselswerdt and Bartels (2015), for example, have shown experimentally that respondents in a representative sample of U.S. adults preferred public benefits described as “reducing their tax burden” to equivalent benefits described as “a cash payment from the government” (2015: 611). We might wonder whether this effect is best interpreted as a response to taxes in general, personal taxes, or income taxes in particular; to assumptions respondents may have made about implementation or administration of the benefits; or to some other perceived difference between a cash benefit and a tax break. Each of these hypotheses could be tested in turn, but an exhaustive experimental search of the policy design space is unfeasible. To narrow the search, we require an analytical description of the relevant dimensions of policy variation.

Still other, observational studies have attempted to specify the relevant dimensions by comparing how people respond politically to multiple categories of taxation (e.g. Williamson 2017; Martin, López and Olsen 2018). Even the best recent studies compare only a few categories at a time. Campbell (2018), for example, compares survey responses given by American adults evaluating the fairness of eight kinds of taxation, corresponding to a conventional classification in revenue statistics; Martin and Gabay (2018) similarly compare frequencies of reported protest against eight kinds of taxation, using nearly the same categories. These studies draw plausible inferences, but their theoretical usefulness may be limited by their reliance on a relatively crude administrative classification of taxes. The curse of dimensionality in such research designs applies to the number of *kinds of taxes* that are compared; even eight categories of taxation is too few to span the theoretical space of potentially relevant policy characteristics.

Social scientists who have relied on theoretical criteria to identify politically relevant dimensions of policy design at a greater level of generality—such as Pierson’s (1993) two dimensions of “visibility” and “traceability” (see Wilensky 1976, Arnold 1991)—have typically abstracted so far from any observable characteristics of policy that their arguments are difficult to operationalize. It is not obvious how to identify which taxes are “visible” or “traceable” from the policy text alone, for example, and scholars applying these abstractions have often relied on untested measurement assumptions (see, e.g., Martin and Gabay 2018; Roosma et al. 2016; Wilensky 2002). At worst, this practice can lead to circular arguments, as when scholars assert that voters resist the personal income tax because it is “visible,” and then adduce, as evidence of its visibility, studies showing that voters resist it (see, e.g., Hibbs and Madsen 1981; Prasad 2005).

We aim to make headway on Rivlin’s question by taking an alternative, inductive approach to describing kinds of taxation that may be more or less popular. We treat the problem as fundamentally one of data availability. As long as political sociologists limit ourselves to comparisons of a few policies at a time, it will be difficult to isolate the features of those policies that make a difference to their popularity, let alone to know whether the features we have isolated are really those that would make the greatest difference outside of the experimental setting or the survey encounter. Our approach is therefore to measure political responses to a large sample of policies that vary along multiple dimensions.

Our focus, like Rivlin’s, is on the search for characteristics of taxes that are popular in the sense that they can be politically sustained in a democracy. We therefore focus on “popularity” as expressed in the share of affirmative voters at the ballot box in a referendum election. The aggregate vote in a policy referendum may differ from other measures of policy attitudes in the adult public: voters are a non-random subset of the population, whose policy attitudes may be unrepresentative of the public at large, and the yes-or-no votes that voters cast at the end of an election campaign also may differ from the policy attitudes that even the same individuals would express in another context, with different information, or with a different array of policy options before them. For all of these reasons, it would be a mistake to assume that the tax that is most popular with voters on election day is the tax that would be most popular in another deliberative context such as an experiment or a public opinion poll. But observing which tax policies survive referendum elections may be more informative than experiments or public opinion polls about which tax policy characteristics can be politically sustained in a democratic system.

A computational analysis of California municipal tax measures

California provides a unique opportunity to implement this approach. In 1978, an amendment to the state constitution (Proposition 13) required most California municipalities to submit certain kinds of tax increases to the voters for approval; in 1996, another such amendment (Proposition 218) effectively generalized this requirement to every category of local taxation in every municipality in the state. In the intervening years, municipal governments have put hundreds of tax policies to a popular vote. Some categories of taxation are not represented in this corpus of proposed policies because state law forbids their use by local government: most notably, local governments may not levy personal income taxes (Cal. Rev. & Tax. Code §17041.5), estate taxes (Cal. Rev. & Tax. Code §13301), or any ad valorem property taxes in excess of the constitutional maximum rate (Cal. Const., art. XIII D, §3). Nevertheless, the creativity of California’s local officials within these constraints has been considerable, and the taxes proposed to voters by California municipalities since 1996 are quite diverse with respect to various policy features that have attracted attention from political sociologists and behavioral economists. The list includes taxes that were described to voters as “service charges,” “fees,” “assessments,” and, of course, “taxes”; taxes on payroll, retail sales, business receipts, parcels of real estate, and utility consumption, among other things; taxes that are earmarked for purposes as diverse as parks and paramedical services, and taxes that are not earmarked; graduated taxes, flat taxes, and lump-sum taxes; taxes on luxury goods and services, whose incidence may be assumed to be progressive, and taxes on utilities, whose incidence may be assumed to be regressive; taxes that are borne by the legally liable taxpayer, and taxes whose costs can be passed along to customers or employees; and taxes that vary in many other respects. If policy characteristics such as these matter to voters, then we might expect to find variation in voters’ willingness to approve municipal taxes in California.

Our analysis focuses on a subset of 850 municipal tax proposals put to voters in 330 cities and towns from 1996 to 2010. The average tax proposal in this sample was favored by 57.5 percent of the voters who cast ballots on the question; the affirmative vote share ranged from 5.8 percent to 96.1 percent, and the standard deviation was 15.4 percentage points. For each proposal, we collected, digitized, and cleaned the “impartial analysis” provided by mail to voters, which is a summary of no more than 500 words, “showing the effect of the measure on the existing law and the operation of the measure,” required by state statute, and prepared, except in unusual cases, by the city attorney (see Cal. Elec. Code §9280). Our sample includes every proposal to introduce, continue, or increase a tax, for which we could locate the impartial analysis in the documentary record that survives in county or city archives. We excluded proposals to issue municipal bonds, because such measures are subject to different legal rules from tax measures, and may elicit different political behavior.

We cast a broad net for relevant policy design features. In contrast to previous research on California ballot measures, which selected a small number of design features for analysis based on an informal review of the literature (e.g. Martin, López and Olsen 2018), we used machine learning methods to identify *any* aspects of policy design revealed in the semantic features of the impartial analysis text that were associated with the affirmative vote share. We preprocessed the corpus of text for analysis in several steps. First, we converted the impartial analysis documents to plain text format, and cleaned the text files by manually retyping sections that were incorrectly read by optical character recognition software. Next, we converted words into tokens, which were standardized in lower case and stemmed into their root or base form (following Porter 1980). We expanded common abbreviations, such as “UUT” for “Utility

Users’ Tax,” and removed “stop words” including the names of cities and individuals, along with a corpus of common English-language stop words provided with the Natural Language Tool Kit (see Bird, Klein and Loper 2009) that provide little semantic information for distinguishing texts. We also removed any term that appeared in fewer than three documents. After all of these preprocessing steps, our corpus contained 1,548 unique terms.

The next step was to convert each text into a vector representation. We computed the frequency of each term in each document, and weighted these term frequencies by the inverse of the frequency of documents containing the term across the corpus, so that each text could be represented as a vector of term-frequency-inverse-document-frequency (TF-IDF) weights. We computed a smoothed TF-IDF weight $k_{t,d}$ as follows:

$$k_{t,d} = f_{t,d} * (1 + \log [(n_t + 1)/(N + 1)])$$

where $f_{t,d}$ is the number of times term t appears in document d ; n_t is the number of documents containing term t ; and N is the total number of documents in the corpus. After computing these weights, the final step was to normalize the weights $k_{t,d}$ within each document d so that the sum of squared TF-IDF weights within each document equals one. This TF-IDF vectorization process discards all syntactical information and treats each text as a set of unordered words, yielding a so-called bag-of-words representation. It takes into consideration the occurrence of each term within a document and the rarity of the term across the entire corpus; words that occur frequently within a document but infrequently across the entire corpus will have the largest weights. This quantitative summary allows us to represent each impartial analysis as a row vector in a 1,548-dimensional semantic space, and the entire corpus of policy text as an 850-by-1,548 document-term matrix of TF-IDF weights. Although our vector representation of policy text associates a quantitative weight with each individual term, we do not attempt to model the effects of individual words. Our approach is to treat TF-IDF weights as input for algorithms that produce a holistic characterization of the policy. The presence and frequency of particular words can provide us with clues as to whether a proposed policy is a hotel tax or a utility users’ tax, for example, and those different kinds of taxes may provoke different political responses.

Our analysis of these data proceeds in two stages. In the first stage, we assess whether this quantitative representation of policy text can yield an interpretable, qualitative measure of features of public policy that are known to be politically salient. We begin with a simple k -means clustering algorithm that identifies k sets of documents whose TF-IDF representations are closest in semantic space. This common clustering algorithm proceeds by randomly assigning each observation to one of k clusters; computing the centroid of each cluster in the vector space of TF-IDF weights; re-assigning each observation to the cluster whose centroid is nearest; and iterating the computation and reassignment steps until a measure of average distance stops improving (see Elkan 2003). We select $k = 5$ clusters, following Martin, López and Olsen (2018), who used interpretive judgment to classify California municipal taxes into five categories. We assess the face validity of the five-cluster classification by inspecting the documents associated with each cluster, and the terms that most distinguish that cluster. The test of face validity is interpretive: whether or not a cohesive cluster of documents in the semantic space corresponds to a legally or administratively meaningful type of tax. We then assess the predictive validity of the classification, by testing whether membership in a cluster is a statistically significant predictor of the share of voters who favor the policy in question. Our aim at this stage is to test our approach to the quantitative representation of tax policies, by ascertaining whether an unsupervised algorithm applied to the matrix of TF-IDF weights can reproduce qualitative patterns that are known from the literature.

In the second stage of our analysis, we exploit the power of machine learning to go beyond what is already known from the literature. Our aim at this stage is to proceed from the classification of policies to a characterization of the dimensions that distinguish them in conceptual space. Presented with a description of a tax policy, can we identify any meaningful dimensions of the policy—i.e., linear combinations of policy features—that reliably predict the share of voters who would favor it? To answer this question, we undertake a supervised principal components analysis (PCA). In contrast to probabilistic topic models, PCA is a simple algebraic technique that entails no assumptions about such matters as the number of latent “topics” or the functional form of their probability distributions. It yields a basis transformation of document-term matrix, so that each text, instead of being represented as a single vector of TF-IDF weights, is represented as a linear combination of vectors called principal components. PCA produces components in decreasing order of variance. *Supervised PCA* (Bair et al. 2006) is an application of PCA to regression analysis. In a high-dimensional setting, when the principal components are intended as predictors in a regression model, the consistency of the regression estimator can be improved by selecting only a subset of terms for inclusion in the PCA, based on the strength of their correlation with the outcome variable; the outcome variable is said to train or “supervise” the PCA.

Our application of supervised PCA proceeds as follows. Let y be the affirmative vote share, and let X be the document-term matrix. First, we measure the association between the vote and each individual term by fitting a univariate regression of y on each column of X considered singly. Second, for each $j = 20, 21, \dots, n-1$, we select the subset of j terms whose univariate regressions yielded the greatest standardized regression coefficients in absolute value, and we compute a PCA of the TF-IDF weights of this subset of text features. (We assume 20 as an arbitrary but reasonable lower bound on the number of terms necessary for meaningful interpretation.) We then fit a univariate regression of y to the first principal component from this PCA. Third, we select the value of j that optimizes model fit in a univariate regression of y on the first principal component (see Bair et al. 2006). Finally, we add additional principal components stepwise to the regression model until the model fit stops improving. The result is a simplified description of the policy feature space in terms of a few dimensions that best predict the popularity of the policy. As with any PCA, the solution is unique except for the arbitrary sign of the components; for interpretive convenience, we report principal components signed so that positive loadings are associated positively with the share of voters affirming the tax.

Our measure of model fit is the square root of the mean error from out-of-sample prediction. In a high-dimensional setting such as this one, conventional measures of model fit such as R^2 are useless for variable selection; with 850 tax measures and 1,548 unique terms, a regression of the affirmative vote share on the TF-IDF weights of any non-perfectly-collinear subset of 849 terms would yield an R^2 of 1.00. Measures of fit based on out-of-sample prediction, by contrast, are comparatively resistant to over-fitting. Our preferred measure of out-of-sample prediction, because it yields a unique solution, is the leave-one-out cross-validation statistic, which is the average of n squared prediction errors computed by fitting the model to n unique subsamples of $n-1$ observations each, and using the results of each of these n implementations to predict the value of y for the omitted observation. We report the square root of this statistic below as the Root Mean Squared Error (RMSE).

Once we have identified meaningful principal components, the final step is to test whether knowledge of these components improves our ability to predict the popularity of a tax. We establish a baseline by fitting linear regression models of the affirmative vote share y on

contextual covariates that include the home ownership rate, the share of the population that is Black, the share of the population that is Latinx, the share of the population with 4-year college degrees or more education, the population (logged), the median household income (logged), and municipal tax revenues per capita (logged) (data from Martin, López and Olsen 2018). Our baseline model also controls for the share of the municipal electorate that voted Republican in the last presidential election, standardized across candidates by subtracting the national Republican share of the vote in the last presidential election, and imputed to off-year elections by linear interpolation; this measure of partisanship has been shown to predict some dimensions of municipal tax policy (see Einstein and Kogan 2016). We also estimate fixed-effects models with implicit city-level intercepts (to adjust for unmeasured, time-constant differences among cities in their propensity to favor certain policies) and a dummy variable for each year (to adjust for unmeasured differences over time that might be assumed to affect all municipalities similarly, such as differences in turnout between presidential and off-year elections). Rivlin’s search for a popular tax aimed to identify a tax policy design that was consistently popular, regardless of partisan context. We approximate this search by testing whether principal components of policy design improve the predictive performance of a linear regression model relative to our baseline regression specification.

The promise of this method is to produce not only a list of popular taxes, of which there are many in this sample, but also an inductive characterization of *dimensions* of policy design that make the greatest difference for whether or not a tax is popular. The null hypothesis that there are no such dimensions is also plausible. A municipal tax referendum presents voters with a dichotomous choice. Assuming that voters compare a proposed tax policy to the status quo, we might expect that it is the popularity of the status quo, rather than the policy of the proposed tax increase, that drives the decision (see Romer and Rosenthal 1978). If a city voter is well satisfied with the current level of municipal public goods provision, for example, then that voter may be motivated to oppose a tax increase, regardless of the policy details. If a voter is dissatisfied with the current level of public goods provision, and wishes for greater municipal spending, then that voter may have no option but to support the tax increase—again, regardless of its precise details. There is no reason to assume *a priori* that *any* qualitative policy details will matter for the popularity of a tax increase, relative to the status quo.

Varieties of tax policy

We find that the vector representation of impartial analysis documents can indeed produce a meaningful classification of tax policies. We begin with a *k*-means cluster analysis of the term-document matrix. Table 1 summarizes our five-cluster solution in a form suitable for interpretation, by listing the terms that have the greatest average TF-IDF weights within each cluster. These clusters correspond clearly to types of taxation that are distinct in California law and in local tax policy discourse (see Ingenito et al. 2007). There is one cluster each corresponding to the categories of sales tax, transient occupancy tax, business license tax, real estate parcel tax, and utility users’ tax.

[Table 1 here]

These categories of taxation are not only legally and semantically distinct from one another, they are also politically salient. To assess the predictive validity of this classification, we regressed the affirmative vote share, measured in percentage points, on a vector of dummy variables representing these clusters of tax policy. Table 2 reports the coefficients of the dummy variables, with utility users’ taxes as the reference category. The first column reports results from

a sparse OLS model that includes only the dummy variable for each cluster. The pattern of coefficients indicates that real estate parcel tax received the greatest share of affirmative votes, followed, in order, by transient occupancy taxes, business license taxes, sales taxes, and utility users' taxes. This rank order corresponds to the pattern reported by Martin, López and Olsen (2018: 16) in a human-classified data set that applied a similar classification scheme to California municipal tax measures over the same period. (Two differences of note are that Martin, López and Olsen [2018] did not classify business license taxes separately, and that they included a residual "other or multiple" category that was second to real estate parcel taxes in popularity.) The second column of Table 2 reports the coefficients of the same dummy variables from a fixed-effects regression model that includes controls for contextual covariates, with dummy variables for year and implicit city-level intercepts. The coefficients of this model can be read as expected differences in the affirmative vote share across policies of different types within the same city and year. The coefficients for parcel tax, transient occupancy tax, and sales tax remain comparatively large and sharply estimated. The type of tax policy that was proposed, as inferred solely from quantitative information about the co-occurrence of terms in the city attorney's impartial description, was predictive of its relative popularity in a referendum election. The most popular type of tax policy proposal, a transient occupancy tax, could expect a greater affirmative share of the vote by 8 percentage points than the least popular type, a utility users' tax (cf. Martin, López and Olsen 2018).

[Table 2 here]

This cluster analysis validates our quantitative representation of policy text, and yields evidence that some kinds of taxation are indeed more popular than others on average. But what is it about those policies that makes them more popular with voters? Having established that a highly simplified quantitative representation can encode legally meaningful and politically salient differences among tax policy proposals, we turn now to supervised PCA to identify the dimensions that distinguish popular from unpopular tax policies.

The dimensions of tax policy

Our supervised PCA algorithm yields a simplified representation of policies in terms of a few dimensions that reliably predict the popularity of a tax. The PCA solution that minimizes out-of-sample prediction error (RMSE = 14.8) includes 195 terms from the policy documents, and summarizes their associated TF-IDF weights with two principal components. We describe these components in Figure 1, which is a scatterplot of all 195 terms in the vector space spanned by the two principal components. Each term is centered on coordinates that correspond to its factor loadings, which may be understood as the multipliers that would be applied to TF-IDF weights to represent the position of a given policy description in this vector space. Words that are closest to each other in this two-dimensional space are words that tend to co-occur in the same tax policy proposal; words that are farthest apart are words that tend most clearly to distinguish tax proposals from one another. We have highlighted illustrative terms to improve the legibility of the scatterplot and facilitate interpretation.

[Figure 1 here]

We label the first principal component *risk pooling*. This concept designates a policy that raises funds from a category of persons who are at common risk for a particular kind of misfortune, and allocates those funds to mitigate the misfortune for those who experience it. The determination that a measure is risk pooling requires information about both how the revenue is raised and how it is spent; risk pooling involves both the categorical allocation of the fiscal

burden and the ostensible earmarking for safety net spending of some kind. The paradigmatic example at the federal level is a social insurance program such as Old Age Insurance: funds are raised from people in certain occupational categories, and spent to mitigate poverty and social exclusion in old age for people in those categories. The opposite of a policy that pools risk would be a policy that assigns the cost of mitigating misfortune to the individual.

At the local level, the component that we call risk pooling distinguishes between taxes that socialize risk and taxes that individualize it. The terms that load positively on this component indicate *special* taxes, a term of art in California law for taxes that are earmarked, and most especially those dedicated to such purposes as *paramedical* and *medical* services. The taxes that score most positively on this dimension also are taxes on real estate, distinguished by such terms as *parcel*, *unit*, and *square feet*. An example of a measure that scores very positively on this dimension would be the taxation of town residents to pay for local ambulance services; such a policy follows a classic risk pooling insurance model, and could reasonably be described as a local form of social insurance. It is a tax on all residents of a community in exchange for concentrated benefits available to those residents who experience medical misfortune.¹ Perhaps unsurprisingly, the term *insure* and its cognates also load positively on this dimension, although this is not among the strongest loadings. By contrast, the terms that load negatively on risk pooling are associated with flat-rate utility users' taxes. They include such terms as *utility*, *user*, *electricity*, *telephone*, *television*, *cable*, and *water*. Utility users' taxes are, in effect, mandatory price increases on goods whose consumption is indispensable for social life. They individualize the costs of social exclusion by requiring households to bear the risk of future price increases for such necessary resources as heat and water. In this sense, they may be conceptualized as the mirror image of social insurance.

We illustrate the dimension of risk pooling qualitatively in the top panel of Figure 2 by contrasting two exemplary impartial analysis texts that score respectively high and low on this dimension. One is the measure that scored highest on risk pooling (Measure G from the Town of Ross, which passed with 77.6% of the vote in 2002). The other is one of the measures that scored *lowest* on risk pooling (Measure E from the City of Scotts Valley, which failed with 37.7% of the vote in 1996). Although this measure was not the very lowest-scoring, it is a convenient example for illustrative purposes because of the concision of its impartial analysis. All of the lowest-scoring measures were utility users' taxes quite similar to this one.

[Figure 2 here]

We label the second principal component *community orientation*. We intend the term to resonate with Tönnies's (1957 [1887]) classic sociological distinction between two systems of positive law, which he called "community" and "society" respectively. Community, in Tönnies's conceptual scheme, expresses the shared values of neighbors who live together in a small settlement; it is "based on land ownership" (1957: 223), is administered with little bureaucracy, and attains its most developed form in the town (1957: 227). The conceptual opposite of community is *society*, a system of law that "derives from the conventional order of trade and similar relations" (p. 223); this form of law is based on commercial relations among strangers, is

¹ Although a parcel tax is nominally a tax on property owners, it is borne by renters as well, who can be expected to pay higher rents where a parcel tax supports valued local services (see e.g. Lang and Sonstelie 2015: 553-554). This difference between legal and real incidence is consistent with classic examples of social insurance; for example, the Social Security payroll tax for covered workers in the U.S. is nominally split between employees and their employers, but the employer's share of the tax is also borne economically by employees.

administered by a bureaucratic state, and emerges “only when the town develops into the city” (1957: 227).

In our data, we apply the name community orientation to a dimension that distinguishes taxes on town residency from taxes on city commerce. The term with the second strongest positive loading on this dimension is *town*; the term with the strongest negative loading is *city*. Legally, these are interchangeable terms in California (see Cal. Gov. Code §34502), but symbolically, they evoke a contrast between the bedroom communities of the exurban fringe and the denser settlements of the urban core. Taxes that score positively on this component tend to target residency rather than commercial activity; all of the thirty measures that are most associated with this component are parcel taxes or utility taxes that target residents. The terms *residential* and *non-resident* both load positively on this dimension: policies that score high on community orientation tend to discriminate explicitly between residents and non-residents. The terms that load positively also indicate *special* purposes such as paramedical services, police and fire protection, street and park maintenance, and deficit reduction. Taxes that are negatively associated with community orientation, by contrast, are general purpose taxes (as indicated by terms such as *general* and *government*) on commercial activity (as indicated by such terms as *percent* and *rental*), especially on forms of commerce that may be morally contested or disreputable (as indicated by such terms as *marijuana* and *cardroom*). This dimension evokes a symbolic contrast between the moral order of the residential town and the amoral legal order of the commercial metropolis.

The bottom panel of Figure 2 offers an illustrative example of a measure that scores very low on community orientation. This is the impartial analysis of one of the lowest-scoring measures, Measure R from the City of Seaside in the County of Monterey, which passed with 58.2 percent of the vote in 2008. The description is exemplary both for its emphasis on taxing transactions and for its emphasis that the tax will be deposited in the general fund and used for general government purposes. Seaside is a small city in the Salinas Metropolitan Statistical Area; it is arguably, to judge by its size and position in the regional commuting network, a suburb. However, even in a small, suburban municipality such as Seaside, tax measures, such as Measure R, that score low on this dimension are oriented towards extracting funds from commercial activity to address problems, such as graffiti and traffic congestion, that may be seen as characteristically urban.

The two dimensions of risk pooling and community orientation allow us to go beyond a simple cluster analysis of tax policy proposals, by representing the relationships *among* types of tax policy as well as relationships among additional policy features that are not captured in our simple five-category typology. Figure 3 illustrates the relationships among our five clusters of tax policy proposals, by plotting each cluster in the space spanned by the two dimensions of risk pooling and community orientation. Despite the radical simplification involved in reducing a 1,548-dimensional representation of policy to a two-dimensional plane, the clusters remain clearly distinct from one another. The graph also reveals systematic relationships among them. As our description of risk pooling has already suggested, this dimension most clearly distinguishes real estate parcel taxes from utility users’ taxes, revealing them to be not only different types, but also, in at least this sense, conceptual opposites to each other. The other types of taxation occupy intermediate positions with respect to risk pooling. Conceptually, these other local taxes neither fully pool risk to support social services, nor completely individualize the costs of social inclusion, although most types of local taxation involve *some* collectivization of risk insofar as they support public goods. The greatest values of community orientation are

exhibited by some real estate parcel taxes—the paradigmatic taxes on town residence rather than urban commerce—but as a group, real estate parcel taxes exhibit a wide range of values with respect to community orientation, and many utility users’ taxes are also positively community oriented. The taxes that are lowest in community orientation are generally sales taxes, transient occupancy taxes, and business license taxes, all of which are, indeed, types of taxation associated with urban commerce. The graph illustrates that types of taxation that we previously found to differ in popularity are also conceptually distant from each other with respect to the two dimensions we have identified.

The representation of tax policy proposals in terms of risk pooling and community orientation also improves our ability to predict which of these proposals will be popular. We fit a series of regression models of the affirmative vote share, treating the two principal components of risk pooling and community orientation as predictors. We report our measure of out-of-sample prediction error, the RMSE from leave-one-out cross-validation, in Table 3. For comparison, we also report the RMSE for models that represent the type of tax policy as a categorical variable, corresponding to the five-category classification we have already introduced. The first column of Table 3 presents the results from baseline models with no control variables. In such a model, the representation of tax policy in terms of risk pooling and community orientation yielded an RMSE of 14.8, compared to 15.0 for a similar model that represented tax policy as a set of five categories. Including both sets of variables yielded RMSE of 14.8, suggesting that there is no advantage to including dummy variables for the type of taxation if one is already measuring the degree of risk pooling and community orientation. The second column of Table 3 reports the results of models that control for the political, social and fiscal context. With this specification, a model that represented tax policy in terms of the two principal components performed comparably to a model that represented tax policy in terms of five categories (RMSE = 14.2). The final column reports results of fixed-effects models that include implicit city-level intercepts and year dummy variables. In this specification, both representations of tax policy performed comparably (RMSE = 10.4).

[Table 3 here]

In short, these two ways of measuring tax policy both have similar predictive validity. Our description of tax policy in terms of the two principal components of risk pooling and community orientation is based on the same documents, and summarizes much of the same information, as our descriptive classification of tax policies into the five categories of utility users’ tax, parcel tax, business license tax, transient occupancy tax, and sales tax. But faced with a choice between these two ways of representing a tax policy, the conceptual representation in terms of the two continuous dimensions of risk pooling and community orientation is more concise, more closely related to theoretical concepts in political sociology, and slightly more predictive of actual political outcomes.

Policy design and political context

Just how closely are risk pooling and community orientation related to voters’ support for tax policy proposals? Table 4 reports coefficients and standard errors from a series of linear regressions of the affirmative vote share on the two principal components in our full population of 850 tax measures. All coefficients are semi-standardized, so that they can be read as percentage points of the vote associated with a one-standard-deviation difference in the relevant principal component; all standard errors in the table are clustered at the city level. The first column represents coefficients from a specification without any control variables. A difference

of one standard deviation in risk pooling is associated with a difference of 3.6 percentage points in the share of voters approving a tax. A difference of one standard deviation in community orientation is associated with a difference of 2.6 percentage points in the share of voters approving a tax. In this specification, both risk pooling and community are associated with the affirmative share of the vote, and both coefficients are estimated with relatively small standard errors.

[Table 4 here]

These differences may be spurious if different policies are proposed in different contexts. In particular, the loadings of the terms “city” and “town” in the component we have labeled community orientation might be taken to suggest that this component is indirectly measuring a difference between places, rather than a theoretically important difference in policy design. The second column of Table 3 reports results from a fixed-effects model that adjusts for differences among places by centering the independent and dependent variables at their group means (i.e., the sample means for the respective municipality). This approach computes the average *within-municipality* association between principal components and the affirmative vote share, or the association net of unmeasured, time-invariant differences among municipalities. In the model with this fixed-effects adjustment, we observe that the coefficient of community orientation is negative and statistically indistinguishable from zero. Tax policies that score high on community orientation may be comparatively popular simply because they are the sorts of policies typically proposed in places where voters are most inclined to support tax increases.

The coefficient of *risk pooling*, by contrast, remains positive, substantial, and sharply measured in every regression specification. In the regressions reported in the third and fourth columns of Table 3, we introduce our vector of contextual covariates. The third column reports results without, and the fourth column results with, a fixed-effects adjustment. Our most conservative model, in the fourth column, adjusts for contextual covariates; for unmeasured, time-invariant characteristics of the municipalities; and for unmeasured, city-invariant factors that vary from year to year. The coefficient in this model can be read as our best estimate of the expected difference in the degree of voter support between two different tax policy proposals that differed only with respect to their degree of risk pooling, if those two policy proposals were presented to the same voters in the same city and year. According to this model, a one standard deviation difference in risk pooling was associated with a 2.3 percentage point difference in the expected share of voters approving a tax policy. A difference of this magnitude might be enough to sway a close election.

Our inductive treatment of tax policy data thus yields two dimensions that are both theoretically meaningful and politically salient. Although we have highlighted the ways in which these dimensions evoke some familiar sociological concepts, we should also note that they are different from the dimensions that might be suggested by prior research. Policy terms that did not load on either factor include, e.g., some that might be associated with the moral worthiness of target populations (such as *contribution*, *senior*, *citizen*, *elderly*, *child*, and *veteran*), the visibility of public spending (*park*, *street*, *build*), and the putative invisibility of taxation (*sales*, *consumption*, *withholding*). All of these terms are represented in the corpus of policy text analyzed here, and some of them are quite amply represented, with TF-IDF weights that exhibit considerable variance. But they do not have a sufficiently strong or consistent association with aggregate voting behavior to be selected by the supervised PCA. In this sample of 850 elections, we have found no evidence that these policy descriptors, or others like them, are politically salient at all.

The search for a popular tax continued

The analysis reported here may shed some light on why Rivlin's question has gone unanswered for so long. Many taxes are popular in context, and some of them are very popular indeed; but it is difficult to offer any general statement about which *kinds* of taxation are most popular that is both true to the data and sufficiently abstract to be useful. On the one hand, we have found that the popularity of a tax policy may be systematically related to the kinds of policy features that are described in a city attorney's impartial analysis. On the other hand, our most predictive models rely heavily on a fixed-effects adjustment for unmeasured differences among cities and over time. A comparison of models within each column of Table 3 reveals that the out-of-sample prediction error is relatively insensitive to whether the regression model includes tax policy characteristics only, contextual covariates only, or both; the election outcomes may be predicted about equally well from information about the policy or information about the social and political context, and combining both kinds of information does not much improve our predictions. We may infer that risk pooling and community orientation are politically salient characteristics of tax policy—but we also may infer that most of the politically salient information contained in these principal components is redundant with the information contained in the contextual covariates. In other words, the taxes with the most popular features are also often proposed in the most favorable socio-political contexts, and it is difficult to distinguish observationally between the effects of policy design and the mere fact of favorable conditions. Popular taxes are popular in context, and the dream of a policy design mechanism that is consistently powerful enough to overwhelm the effects of political and social context is likely to prove chimerical.

Still, if no kind of tax is universally the most popular, some kinds are, on average, more popular than others. It would be wise to focus more of our theoretical attention on the dimensions that distinguish these taxes. We identify two such dimensions, risk pooling and community orientation, that differentiate popular from unpopular taxes in a sample of 850 municipal referendum elections in California. Both of these dimensions concern *social* dimensions of taxation: they distinguish among policies according to their social purposes and according to how broadly or narrowly they draw the circle of those who are obligated to pay. The most popular taxes in this sample are those that assign a common obligation to everyone who owns a parcel in the residential community, and that earmark the funds for safety net services available to everyone in the residential community. The least popular taxes are those that increase pro-rated charges for basic utilities that are necessary for survival and social inclusion, and that do not reserve the resources for those in the pool of contributors. These patterns are consistent with the hypothesis that social insurance is among the most popular functions of fiscal policy, and more generally with the idea that taxation can be popular when it is bound to membership in a civic community (see Williamson 2017). Other dimensions of tax policy—whether or not it is coupled with sales transactions, or earmarked for especially morally worthy populations, or symbolically tied to the virtue of work, or labeled a “contribution,” or split between employers and employees; whether it is collected in frequent or infrequent increments; whether it is painfully visible or deducted invisibly from another transaction—do not make much of a difference, to judge from the tax policy descriptions in this sample.

We should not assume that California municipalities at the turn of the 21st century are representative of other polities, of course, but neither should we assume that these findings are

only relevant to California, or to municipal government, or to the turn of the 21st century. The sampling design of the present study shares a weakness of all of the most influential works on how voters respond to public policies, none of which has examined a representative sample of policies or politics (but see Burstein 2014). But the approach taken here—to examine a large population of policies, to adopt a coding protocol that sticks close to the textual data on which it is based, and to focus on observational measures of consent that do not depend on features of an artificial experimental setting or survey encounter—may help to overcome other weaknesses of previous studies. Our approach exploits the distinctive strengths of modern computational social science. The data for future such studies are richly available, and we think this approach is worth pursuing further.

If Rivlin was right, then the search for a popular tax will continue. The good news from the present study is that the most difficult part of the search is over. Popular taxes are quite common. The challenge is no longer to find one. It is to analyze the richly available data concerning the large numbers of popular tax policies that already exist, in order to discover principles that can inform the design of the *next* one. Although this approach is unlikely ever to discover a silver bullet that will end the bitter struggle over how to pay for public goods, we think it can inform the design of tax policies that are more likely to survive the democratic test.

References

- Arnold, R. Douglas. 1991. *The Logic of Congressional Action*. New Haven: Yale University Press.
- Bair, Eric, Trevor Hastie, Debashis Paul, and Robert Tibshirani. 2006. "Prediction by Supervised Principal Components." *Journal of the American Statistical Association*. 101(473): 119-137.
- Bair, Eric, Trevor Hastie, Debashis Paul, and Robert Tibshirani. 2006. "Prediction by Supervised Principal Components." *Journal of the American Statistical Association*. 101 (473): 119-137.
- Barnes, Lucy. 2014. "The Size and Shape of Government: Preferences over Redistributive Tax Policy." *Socio-Economic Review* 13 (1): 55-78.
- Bird, Steven, Ewan Klein and Edward Loper. 2009. *Natural Language Processing with Python*. Sebastopol, California: O'Reilly.
- Burstein, Paul. 2014. *American Public Opinion, Advocacy, and Policy in Congress: What the Public Wants and What It Gets*. New York: Cambridge University Press.
- Campbell, Andrea L. 2003. *How Policies Make Citizens: Senior Political Activism and the American Welfare State*. Princeton: Princeton University Press.
- Campbell, Andrea L. 2012. "Policy Makes Mass Politics." *Annual Review of Political Science* 15: 333-351
- Campbell, Andrea L. 2018. "Tax Designs and Tax Attitudes." *The Forum* 16 (3): 369-397.
- Campbell, Andrea Louise. 2014. "Tax Attitudes in the Obama Era." *Tax Law Review* 67: 647-668.
- Campbell, John L. 1993. "The State and Fiscal Sociology." *Annual Review of Sociology* 19: 163-185.
- Dodson, Kyle. 2017. "Economic Change and Class Conflict over Tax Attitudes: Evidence from Nine Advanced Capitalist Democracies." *Social Forces* 95 (4): 1509-1538.
- Einstein, Katherine Levine and Vladimir Kogan. 2016. "Pushing the City Limits: Policy Responsiveness in Municipal Government." *Urban Affairs Review* 52(1): 3-32.
- Elkan, Charles. 2003. "Using the Triangle Inequality to Accelerate *k*-Means." Pp. 147-153 in *Proceedings of the Twentieth International Conference on Machine Learning (ICML-2003)*. Washington, D.C.: AAAI Press.
- Graetz, Michael and Ian Shapiro. 2005. *Death by a Thousand Cuts: The Fight over Taxing Inherited Wealth*. Princeton: Princeton University Press.
- Grimmer, Justin and Brandon M. Stewart. 2013. "Text as Data: The Promise and Pitfalls of Automatic Content Analysis Methods for Political Texts." *Political Analysis* 21 (3): 267-297.f
- Gulick, Luther. 1941. "Memorandum on Conference with FDR Concerning Social Security Taxation, Summer 1941." Retrieved August 3, 2018 from <https://www.ssa.gov/history/Gulick.html>
- Haselswerdt, Jake and Brandon L. Bartels. 2015. "Public Opinion, Policy Tools, and the Status Quo: Evidence from a Survey Experiment." *Political Research Quarterly* 68 (3): 607-621.

- Hibbs, Douglas A. and Henrik Jess Madsen. 1981. "Public Reactions to the Growth of Taxation and Government Expenditure." *World Politics* 33 (3): 413-435.
- Ingenito, Robert, Marianne O'Malley, Marcus Stanley, and Crystal Taylor. 2007. *California's Tax System: A Primer*. Sacramento: California Legislative Analyst's Office.
- Ikenberry, G. John and Theda Skocpol. 1987. "Expanding Social Benefits: The Role of Social Security." *Political Science Quarterly* 102 (3): 389-416.
- Isaac William Martin and Nadav Gabay. 2018. "Tax Policy and Tax Protest in 20 Rich Democracies, 1980-2010." *British Journal of Sociology*. 69 (3): 647-669.
- Jacobs, Lawrence R. and Robert Y. Shapiro. 2000. *Politicians Don't Pander: Political Manipulation and the Loss of Democratic Responsiveness*. Chicago: University of Chicago Press.
- Lang, Bree J. and Jon Sonstelie. 2015. "The Parcel Tax as a Source of Revenue for California Public Schools." *National Tax Journal* 68 (3): 545-572.
- Larsen, Erik Gahner. 2018. "Policy Feedback Effects on Mass Publics: A Quantitative Review." *Policy Studies Journal*. Published online before print.
- Martin, Isaac William. 2010. "Redistributing toward the Rich: Strategic Policy Crafting in the Campaign to Repeal the Sixteenth Amendment, 1938-1958." *American Journal of Sociology* 116(1): 1-52.
- Martin, Isaac William and Monica Prasad. 2014. "Taxes and Fiscal Sociology." *Annual Review of Sociology* 40: 331-345.
- Martin, Isaac William, Ajay K. Mehrotra, and Monica Prasad. 2009. "The Thunder of History: The Origins and Development of the New Fiscal Sociology." Pp. 1-28 in Isaac William Martin, Ajay K. Mehrotra, and Monica Prasad, eds., *The New Fiscal Sociology: Taxation in Comparative and Historical Perspective*. New York: Cambridge University Press.
- Martin, Isaac William, Jane Lilly Lopez, and Lauren Olsen. 2017. "California Local Tax Ballot Measures." UC Library Digital Collections. doi:10.6075/J09P2ZT7.
- Martin, Isaac William, Jane Lilly López, and Lauren Olsen. 2018. "Policy Design and the Politics of City Revenue: Evidence from California Municipal Ballot Measures." *Urban Affairs Review*. Published online before print, January 22.
- McCaffery, Edward J. 1994. "Cognitive Theory and Tax." *UCLA Law Review* 41 (7): 1862-1996.
- McCaffery, Edward J. and Jonathan Baron. 2006. "Thinking about Tax." *Psychology, Public Policy, and Law*. 12 (1): 106-135.
- Mill, John Stuart. 2004 [1871]. *Principles of Political Economy with Some of Their Applications to Social Philosophy*. Indianapolis: Hackett Publishing Company.
- O'Brien, Rourke L. 2017. "Redistribution and the New Fiscal Sociology: Race and the Progressivity of State and Local Taxes." *American Journal of Sociology* 122 (4): 1015-1049.
- Pampel, Fred, Giulia Andrighetto, and Sven Steinmo. 2019. "How Institutions and Attitudes Shape Tax Compliance." *Social Forces* 97 (3): 1337-1364.
- Pearson, Elizabeth. 2014. "Saying Yes to Taxes: The Politics of Tax Reform Campaigns in Three Northwestern States, 1965-1973." *American Journal of Sociology* 119(5): 1279-1323.
- Pierson, Paul. 1993. "When Effect Becomes Cause: Policy Feedback and Political Change." *World Politics* 45:595-628.

- Pierson, Paul. 1994. *Dismantling the Welfare State? Reagan, Thatcher, and the Politics of Retrenchment*. New York: Cambridge University Press.
- Porter, M. F. "An Algorithm for Suffix Stripping." *Program* 14 (3): 130-7.
- Prasad, Monica. 2005. "Why is France So French? Culture, Institutions, and Neoliberalism, 1974-1981." *American Journal of Sociology* 111 (2): 357-407.
- Rivlin, Alice. 1989. "The Continuing Search for a Popular Tax." *American Economic Review* 79 (2): 113-117.
- Romer, Thomas, and Howard Rosenthal. 1978. "Political Resource Allocation, Controlled Agendas, and the Status Quo." *Public Choice* 33 (4): 27-43.
- Roosma, Femke, Wim van Oorschot, and John Gelissen. 2016. "A Just Distribution of Burdens? Attitudes Toward the Social Distribution of Taxes in 26 Welfare States." *International Journal of Public Opinion Research* 28 (3): 376-400.
- Rousseeuw, Peter J. 1987. "Silhouettes: A Graphical Aid to the Interpretation and Validation of Cluster Analysis." *Journal of Computational and Applied Mathematics* 20 (1): 53-65.
- Schattschneider, E. E. 1974 (1935). *Politics, Pressures, and the Tariff: A Study of Free Private Enterprise in Pressure Politics, as Shown in the 1929-1930 Revision of the Tariff*. New York: Arno Press.
- Steensland, Brian. 2008. *The Failed Welfare Revolution: America's Struggle over Guaranteed Income Policy*. Princeton: Princeton University Press.
- Tönnies, Ferdinand. 1957 [1887]. *Community and Society*. Charles P. Loomis, trans. New York: Harper and Row.
- Wilensky, Harold L. 2002. *Rich Democracies: Political Economy, Public Policy, and Performance*. Berkeley and Los Angeles: University of California Press.
- Wilensky, Harold. 1976. *The 'New Corporatism', Centralization, and the Welfare State*. Beverly Hills: Sage Publications.
- Williamson, Vanessa. 2017. *Read My Lips: Why Americans Are Proud to Pay Taxes*. Princeton: Princeton University Press.

Table 1

Label	Distinguishing terms
Business license tax (N=70)	busi licens gross receipt citi fee measur ordin cannabi gener marijuana rate per increas annual employe approv year adopt vote
Sales tax (N=139)	sale citi transact use measur gener retail percent revenu cent state vote propos increas servic ordin rate approv voter purpos
Utility users' tax (N=168)	util user citi servic measur telephon gener electr ordin voter vote rate approv council collect water exempt fund revenu provid
Transient occupancy tax or hotel tax (N=179)	hotel occup transient citi tot room motel measur increas percent rate gener rent guest vote charg approv voter oper ordin
Parcel tax (N=294)	citi parcel special measur per year servic properti assess unit librari levi residenti approv paramed town vote fund increas amount

Table 2

Coefficients (and standard errors) from regressions of affirmative vote share on category of taxation, N=850 tax measures

	OLS, no controls	FE, with controls for contextual covariates
Parcel tax	10.3 (1.7)	5.4 (2.0)
Business license tax	9.5 (2.1)	8.0 (2.4)
Transient occupancy (hotel) tax	6.3 (1.6)	6.7 (2.0)
Sales tax	4.0 (1.6)	3.8 (2.1)

Note:

Utility users' tax is the reference category. Contextual covariates include the home ownership rate, the share of the population that is Black, the share of the population that is Latinx, the share of the population with 4-year college degrees or more education, the population (logged), the median household income (logged), and municipal tax revenues per capita (logged). "FE" refers to models with dummy variables for years and a "fixed effects" adjustment that introduces an implicit, fixed intercept for each city.

Table 4

Coefficients (and standard errors) from regression of affirmative vote share on two principal components, N=850 tax measures

	No controls		With controls for contextual covariates	
	Without FE	With FE	Without FE	With FE
Risk pooling	3.6 (0.6)	2.3 (0.8)	2.8 (0.6)	2.3 (0.8)
Community	2.6 (0.5)	0.7 (0.8)	1.7 (0.5)	0.8 (0.8)

Note: Contextual covariates include the home ownership rate, the share of the population that is Black, the share of the population that is Latinx, the share of the population with 4-year college degrees or more education, the population (logged), the median household income (logged), and municipal tax revenues per capita (logged). “FE” refers to models with dummy variables for years and a “fixed effects” adjustment that introduces an implicit, fixed intercept for each city.

Table 3

RMSE from regression models of affirmative vote share with and without fixed effects (FE),
N=850 tax measures

	Without controls	With controls for contextual covariates	
		Without FE	With FE
Five cluster indicators only	15.0	14.2	10.4
Two principal components only	14.8	14.2	10.4
Principal components and cluster indicators	14.8	14.2	10.4
Contextual covariates only	...	14.5	10.5

Note: Contextual covariates include the home ownership rate, the share of the population that is Black, the share of the population that is Latinx, the share of the population with 4-year college degrees or more education, the population (logged), the median household income (logged), and municipal tax revenues per capita (logged). “FE” refers to models with dummy variables for year and a “fixed effects” adjustment that introduces an implicit, fixed intercept for each city.

Figure 2.

Illustrative examples of tax measures that score high and low on *risk pooling* and *community orientation*

Low risk-pooling, high community orientation: Utility users' tax for the general fund	High risk-pooling, high community orientation: Parcel tax for paramedic services
<i>Impartial Analysis by City Attorney City of Scotts Valley, Measure E</i> <i>Measure E seeks voter approval of a utility users tax measure. A "yes" vote on Measure E will repeal the existing utility tax in the City of Scotts Valley and replace it with a new utility users tax of five percent (5%) which is applicable, with some exceptions, to all charges made to users of gas, electric and telephone services in Scotts Valley. A "no" vote on Measure E will reject the new utility users tax and have no affect on the existing utility users tax.</i>	<i>Special Tax for Paramedic Services Impartial Analysis by Town Attorney Measure G</i> <i>If this measure is approved by a two-thirds vote, the special tax for paramedic services administered by a Joint Powers Authority authorized by the voters in 1983, 1991, 1994 and 1998, will be continued with a maximum annual levy of \$30.00 on each living unit and \$30.00 per year per 1,500 square feet of structure on each developed parcel in non-residential use. The applicable appropriations limit established pursuant to Article __ of the California Constitution will be increased in an amount equal to the proceeds of the special tax. The special tax and the corresponding increase in the appropriations limit expire on June 30, 2007.</i>
Low risk-pooling, low community orientation: sales tax for general government	
<i>City of Seaside, Measure R City Attorney's Impartial Analysis</i> <i>The City Council of the City of Seaside placed Measure R on the February 5, 2008 ballot to ask voters to consider an ordinance enacting a general transactions and use (sales) tax of one percent (1%) to be administered by the State Board of Equalization. The Seaside Vital City Services Measure R proposes to raise revenue for general purposes, including police and fire personnel, recreation programming, street and traffic improvements and other general services, for the City of Seaside. The one percent tax would be paid in addition to current sales taxes and would be collected at the same time and in the same manner as existing sales taxes. This is a local tax that will be collected and placed in the City of Seaside's General Fund. This tax cannot be taken away by the State of California or the County of Monterey. The City of Seaside Vital City Services Measure R states that the one percent sales tax is to be used for City services such as</i>	

hiring additional police officers and firefighters, purchasing breathing apparatuses and emergency radio communications equipment, improving 9-1-1 police and fire response times, fixing potholes, improving and maintaining City streets, reducing traffic congestion, expanding youth and senior programs and improving parks, increasing code enforcement, expanding anti-gang and drug enforcement programs, removing graffiti and blight, and other general City services. This sales tax would be a general tax; therefore the tax would go into the City's General Fund and could be used for any legal municipal purpose. California Revenue and Taxation Code Section 7285.9 authorizes the City Council to levy a transaction tax at a rate of 0.25 percent, or a multiple thereof, provided the tax is approved by a majority of the voters voting in an election on that issue. The proposed tax will increase the local sales tax rate in City of Seaside from the current 7.25% to 8.25%. Measure R includes a requirement that the City of Seaside have published, independent annual audits available for public review that will specify how the funds raised by Measure R are spent for City Services. By placing Measure R on the ballot, the City complies with Article of the California Constitution, which requires the voters to approve an ordinance which imposes a general tax.

Figure 1. Scatterplot of 195 tax policy terms in a vector space spanned by the principal components *risk pooling* and *community orientation*, with illustrative terms highlighted, N=850 municipal tax policy proposals




Figure 3. Scatterplot of 850 tax policy proposals in a vector space spanned by the principal components *risk pooling* and *community orientation*, with clusters of tax policy highlighted

