

UC Santa Barbara

UC Santa Barbara Electronic Theses and Dissertations

Title

Defining the Community of Interest as a Criterion for Boundary Drawing of Electoral Districts

Permalink

<https://escholarship.org/uc/item/2461k332>

Author

Phillips, Daniel Wesley

Publication Date

2016

Peer reviewed|Thesis/dissertation

UNIVERSITY OF CALIFORNIA

Santa Barbara

Defining the Community of Interest as a Criterion for Boundary Drawing of Electoral
Districts

A Thesis submitted in partial satisfaction of the
requirements for the degree Master of Arts
in Geography

by

Daniel W. Phillips

Committee in charge:

Professor Daniel R. Montello, Chair

Professor Stuart H. Sweeney

Professor Heather Stoll

September 2016

The thesis of Daniel W. Phillips is approved.

Stuart H. Sweeney

Heather Stoll

Daniel R. Montello, Committee Chair

August 2016

ACKNOWLEDGEMENTS

My thanks to Edgar Basave for serving as an extremely competent and helpful research assistant during the survey administration. I am very grateful for the effort and attitude that he showed. I would also like to thank Adam Davis for allowing me to use his Python script to compute the degree of overlap between polygons. I also acknowledge my officemate Crystal Bae for inspiring me with her Master's work and offering helpful feedback and guidance. I value having her as a colleague and friend. Finally, I appreciate the outstanding mentorship, advice, and counsel offered by my advisor Dan Montello. Even while busy serving as the chair of the department he made time to give me the help I needed. Thank you very much, Dan.

ABSTRACT

Defining the Community of Interest as a Criterion for Boundary Drawing of Electoral Districts

by

Daniel W. Phillips

When deciding where to draw the boundaries for electoral districts, officials often strive to ensure that communities of interest are not split up but kept together within a single district. What exactly constitutes a community of interest is somewhat vague, with legal and academic sources describing either a thematic region with shared demographic and land use traits or a cognitive region that is meaningful to people and commonly agreed upon. This research seeks to identify communities of interest at the sub-city level in both the thematic sense—by clustering Census tracts and land parcels according to classes of important variables—and the cognitive sense—by surveying residents about the location and extent of their community and finding areas of highest agreement. Then the degree to which the two senses of communities overlap is assessed; the more overlap, the more evidence there is that the two definitions correspond closely. Finally, the amount of overlap between the different communities and existing electoral districts is determined, to see which of the two types receives more attention from boundary drawers. The study finds that the two types of communities of interest correspond relatively well, and that the thematic type corresponds with the electoral districts better than the cognitive type.

TABLE OF CONTENTS

I. Introduction	1
II. Literature Review.....	4
III. Overview of the Research.....	10
IV. Methods.....	12
A. Cases	12
B. Materials.....	14
1. Thematic Communities of Interest	14
2. Cognitive Communities of Interest.....	19
C. Procedure	21
D. Analysis	23
1. Thematic Communities of Interest	23
2. Cognitive Communities of Interest.....	28
V. Results.....	30
A. Thematic Communities of Interest	30
B. Cognitive Communities of Interest.....	33
C. Correspondence Between Communities of Interest and Electoral Districts	41
VI. Discussion.....	43
VII. Conclusion.....	53
References.....	54

LIST OF FIGURES

Figure 1. Map of Santa Barbara city council districts as of 2015.....	3
Figure 2. Conceptual diagram of the aims of this research	11
Figure 3. Sampling grids for the three districts surveyed.....	13
Figure 4. The front side of the sheet of paper given to survey respondents	20
Figure 5. The back side of the sheet of paper given to survey respondents	21
Figure 6. Map of demographic clusters in Santa Barbara	25
Figure 7. Map of land-use clusters in Santa Barbara.....	26
Figure 8. Map of thematic clusters in Santa Barbara.....	27
Figure 9. Lines drawn by residents of District 1	29
Figure 10. Thematic communities of interest associated with each district	31
Figure 11. Cognitive agreement and centroids for District 1	36
Figure 12. Cognitive agreement and centroids for District 3	37
Figure 13. Cognitive agreement and centroids for District 2	38

I. Introduction

Many democracies elect their representatives from carefully crafted districts, but the methods that governments use to draw their boundaries vary substantially. While many jurisdictions allow their public officials to tweak the lines to serve partisan interests, others opt to use a set of nonpartisan criteria to create districts that are more representative. One such criterion, referred to as “respecting the community of interest,” is the degree to which district boundaries unite—rather than separate—a *community of interest*, defined as a group of people with shared values, concerns, and cultural traits. The fact that dozens of polities utilize this criterion demonstrates the wide belief that respecting communities of interest is critical to ensuring effective and fair representation for members of these groups; when these individuals are kept together in a single district, the resulting homogeneity enables its representative to better focus on advocating for and catering to that group’s interests.

While a general consensus exists on the importance of setting districts that respect the community of interest, there is little agreement on any of the specific traits that characterize such a community. Unless this criterion is more precisely defined, when officials draw district boundaries, they may fail to uphold it as they would intend. Members of those communities will find themselves poorly represented as a result. Here I explore defining a community of interest according to thematic attributes and according to people’s beliefs about their community; I then compare the two approaches in the hope that both together might reveal whether a distinct community of interest exists, where it is located, and how far it extends. I delineate communities of interest both thematically, by examining demographic and land-use attributes, and cognitively, by evaluating agreement among people’s responses from surveys about what they consider to be their community of interest.

Then I analyze how communities of interest defined in these two ways correspond with one another and with existing electoral districts.

The city of Santa Barbara, California provides an excellent opportunity for investigation of this topic. Since 1968, the city has elected its city council members through at-large elections across the city. However, a series of recent events has brought a marked change to the status quo. In July 2014 a group of plaintiffs advocating for Hispanic voting rights sued the city over its electoral system, demanding that it use district elections instead. They argued that the at-large system had failed to adequately represent Hispanics, most evidenced by the fact that few of their number had been elected to the office. On February 24, 2015, the city settled the lawsuit by agreeing to switch its city council elections to a district system, initiating a relatively quick month-long process to create six single-member districts in time for that November's city council elections (Potthoff 2015; the districts are shown in Figure 1). The actual boundary drawing task was contracted out to a private company called National Demographics Corporation (NDC), which produced three draft plan maps for consideration by citizens at public meetings. Beyond the paramount requirement of roughly equal populations and a settlement stipulation that two majority-Hispanic districts be created, NDC aimed to conform to the traditional criteria of boundary adherence, contiguity, compactness, and respect for perceived communities of interest (Johnson 2015). Much of the discussion at the public forums involved which neighborhoods seemed to "belong" in the same district as others, for example whether "Oak Park had more connection to downtown" based on similar population density and renter to owner ratio (NDC 2015b, 4). After public comment made it clear that Draft Plan 3 "had proven to be the

most popular choice,” the city council voted on March 30, 2015 to adopt a version of that plan for official use (City of Santa Barbara 2016a).

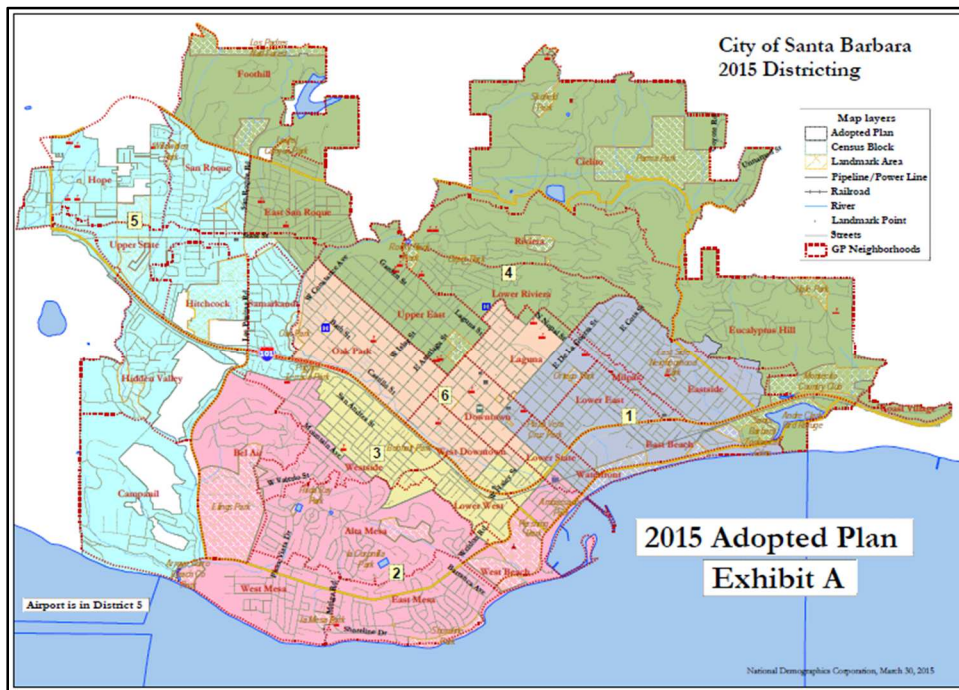


Figure 1. Map of Santa Barbara city council districts as of 2015. Source: http://www.santabarbaraca.gov/gov/vote/district_elections.asp

In light of these developments, Santa Barbara looks to be an informative and timely place to attempt to explore thematic and cognitive communities of interest. First, I believe the results from this city can be effectively generalized to other urban areas of similar size due to the great diversity in many of its thematic attributes that may distinguish certain communities of interest, for example ethnicity, income, education, land use, etc. To illustrate, Santa Barbara can serve as an example for cities with much disparity in education among their communities of interest, but it can also serve as an example for those cities whose communities of interest are more differentiated by land use. Second, the fact that these districts are brand new means that few residents are aware of their existence, let alone have informed opinions about them, so cognitive communities of interest would likely be

conceived apart from the influence of authorities' boundary decisions. Finally, the analysis in my research can assess how well the city carried out its hurried districting project vis-à-vis the community of interest criterion so as to inform those municipalities that do the same, because of litigation or otherwise.

II. Literature Review

In a representative government, the people delegate their governing powers to a professional who can dedicate most of his or her time and energy to making informed decisions on the conduct of the state. Most such governments feature either a proportional system, where the goal is to produce an inclusive government reflecting the diversity of opinion, or a majoritarian system, where the aim is to produce a decisive government identifiable to voters as deserving of rewarding or punishment at the polls (Htun and Powell 2013). In majoritarian countries elections are usually structured to give one party a convincing majority over a sizeable opposition party in the legislature, and rules of cabinet appointments and committee assignments are designed to enable that majority to dominate and push through its agenda essentially unchallenged. In contrast, elections in proportional countries are generally structured to result in a multiparty legislature with no party achieving a majority, thus necessitating coalition agreements and their ensuing bargaining and compromising (Powell 2000). Majoritarian systems are commonly designed by confining representatives to single-member districts in which victory goes to the candidate with the plurality of votes, known as first-past-the-post (FPTP) voting. Four notable countries still utilize FPTP voting for at least the lower house of their respective legislatures: Canada, India, the United Kingdom, and the United States. In these countries the importance of the results in individual electoral districts is inflated due to the small size of the districts

(just one seat for each one). This often leads to a disparity between votes won and seats won for each party and, therefore, a disproportional legislature (Powell 2004).

Because single-member districts play such a central role in the election results in these nations, the following attributes hold great significance to the final outcome: the total number of districts; how much the size of population in each district may deviate from the average figure for all districts; the geographical location and extent of each district; and the demographic, socioeconomic, and political characteristics of the population in each district. Constitutions, statutes, and/or courts commonly establish the first two of these qualities, but the third is subject to continual readjustment, which in turn affects the fourth. This readjustment is known as redistricting, and it has formed an important component of the political process in these countries ever since the courts there have deemed that districts should have roughly equal population sizes; when population shifts result in intolerable deviations from the equal population standard, redistricting becomes necessary.

However, this requirement raises the critical issue as to who will decide the geography of these districts and what criteria they will use to make their decisions. Placing the power to draw or redraw district boundaries in the wrong hands may well result in egregious abuses of the redistricting process, as those responsible will opt to base their decisions on partisan gain rather than an objective set of criteria. Many political science scholars cite the partisan gerrymandering carried out by many state legislatures in the United States as a prime example of such misconduct (Vickrey 1961; Handley 2008; Winburn 2008). These abuses suggest the need for two important elements: an impartial body of key actors to perform the boundary drawing task, and the use of a fair and optimal set of criteria to guide these actors in making districts that are as representative as possible.

However, representation is not a simple concept that can be boiled down to a single meaning. Rather, the idea incorporates multiple senses, two of the most prominent being the descriptive and substantive meanings. Descriptive representation concerns how closely the representative body reflects the constituents in terms of external characteristics like race, ethnicity, or sex. Substantive representation, on the other hand, involves how well that body acts in the interests of those constituents (Pitkin 1967). Whether districts should be made to cater to one or both of these senses of representation will be considered below.

Assuming the first elements of impartial drawers of district boundaries is in place—and resolving what particular body with what composition of members to charge with this task is no straightforward exercise—the next step involves what specific criteria should inform the drawing. Four specific criteria stand out for their frequent appearances as stated goals in district boundary drawing around the world: consideration of existing local administrative boundaries, contiguity of shape, compactness of shape, and respect for communities of interest (Mann 2005; Handley 2008). Most of these are defined easily enough: The first criterion involves making district lines correspond to administrative ones as closely as possible; the second refers to keeping a district as a single coterminous shape instead of disconnected pieces; and the third concerns ensuring that a district has a rounded, sensible shape instead of a sinuous, convoluted one. Of these three, the compactness criterion is certainly the trickiest to measure due to there being a number of possible ways to quantify it (MacEachren 1985). Still, experts agree on what makes a district compact: “the sum of distances from any location to the center of a district is minimized” (Fan et al. 2015, 737). A consensus definition has eluded the fourth goal, however, as what exactly constitutes a community of interest has remained very nebulous (Cain, Mac Donald, and

McDonald 2005; Courtney 2008; Medew 2008). Whatever the definition may be, the objective with this criterion is to respect communities of interest by ensuring as best as possible that district boundaries keep together—rather than split apart—these entities.

This ambiguity in the concept of the community of interest has not prevented certain districting stipulations from attempting a description. A few examples deserve mention. Australian law defines a community of interest in sociological terms by referring to “economic, social, and regional interests,” as well as accessibility of communication and travel (Medew 2008, 103). The state of Colorado makes mention of “ethnic, cultural, economic, trade area, geographic, and demographic factors” (Cain, Mac Donald, and McDonald 2005, 18). Alaska attempted to elucidate the concept by interviewing scores of residents on their commonalities, thus taking a more ethnological approach (ibid). One particularly detailed example of a definition for a community of interest comes from the California Constitution’s list of criteria for the state’s independent redistricting commission:

A community of interest is a contiguous population which shares common social and economic interests that should be included within a single district for purposes of its effective and fair representation. Examples of such shared interests are those common to an urban area, a rural area, an industrial area, or an agricultural area, and those common to areas in which the people share similar living standards, use the same transportation facilities, have similar work opportunities, or have access to the same media of communication relevant to the election process. Communities of interest shall not include relationships with political parties, incumbents, or political candidates. (California State Constitution, Article XXI, Section 2-d-4)

While all of the above definitions tend to focus on thematic attributes that come from *outward* observations of the people making up these communities, there may be another way to understand the concept that examines cognitive attributes coming from *inward* observations of those same people. Montello (2003) discussed this thematic versus cognitive distinction in the context of regions, describing the former as being “formed by the measurement and mapping of one or more observable content variables or themes” and the latter as being “produced by people’s informal perceptions and conceptions” (177). Some scholars have suggested that a human cognitive element should come into play when considering where to draw district boundaries. When describing the principles many believed should guide boundary reorganization of local government areas in England, Prescott (1965) referenced one stating that “the boundary should be drawn to cater for local sentiment and regional patriotism” (173). Grofman (1993) introduced an idea that he called the “cognizability principle,” which refers to the ability of residents to cognize their district by being aware of the general configuration of the boundaries, thereby facilitating their “identification of and with the district” (1263). These two principles are not used to define communities of interest per se. However, since they are cited as guidelines for (re)districting and involve individuals’ impressions about and attachments to their local community, they serve as potentially informative ways to understand what gives a community of interest meaning. They also raise the interesting theoretical question of how well cognitively defined communities of interest will correspond to thematically defined ones.

What rationale lies behind taking communities of interest into account? Handley (2008) explained how many authorities and citizens believe that “electoral districts should be cohesive units with common interests related to representation” so as to make the

representative's job easier (275). Yet without a better, clearer definition for a community of interest, this aim may go unfulfilled. Much of the public may instead find themselves living in districts where their interests will not be sufficiently represented, and these interests may go unattended as result. The quality of representation in the context of defining communities of interest may be further informed by considering the descriptive and substantive senses of the concept discussed above. For instance, the thematic definition can serve to ensure that constituents will be well represented in the descriptive sense, as the representative is more likely to be a member of the community of interest if all its thematic attributes are kept together within a given district. On the other hand, the cognitive definition can serve to make sure that constituents will be well represented in the substantive sense, as the representative should have a clear idea of what his or her constituents care about most and thereby effectively advocate for them if most people in the district share a common sense of identity and belonging. For these reasons among others, bringing clarity to the vague idea of communities of interest stands to benefit representative democracy in important ways.

Political science scholars differ as to whether the descriptive or the substantive senses of representation deserve the most focus and emphasis. Pitkin (1967) came down on the side that the substantive sense is paramount over all the others, maintaining that the heart of representation lies in "acting in the interest of the represented, in a manner responsive to them" (209). Phillips (1995), however, contended that the descriptive meaning has an important role to play and that this "politics of presence" must not be isolated from the "politics of ideas" that pertains more to the substantive sense. She makes a strong case that neglecting the descriptive "politics of presence" will tend to result in the political exclusion of women and racial and ethnic minorities as little to no efforts will be made to increase

their presence in representative bodies. Given that the whole purpose behind the move to district elections in Santa Barbara lay in the plaintiffs' desire to fight the political exclusion of Hispanics, and that they explicitly referred to the lack of Hispanic city council members, it seems most appropriate to side with Phillips and consider both the descriptive and substantive senses of the definition of representation in this project.

Therefore, this research will seek to investigate whether the Santa Barbara city council districts deliver quality representation in both of these two senses by empirically assessing both the thematic (from the descriptive) and cognitive (from the substantive) understandings of communities of interest. In addition, the study will explore the feasibility of two approaches of assessing the cognitive meaning in the context of regions: demarcating one's own cognitive region versus rating the cognitive attachment to an existing administrative region (Montello 2003). While plenty of efforts have been made to define communities of interest in the former sense, to date no one has attempted to examine them in the latter sense, let alone both of them together. This study looks to provide a unique contribution to the existing literature by taking both senses into account. It may turn out that surprisingly disparate results come out of these two analyses.

III. Overview of the Research

This research seeks to delineate communities of interest by coupling thematic data with cognitive surveys in a quantitative study within the city of Santa Barbara, California. Demographic and land-use data are linked together as thematic because of their association in the California Constitution's definition of a community of interest. First I take Census data and search within the city limits for clusters of a number of key variables, such as percent who are Hispanic and percent who earn a certain income, in order to identify

communities based on those variables; I then cross that with land-use data. Then I survey residents to obtain their perspective on the extent of their own community of interest as well as on how well the existing boundaries reflect that perspective. Finally, I compare the thematically- and cognitively-derived results, analyzing how well the two sets of results line up with one another and how they match up with the existing city council districts (Figure 2). Whether the thematic or cognitive characterizations of communities of interest appear to better reflect these districts may indicate that one serves as a better definition than the other. Correspondence between the communities of interest derived from this study and the districts may indicate how well citizens are being represented, and so whether the switch to district elections served its objective.

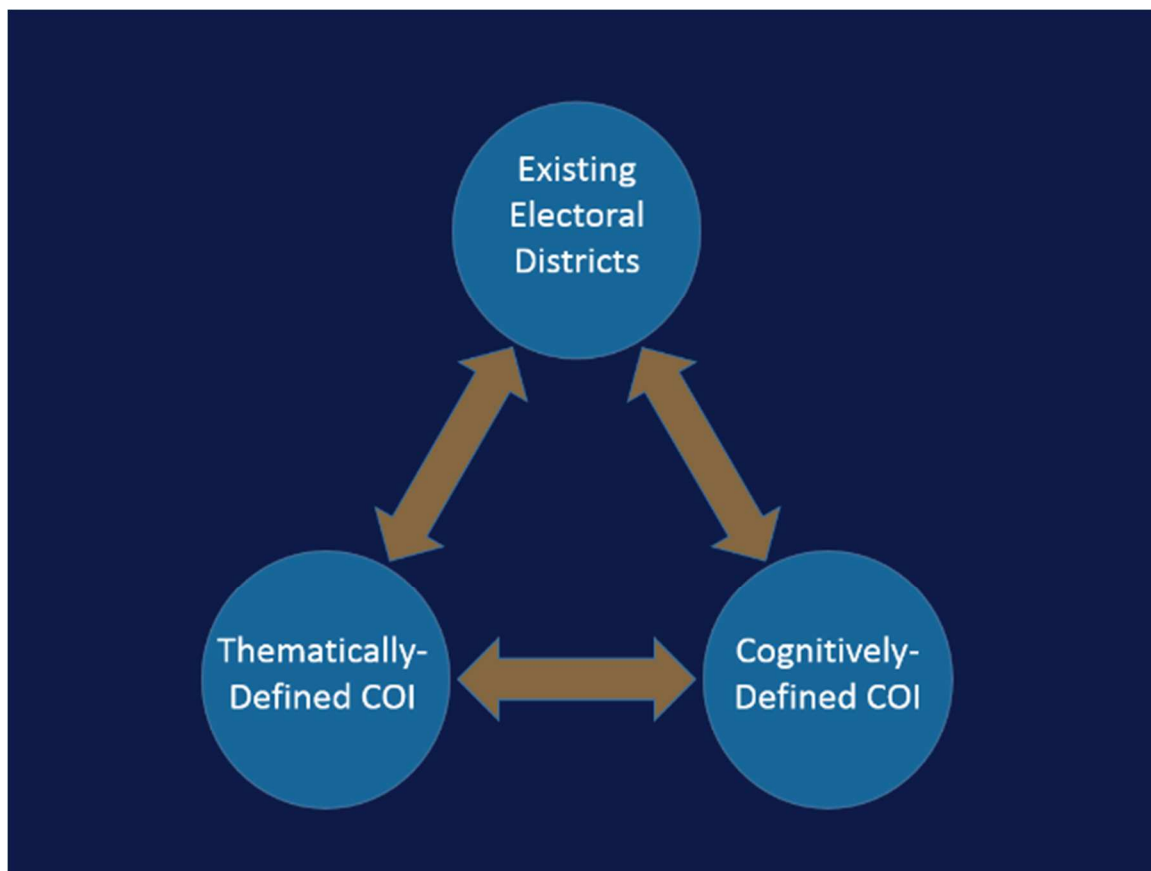


Figure 2. Conceptual diagram of the aims of this research, to investigate the relationships between the three types of regions.

IV. Methods

A. Cases

The cognitive analysis employed primary data that I collected from surveying residents of the first, second, and third city council districts of Santa Barbara (henceforth referred to as Districts 1, 2, and 3). Of course it would have been better to survey across all six districts in the city, but given the limited time and effort I had available to expend I opted to study just those three. Districts 1 and 3—also known as the “Eastside” and “Westside” districts, respectively—are the two majority-Hispanic districts that were the main intended outcome of the effort to initiate district elections in the city in the first place, with a 69% proportion of Hispanics in each district (NDC 2015a). These two districts also have attributes beside that of ethnicity that set them apart from the others and make them worthy of focus, for example their relatively high population densities of 2,997.0 and 6,481.4 people per square kilometer, respectively. In contrast, District 2—also known as the “Mesa district”—features many aspects that differ greatly from those of the majority-Hispanic districts, especially its high share of non-Hispanic Whites at 73% and lower population density of 1,468.9; it therefore serves as a nice counterbalance to the other two. Moreover, these districts are the ones in which city council elections were about to take place at the time of survey collection, meaning that their residents were potentially more open to participating since the questions asked held more immediate relevance for them.

Geographical cluster sampling was used to select houses to approach for an interview at regular intervals throughout each district. This was done to ensure that responses were drawn from residents across the entire district, in the hopes of removing any biases particular to certain neighborhoods. Each district was tessellated into a grid with 10

rows and 10 columns, with the centroid for each grid cell marked. Due to size and shape differences between the three districts, there were 55 such centroids within the boundaries of District 2, but only 45 and 35 in Districts 1 and 3, respectively. To compensate for the relative lack of available center points in Districts 1 and 3, I produced a series of randomly generated points across each of those districts, 6 for District 1 and 16 for District 3. This resulted in a total number of points—both centrally located and randomly generated—of 51 for both of these districts, much closer to the 55 center points in District 2 (Figure 3). I did not expect to collect a response for each one of these points, however, as I realized that many of these cells would only contain a handful of residences if any; some would mostly cover open space or some other non-residential land use. Therefore, I intentionally generated more points than the amount of responses I aimed to collect, which was 35 from each district for a total of 105. Overall 275 residences were approached for a response for this survey; of those, 188 had someone come to the door and 114 agreed to participate, for a total response rate of 60.6%. More detailed information is given in Table 1.

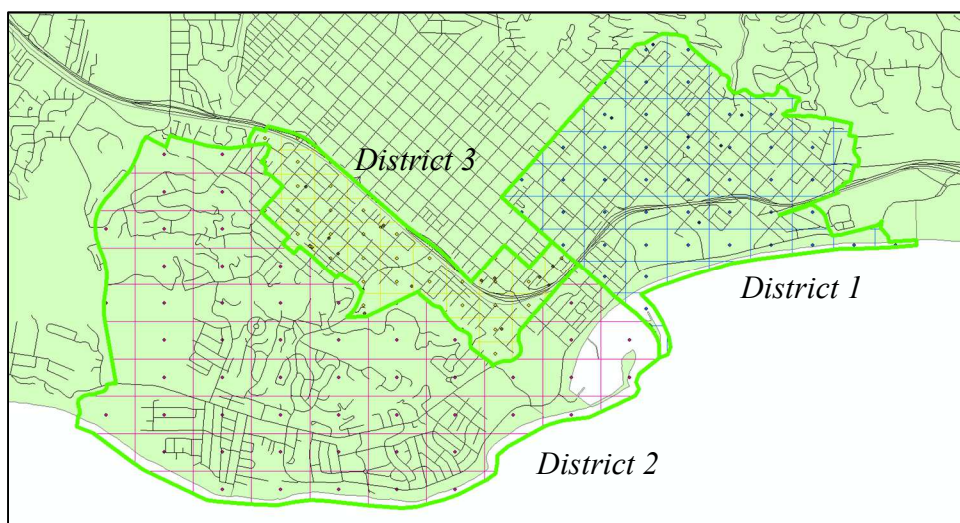


Figure 3. Sampling grids for the three districts surveyed (district boundaries in green), with a center point for each cell and randomly-generated points scattered in Districts 1 and 3.

Table 1. Demographic Summary of Survey Respondents (Compared to 2010 Census and 2009-2013 ACS Data; NDC 2015a).

	Potential Respondents	Participating Respondents	Age in Years	Years in SB	Race/Ethnicity	Sex
District 1	60 (14,865)	35	46.5 (58% 20-60, 16% 60+)	26.9	51.4% Hispanic, 40.0% NH White, 8.6% Others (69% Hispanic, 25% NH White)	60.0% Female
District 2	68 (14,924)	40	51.5 (57% 20-60, 20% 60+)	28.8	90.0% NH White, 10.0% Hispanic (73% NH White, 18% Hispanic)	52.5% Female
District 3	60 (14,324)	39	45.0 (61% 20-60, 10% 60+)	21.1	56.4% NH White, 41.0% Hispanic, 2.6% Others (69% Hispanic, 26% NH White)	59.0% Female
Total/Average	188 (44,113)	114	47.8 (59% 20-60, 15% 60+)	25.6	63.2% NH White, 33.3% Hispanic, 3.5% Others (52% Hispanic, 41% NH White)	57.0% Female

B. Materials

1. Thematic Communities of Interest

The thematic analysis required the acquisition of both demographic and land-use data for those areas within the Santa Barbara city limits. Obtaining the former presented a special challenge because the US Census American Community Survey (ACS), which provides the most comprehensive and detailed demographic survey data for communities across the nation, has some crucial limitations. First among these is the fact that, for privacy concerns, ACS data is aggregated into regions of varying size at two levels of Census

analysis: first the block group which is then nested into the tract. This brings about the issue of the modifiable areal unit problem (MAUP), the fact that areal units are “arbitrary, modifiable, and subject to the whims and fancies of whoever is doing, or did, the aggregating” (Openshaw 1983, 3). Had the boundaries of Census tracts or block groups been drawn differently, the results from those units would have been slightly if not substantially different, yet there is no way around this since this is only way the data are made available. The second critical limitation with the ACS is the fact that it is not a census but a sample, which means that a fair amount of uncertainty about the data is unavoidable. In some areas the level of uncertainty for the estimate of a certain variable, as measured by the margin of error, can exceed the entire estimate. There are methods to deal with this uncertainty, however. One such technique proposed by Spielman and Singleton (2015a) looks especially promising. They advocated taking a multivariate approach in which a large number of different variables for a given area are evaluated together in order to classify that area into a certain group. The benefit in doing what they termed a *geodemographic classification* is that the errors for these largely independent variable-specific estimates tend to cancel each other out, thereby mitigating the overall uncertainty effect. Nevertheless, the uncertainty involved is still substantial enough that this method only works well at the Census tract level, meaning that these units are the smallest level of resolution in which to reasonably conduct this type of analysis.

Spielman and Singleton (2015a) took on the task of classifying each tract in the conterminous United States based on a geodemographic analysis of 136 variable measures taken by the ACS. A full list of these measures is available at their online Github page (Spielman and Singleton 2015b). They selected these measures by conceiving three broad

concepts which they believe differentiate Census tracts: population, environment, and economy. These concepts were then broken down into the domains of age, race, education, family structure, and language for population; moving stability, housing, and population density for environment; and commuting, occupation, and wealth for economy. The individual measures were drawn from each of these domains. Though they admitted this selection task was a somewhat arbitrary exercise, they provided a persuasive rationale for why each measure captured a valuable and unique aspect of what distinguishes one tract from another. They aimed to “select variables that measured each domain, taking into account practical considerations such as coverage, margins of error, redundancy with other variables in the model, and balancing of domains” (2015a, 1010). After selecting the variable measures, they standardized them, input them into a *k*-means algorithm, and fed the output of that algorithm into a Ward’s hierarchical cluster analysis, resulting in a dendrogram showing 250 classes of tracts and the relations between them.

In order to simplify the classification scheme, Spielman and Singleton (2015a) used a statistic called the average silhouette width to determine the best places to partition the dendrogram, which turned out to be at the 10- and 55-class levels. This means that a tract can be justifiably classified into one of 10 categories or into one of 55 categories nested within those 10. Having classified every Census tract according to this procedure, they were thereby able to produce a shapefile of tracts nationwide categorized at both levels, which they made publicly available on their Github (2015b). They also tested the validity of the typology they created by assessing the differentness of their classes and whether such differences were meaningful. By calculating a Gini index for FEC data on campaign contributions and data on crime in Chicago and comparing the index scores across classes,

they were able to show that substantial and meaningful differences did indeed exist, thus validating their typology.

With such a comprehensive, detailed, and valid demographic dataset already completed and freely usable, it made sense to incorporate Spielman and Singleton's results into my study as the demographic component of my analysis. That does not mean that no problems or limitations remain, however. For one, the Census tract is not an ideal unit of analysis for the scale of my study of Santa Barbara city council districts. Tracts are rather large units encompassing entire neighborhoods, so that just two or three of these could approximate the size of an individual district. A particular tract might overlap substantially with a given district but still extend into distant outlying areas, thus detracting from a comparison made between the two regions. Nevertheless, the high uncertainty values associated with the block group level makes that level too suspect to draw any conclusions from it, making the tract the best one can do. Furthermore, the tract boundaries happen to deviate only somewhat from those of the city council districts, thus mollifying the threat to the comparison between the two. Lastly, overlapping the Census tracts with land-use areas will yield a resulting series of smaller intersecting polygons with a resolution more amenable to clustering into a larger polygon that can reasonably be compared with the city council district. Together these considerations make the tract much more appealing than the block group as a unit of analysis for my study.

One other important concern is the fact that the Spielman and Singleton selected their variables for the purpose of differentiating neighborhoods, not to identify communities of interest. Yet the two concepts do not differ very much. The fact that Spielman and Singleton used "community" as a synonym for "neighborhood" several times reveals the

near-interchangeability between the two. A neighborhood is usually considered to be a small-scale community. Though larger communities may exist, they are not a consideration for this intra-city analysis, so there is no need to distinguish the concepts much. When narrowing the focus to the community of interest, the definition provided by the California Constitution offers some good guidelines for the types of commonalities that thematically define a community of interest, as it explicitly references common living standards, means of transportation, work opportunities, and means of communication.

The first three of these four mentions are covered by the Spielman and Singleton data. Common living standards are well captured by the wealth domain, which includes measures such as income, public assistance, and vehicle ownership. Common means of transportation are represented by the commuting domain, featuring measures like duration of commuting and whether people do so by public transport. Common work opportunities are indicated by the occupation domain where measures of types of employment are found, as well as the education domain, since one's level of education highly predicts for what employment one qualifies. The final example given by the Constitution is common means of communication; however, because everyone in Santa Barbara is served by the same television stations, radio stations, and newspapers, this is not useful for the small scale of this analysis. Though the Constitution does not reference the domains of age, race, family structure, language, and moving stability, they nonetheless merit consideration since those attributes coincide highly with the "common social and economic interests" referred to in its definition. Finally, the environmental domains of housing and population density are alluded to in the definition's remarks on land use, as these elements help dictate whether an area is more urban or suburban. In sum, the measures used by Spielman and Singleton are

more than adequate demographic indicators for communities of interest within Santa Barbara to use in my analysis of thematic communities of interest.

Having confirmed the suitability of Spielman and Singleton's demographic data, I uploaded the shapefile they created into a GIS and extracted those Census tracts completely or partially within Santa Barbara, 26 in all. On top of this, I included land-use data because the California Constitution cites common types of land use, such as industrial or agricultural, as a marker of a community of interest. I therefore obtained a dataset provided by the city government website showing the Santa Barbara general plan and uploaded it into the same GIS. This general plan apportions all of the parcels in the city into certain classes, such as high-density residential and industrial (City of Santa Barbara 2013). Obviously the parcel is a much smaller scale of analysis than the Census tract, but the city aggregates these parcels into 25 land-use classes that can be further grouped into larger units that more closely approximate the scale of the tracts and city council districts.

2. Cognitive Communities of Interest

The cognitive survey was administered both orally and in pencil. The oral part involved asking the following open-ended question: "What criteria do you think are important in defining a community?" (I refrained from using the term "community of interest," which I assumed is less well understood). This question was asked to learn about people's understanding of what defines a community. The pencil part involved participants drawing and filling in bubbles on a double-sided sheet of paper, the instructions of which were redundantly printed in both English and Spanish. The front side of the sheet featured a plain street map of the city with major streets labeled (Figure 4). Two items were typed out below this map. The first was the following cognitive definition of a community of interest:

“A community of interest is a group of people who live next to each other and share a common identity and sense of belonging. For the purposes of this study, consider a community of interest to be about the size of a city district or large neighborhood” (I included this instruction to prompt respondents to draw communities of interest to approximate the size of the administrative and thematic regions). The second was the following request: “Please draw on the map above the boundaries of the area containing what you believe to be your community of interest within Santa Barbara.” On the back side of the sheet was the same street map as the front side but with superimposed city council district boundaries (Figure 5). Below that map was the following prompt: “The map shows the districts that the city of Santa Barbara has recently created for city council districts. How well do the boundaries of the district in which you live reflect what you believe to be the boundaries of your community of interest? Please mark one bubble.” The bubbles represented options on a five-point scale from “Very well” to “Very poorly.”

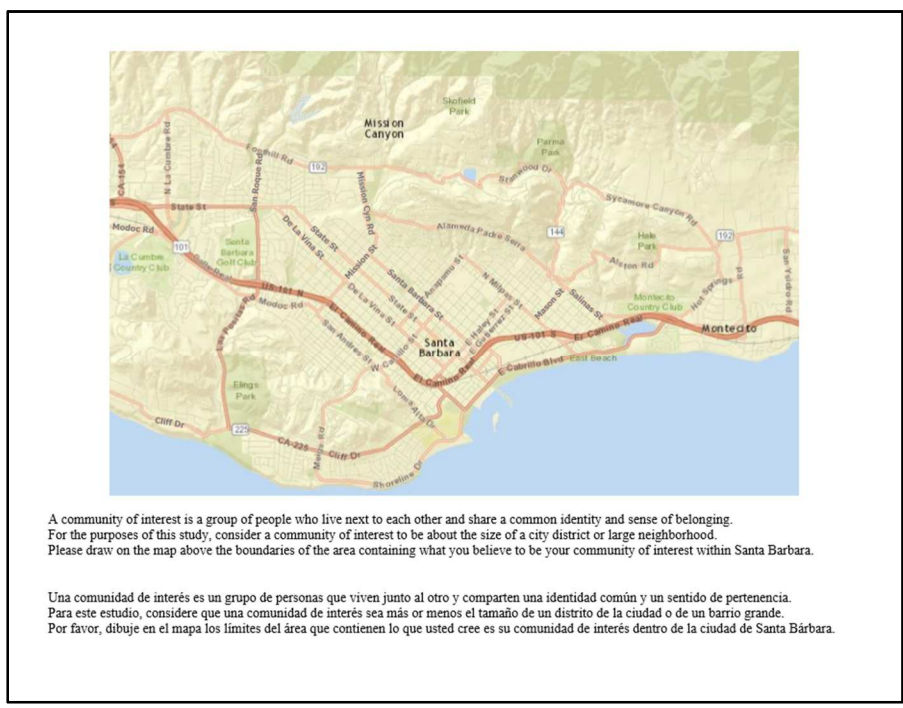


Figure 4. The front side of the sheet of paper given to survey respondents.

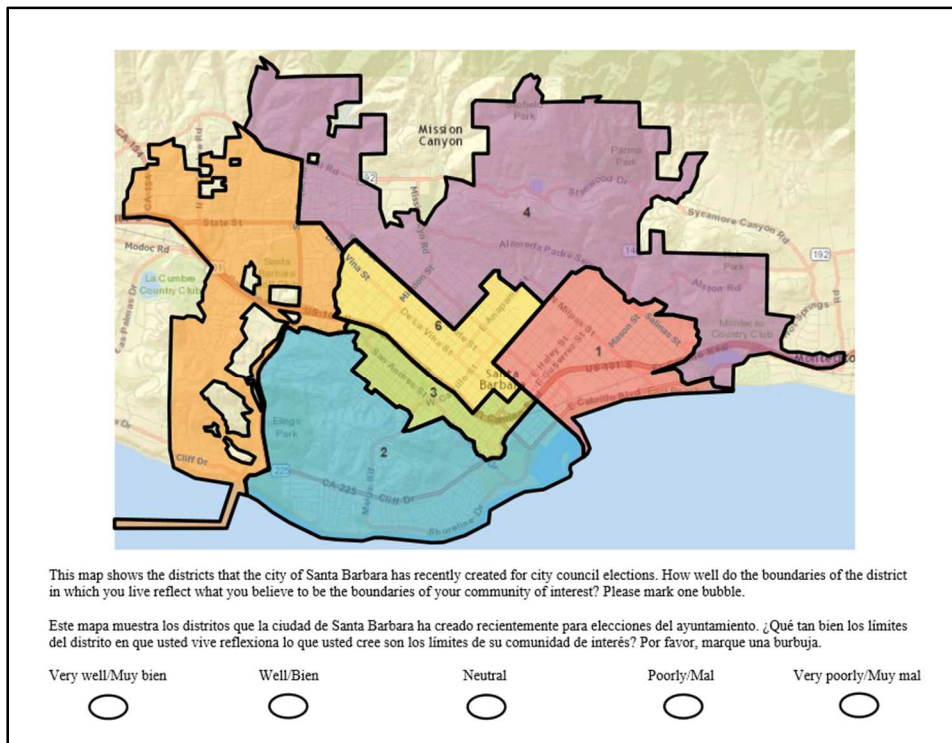


Figure 5. The back side of the sheet of paper given to survey respondents.

C. Procedure

The cognitive survey was conducted using the sampling grid that I devised, so that each cell which contained at least some residences was visited. Since many native Spanish speakers live in Santa Barbara, I enlisted the help of a research assistant who was able to communicate in that language. We first approached the house that we judged to be nearest to the sampling point in the given cell and sought a response from there. If the resident did not appear to be home or refused to take part in the survey, we moved on to the house to the right. This process continued around the street block until we obtained a response, provided the block did not extend into another cell. If a block was exhausted, we would then find the house across the street from the first one and repeat the process for that block. Once someone agreed to participate, we checked off the cell and moved on to the next one. If we could find no one in the cell to participate, we crossed it off as unsuccessful and moved on.

The survey was carried out over a span of six weeks from August 18th to September 22nd, 2015, three days per week. This period proved to be the most convenient time for my research assistant and me, and provided an ample amount of daylight. We did most of our survey collection in a two- to three-hour time block during the late afternoon and early evening hours of weeknights, with the exception of a midday Saturday outing. These times seemed to be the best opportunities to catch residents while they were at home but not yet gathered for dinner. The two of us spent the first two weeks going through District 1 together so that he could help me if I encountered a Spanish speaker, and I could show him how to conduct the survey when the time came for him to go out on his own. After we finished surveying District 1 we parted ways and surveyed Districts 2 and 3 separately but simultaneously; he covered the latter due to its high proportion of Spanish speakers, while I did the former. By surveying both of these districts at the same time we aimed to remove date as a confounding factor for differences observed between the two. Unfortunately we could not do this for District 1, but the fact that Districts 1 and 3 are very similar in composition hopefully mitigates this shortcoming.

For each residence that one of us visited where someone was home, we introduced ourselves to an adult resident as a UCSB student conducting research and asked whether he or she would be willing to participate in a survey about Santa Barbara. When a given resident agreed to participate, we first asked the open-ended question about their definition of a community, recording their answer with an application on our smartphones. We then presented the respondent with the plain street map, on which he or she drew the boundaries of his or her cognitive community of interest. After that we requested that the participant examine the city council districts shown on the other street map and rate the degree to which

the district in which he or she lives reflects what they believe to be their community of interest, using the five-point scale described above. We finished by asking respondents their age and how long they had been living in Santa Barbara; we noted their address and race or ethnicity as we left the house.

D. Analysis

1. Thematic Communities of Interest

The main task of the thematic analysis consisted of grouping both the Census tracts from Spielman and Singleton (2015) and the land-use classes from the city into meaningful clusters, in order to compare those clusters to the existing districts as well as areas of cognitive agreement. Grouping the Census tracts was done easily enough because Spielman and Singleton had already developed the classification scheme at the 10- and 55-class levels, but only one of these two were needed for my analysis. In the 10-class scheme, 7 classes are represented by tracts in the city, while in the 55-class scheme, 8 are represented. These do not differ much, but the latter makes some crucial distinctions between certain neighborhoods that the former does not, which makes it more useful in this context. This is most exemplified by the fact that the 55-class scheme differentiates the Riviera and East Side neighborhoods, which is appropriate given that these areas are widely viewed as distinct from each other, as evidenced by the city's own neighborhood designations (City of Santa Barbara 2016b). Spielman and Singleton developed names for the members of the 10-class scheme but not those of the 55-class one, so fresh identifiers are needed in order to use the latter. However, only one 10-level class in Santa Barbara splits at the 55 level, the class named "Wealthy Nuclear Families." Since this class branches into two sub-classes, one situated at or around the downtown area and the other located in the outer suburbs, it makes

sense to refer to these as “Wealthy Nuclear Families – Inner City” and “Wealthy Nuclear Families – Outer City,” respectively. The remaining classes can be referred to with just their 10-level designations, without any hyphens.

Having settled on the number of classes to use, tracts of the same class could then be grouped into clusters based on their contiguity. For example, if a tract of a certain class shares a border with a tract of the same class, they form a cluster of that particular class. One exception to this procedure bears mentioning though. I did not group tracts of the same class into a single cluster if their common boundary follows the 101 Freeway. This is due to the fact that the freeway divides these areas as a major edge feature (Lynch 1960). They can thus be viewed as separate, though similar, demographic clusters. The end result is a total of 13 clusters of Census tracts in Santa Barbara, representing 8 classes. There are 3 clusters of “Old Wealthy Whites” (OWW), 3 of “White Nuclear Family – Outer City” (WNF-OC), 2 of “White Nuclear Family – Inner City” (WNF-IC), and 1 each of “Hispanic and Kids” (H&K), “Low Income Diverse” (LID), “Middle Income Single Family Households” (MISFH), “Residential Institutions and Young People” (RI&YP), and “Wealthy Urbanites” (WU). Only 8 of these clusters overlap with the districts under study, however (Figure 6).




Figure 6. Map of demographic clusters in Santa Barbara (all of which are labeled), with each differentiated by hue (city limits in red; boundaries of Districts 1, 2, and 3 in green).

Grouping the 25 land-use classes outlined by the city required a categorization scheme that fit the community of interest context. The city itself groups these classes into the broad categories of open space, hillside, suburban, general urban, and institutional and related land uses. These categories did not appear to reflect the land-use distinctions made in the Constitution’s definition of a community of interest, however. For example, the definition explicitly refers to industrial land use, but this class is swallowed up by the city’s general urban category. Furthermore, a hillside category makes little sense because the same type of land use can exist on both flatland and highland; the terrain does not necessarily dictate what can be done on that land. Therefore, it seemed prudent to make new categories for the city’s land-use classes: low density residential, medium density commercial and residential, high density commercial and residential, industrial/harbor, and open space/other

uses (there is no land that is low density commercial). This scheme keeps industrial land use as a distinct type while dividing the rest of the urban land use, both residential and commercial, into density groups that reflect the city's own emphasis on density (almost all of its land-use classes are described by their density level). The final step was to group these land-use categories into clusters based on contiguity. This yielded 11 clusters of low-density (LD) land use, 21 of medium-density (MD), 19 of high-density (HD), 3 of industrial/harbor (IH), and 58 of open space/other uses (mostly small parks and schools) (O). A map of these land-use clusters is given in Figure 7.

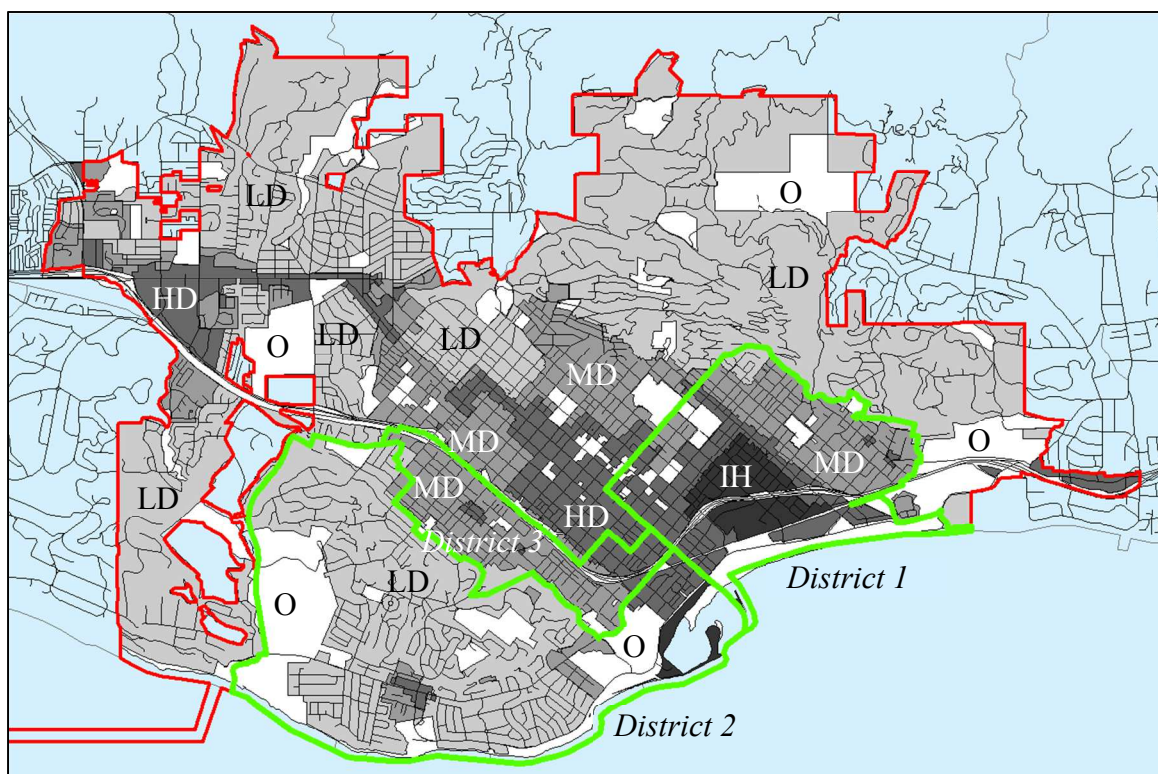


Figure 7. Map of land-use clusters in Santa Barbara (some of which are labeled), with each differentiated by grayscale lightness.

Having clustered the demographic and land-use classes, the two types of clusters could then be overlaid to produce comprehensive thematic clusters reflecting both aspects. The process of intersecting the overlapping cluster types yielded 177 new thematic clusters,

a series of smaller units of finer resolution with which to compare the existing districts and areas of cognitive agreement. These new clusters range from 1 that is a combination of “Wealthy Urbanites” and high-density land use (WU/HD), to 11 that are a cross between “Old Wealthy Whites” and medium-density (OWW/MD), to 28 that cross “Wealthy Nuclear Families – Outer City” with open space and other uses (WNF-OC/O). Figure 8 gives a map of all the thematic clusters, a selection of which are identified.

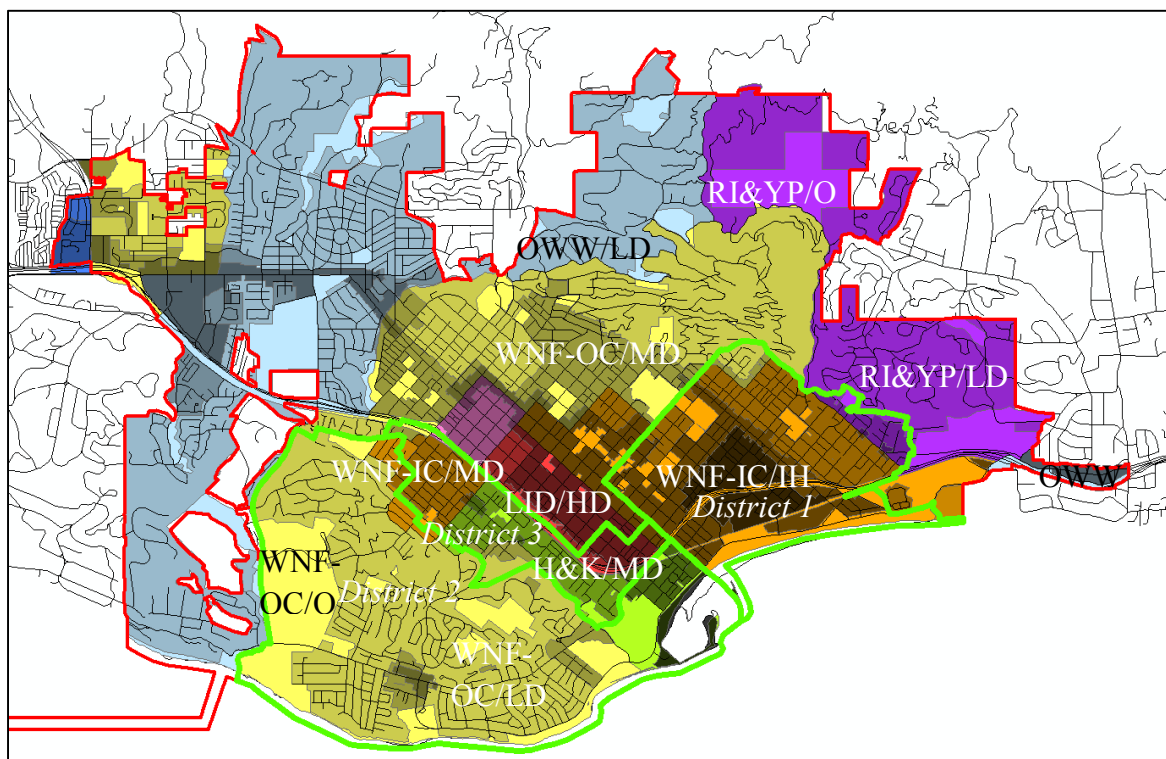


Figure 8. Map of thematic clusters in Santa Barbara (some of which are labelled), with hue indicating its parent demographic cluster and lightness its parent land-use cluster.

With a new series of comprehensive thematic clusters now in hand, I could then determine a thematic community of interest to associate with each district for the purpose of comparison. In order to decide whether a given cluster should be grouped into the thematic community linked with a certain district, three simple rules were followed (in addition to the one already mentioned about the freeway serving as an edge feature). First, if the majority of

a cluster's area fell within the boundaries of a district, that entire cluster was grouped into the community. This rule ensured that a particular thematic cluster was kept whole, just as the district to which it would be compared is retained whole. Second, I considered land-use clusters that were completely contained by open space/other land-use clusters to be a part of the latter. And third, I considered each and every open space/other land-use cluster to be part of the immediately surrounding cluster unless it was land on which the freeway was built. These last two rules enabled the grouping of open space/other land-use clusters with the urban clusters that make the most use out of them, for example a park with the residential neighborhood around it. Following this procedure resulted in a contiguous thematic community of interest identified with each district.

2. Cognitive Communities of Interest

The cognitive analysis involved three exercises: coding the open-ended responses for the community definition, digitizing the drawn polygons and determining their areas, and calculating degree of agreement among the polygons. When coding responses to the question asking participants to define a community, I looked for common words or phrases given among all the respondents and grouped them into categories. For example, the mention or allusion to interaction among people, including use of the word "together," led to the creation of an "Interaction" category. Once a category was determined, I could then tally up the number of respondents whose definitions included wording that fell into that category. The most popular categories represented the criteria that people most often took into account when considering what community means to them. Altogether, twelve categories were created (Table 6).

Next the boundary lines drawn by respondents were analyzed to determine the area of each polygon they formed. This was done by first scanning all the drawings and then digitizing the lines in a GIS to create a series of overlapping polygons for each district; this allowed for calculating the area of each polygon. In addition to one case thrown out due to an error made when administering the survey, anomalies found in people's drawings led to the exclusion of six more cases. Three individuals chose not to draw any polygons, while two drew so many polygons with so much overlap among them that their drawings were incomprehensible. Finally, one person drew a polygon with an opening on one of its sides that prevented a confident determination of its area. This winnowing left 107 cases for areal analysis. A selection of community polygons drawn by residents is presented in Figure 9.

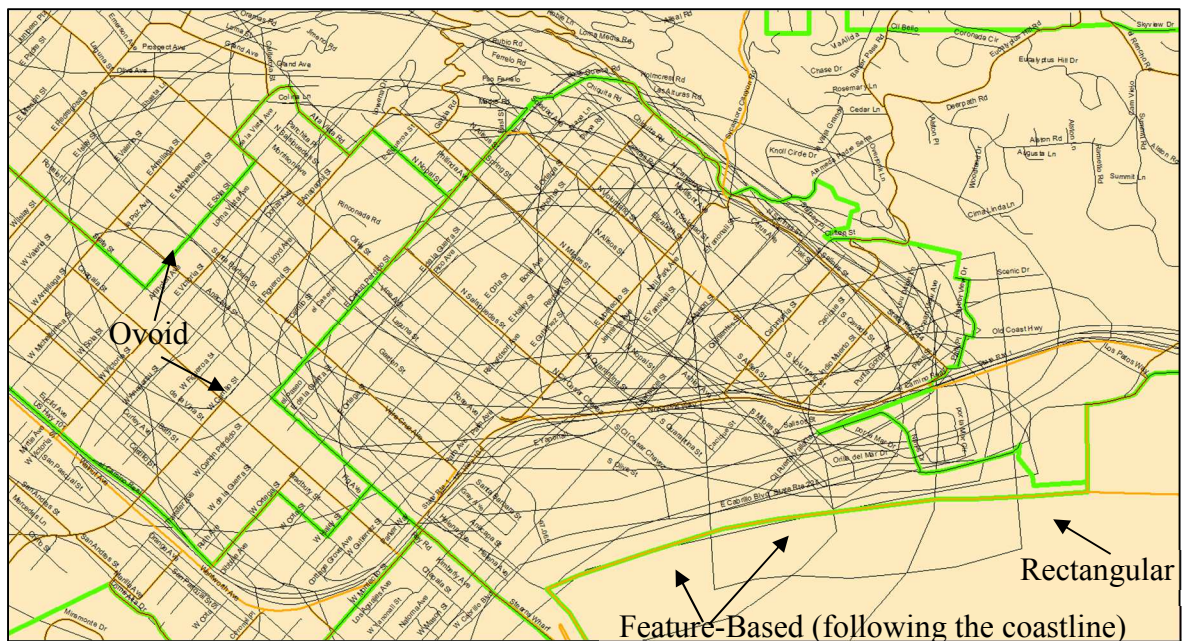


Figure 9. Lines drawn by residents of District 1, with examples of different classes of shape indicated as such.

As a thematic community of interest could be associated with each district based on which thematic cluster overlapped with a given district, so too could a consensus cognitive community of interest be created for each district based on the polygons drawn by their

residents. These communities were determined by the degree to which residents of a given district agreed about the location and extent of their community of interest, so that for each community a gradation from a lesser-agreed-upon periphery to a greater-agreed-upon core would be visible, rather than a monolithic average polygon. To determine level of agreement, all the polygons were merged into one shapefile, which served as the input for two operations. First, I computed a count of the overlapping polygons at each point in space. Second, I used that count to produce an output raster with 25×25 meter cells (deemed to be adequate resolution). This output raster could then be classified based on degree of agreement across points in space. Agreement could range from 0% at points in space contained by no resident's polygon to 100% at points in space contained by all residents' polygons (Woodruff 2012). This process resulted in maps of the cognitive communities of interest salient within each district, with a light yellow to dark brown color scheme showing lesser to greater agreement.

V. Results

A. Thematic Communities of Interest

I first examine and profile the thematically-defined communities of interest created by combining all thematic variables. . One such community of interest was fashioned for each of the three districts under study (Figure 10). The community centered on District 1 is chiefly formed by the WNF-IC demographic cluster, with a slice of RI&YP that is medium density. High and medium density land-use areas make up the bulk of this region, as well as the only concentration of industrial use in the city; parks, beaches, and the freeway fill the rest of the area (Table 2). The thematic community associated with District 2 is largely defined by WNF-IC/LD and WNF-IC/O (mostly parks and open space used by residents of

the first class); together these two classes make up 83.7% of the community's area. Harbor facilities, H&K-HD uses connected to beach tourism, and Santa Barbara Community College occupy most of the remainder (Table 3). Finally, the community of interest connected to District 3 consists largely of medium density uses in the H&K or WNF-IC classes; when combined with the LID/O cluster containing the freeway and adjoining land, these take up 85.3% of the area of the community. Only small commercial clusters and parks or schools remain (Table 4). These three thematic communities of interest represent what the city council districts might look like if boundary drawers only paid attention to a community of interest criterion defined solely by demographic and land use attributes.

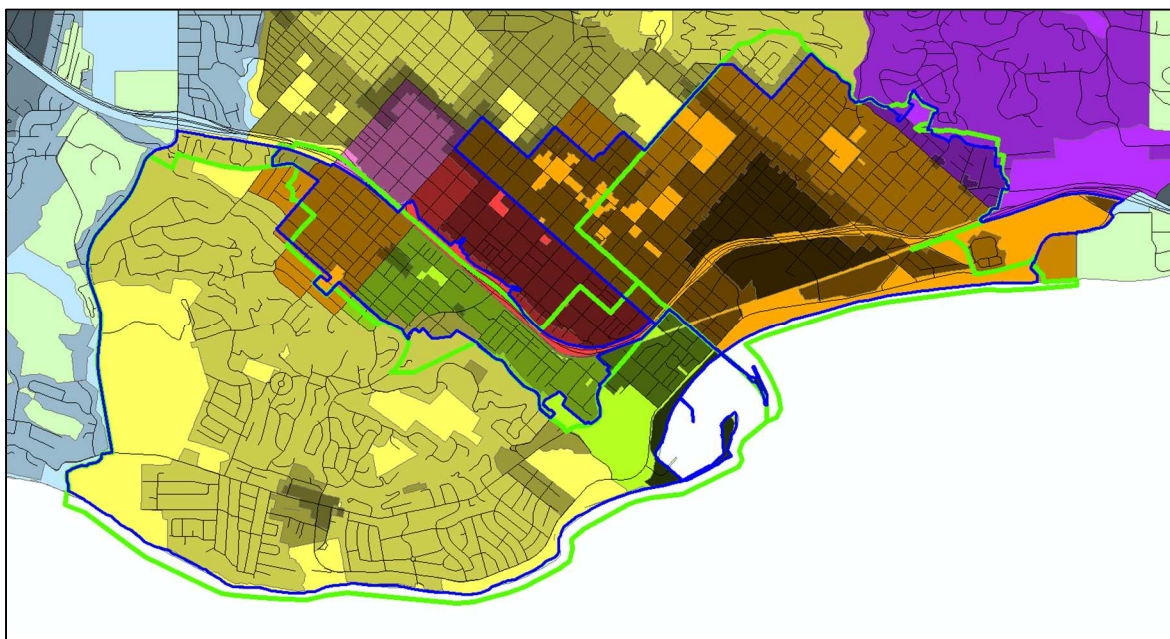


Figure 10. Thematic communities of interest (boundaries in blues) associated with each district.

Table 2. Descriptive Statistics for Thematic Community of Interest Based on District 1

Cluster Class	Number of Clusters	Area (sq km)	Area Percentage
WNF-IC/HD	9	1.95	34.6%
WNF-IC/MD	5	1.38	24.5%
WNF-IC/O	10	1.26	22.4%
WNF-IC/IH	2	0.77	13.6%
RI&YP/MD	1	0.23	4.1%
RI&YP/O	1	0.04	0.6%
RI&YP/HD	2	0.01	0.2%
Total	30	5.65	100.0%

Table 3. Descriptive Statistics for Thematic Community of Interest Based on District 2

Cluster Class	Number of Clusters	Area (sq km)	Area Percentage
WNF-OC/LD	1	5.91	61.5%
WNF-OC/O	7	2.14	22.2%
H&K/O	3	0.39	4.1%
WNF-OC/MD	5	0.36	3.7%
H&K/HD	3	0.31	3.3%
WNF-IC/LD	2	0.26	2.7%
H&K/IH	1	0.13	1.3%
WNF-OC/HD	2	0.07	0.7%
H&K/LD	1	0.04	0.4%
Total	25	9.62	100.0%

Table 4. Descriptive Statistics for Thematic Community of Interest Based on District 3

Cluster Class	Number of Clusters	Area (sq km)	Area Percentage
H&K/MD	1	0.77	44.3%
WNF-IC/MD	1	0.53	30.8%
LID/O	1	0.18	10.2%
WNF-OC/MD	2	0.13	7.5%
H&K/HD	2	0.03	1.7%
H&K/LD	2	0.03	1.6%
WNF-IC/O	1	0.02	1.4%
H&K/O	3	0.02	1.2%
LID/MD	1	0.01	0.5%
WNF-IC/HD	1	0.01	0.5%
LID/HD	1	0.00	0.2%
Total	16	1.73	100.0%

The District 1 thematic community of interest has an area of 5.65 square kilometers, as compared to 4.96 for the district itself. This 0.69 square kilometer deviation owes to the fact that the WNF-IC/HD area takes in the whole downtown section centered just outside the District 1 boundaries to the northwest. The District 2 community totals 9.62 square kilometers, as compared to 10.16 for the district itself. A quick look at the map reveals that this 0.54 square kilometer discrepancy is mostly a consequence of the district boundary extending farther into the sea than the community boundary (originally the tract boundary), especially in the harbor. Lastly, the District 3 community has an area of 1.73 square kilometers, as compared to 2.21 for the district itself. This 0.48 square kilometer difference is due to the exclusion from the community of the small section of the district on the other side of the freeway, which is mostly occupied by a cluster centered in another district.

B. Cognitive Communities of Interest

Next I examine the cognitively defined communities of interest. Drawn polygon areas average 6.94, 13.80, and 7.62 square kilometers for residents of Districts 1, 2, and 3, respectively. That these figures are well above the areas of each electoral district—4.96, 10.16, and 2.21 square kilometers, respectively—reflects the fact that some of the polygons had areas far exceeding that of the district. In spite of the instructions to “consider a community of interest to be about the size of a city district or large neighborhood,” a few people insisted on drawing a polygon encompassing almost the entire map area, which pulled the averages upward. This is borne out by the huge standard deviations of polygon areas of 13.55, 20.51, and 12.74 square kilometers. Because of the undue influence exerted by these outliers, it seemed informative to examine the polygon areas after excluding them. This was done for polygons more than two standard deviations larger than the mean of each

district, of which there are six (three of them in District 2); that left 101 for analysis. As a result, polygon averages for the three districts fall to 4.77, 9.28, and 5.18 square kilometers, with much smaller standard deviations of 6.15, 13.06, and 7.16. These numbers more closely approximate those of the district areas and offer a clearer picture of how the communities conceived by most participants compare with the boundaries the city drew.

Residents of the three districts do not agree very much about the location and extent of their community of interest; in no district is even a single point in space contained in the polygons of 70% or more of its residents. Within District 1, the area contained in the largest percentage of residents' polygons is done so by 68% of residents; it is 2.1% of the total district area. For District 2, that area is contained by 60% of residents and represents a scant 0.2% of the district area. For District 3, that area is contained by 68% of residents and totals an even smaller 0.1% of the district area. Given these patterns, I depict agreement at three levels: 40%+, 50%+, or 60%+ agreement. For example, an area at the 50% level is contained in the polygons drawn by at least 50% of the respondents. As expected, the areas of agreement within each class decline in size as one moves toward greater agreement (Table 5). For this analysis I included the six outlier polygons I had excluded for the areal analysis; their large sizes had no skewing effect here since only their innermost parts that overlapped with other polygons are taken into account. Thus the outlier-skewed average areas for the drawn polygons are given in the table below to reflect the fact that all 107 are being included.

Table 5. Areas of Individual Drawn Polygons vs. Cognitive Agreement Areas, in sq. km.

	District Area	Average Area of Drawn Polygons	40%+ Agreement Area	50%+ Agreement Area	60%+ Agreement Area
District 1	4.96	6.94	3.04	1.43	0.35
District 2	10.16	13.80	8.91	3.53	0.02
District 3	2.21	7.62	2.52	0.97	0.22

The majority of residents in Districts 1 and 3 agree to the existence of a single community of interest situated almost entirely within the boundaries of their respective districts. In District 1 the 50%+ agreement region spans most of the inhabited part of the district, with a 60%+ core region centered along Milpas Street, commonly considered the main street of the area. The 40%+ region spills northwest into areas outside the district but remains firmly bounded on the south by the 101 Freeway; very few people live on the other side of this prominent edge feature (Figure 11). Taking the centroids of the individual drawn polygons (again excluding the large outliers) reveals that most of them concentrate in the 40%+ agreement region, though a good number are located to the west in the downtown area (perhaps reflecting people’s workplaces), which pulls the mean centroid westward so that it falls to the west of the 60%+ core though still barely remains in the 50%+ region. The standard deviational ellipse of the centroids reflects this westward shift to cover a fair amount of land outside the boundary, much of it in downtown.

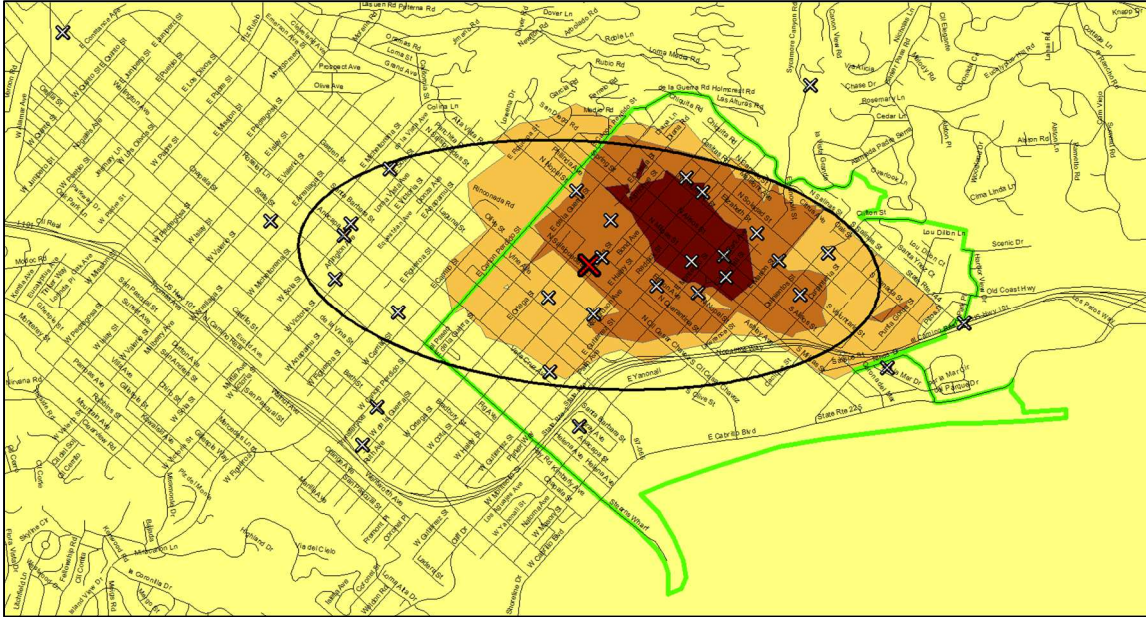


Figure 11. Cognitive agreement for District 1 (green boundaries), with yellow representing the 0-39% class, orange the 40-49% class, brown the 50-59% class, and dark brown the 60%+ class. Also shown are the centroids of the drawn polygons (marked by the black and white X's, the mean centroid (large black and red X), and the standard deviational ellipse of the centroids (black line).

Similar to District 1, in District 3 the 50% agreement region includes much of the inhabited part with the notable exception of the residential areas southeast of West Carrillo Street, and its 60%+ core region stretches along the main artery of San Andres Street. The 40%+ region does take in the areas beyond West Carrillo Street but overflows the district boundaries on both sides (Figure 12). Again there is a pattern of centroids mostly lying within the 40%+ region but several scattered in the downtown area. In this case the result is the mean centroid drifting from the core toward the east, landing right on the edge of the area of majority agreement. As such the standard deviational ellipse also extends eastward to envelop many of these downtown locations. Despite both being somewhat pulled toward downtown, the cognitive communities of interest in Districts 1 and 3 can be viewed as relatively compact, cohesive, and unitary.

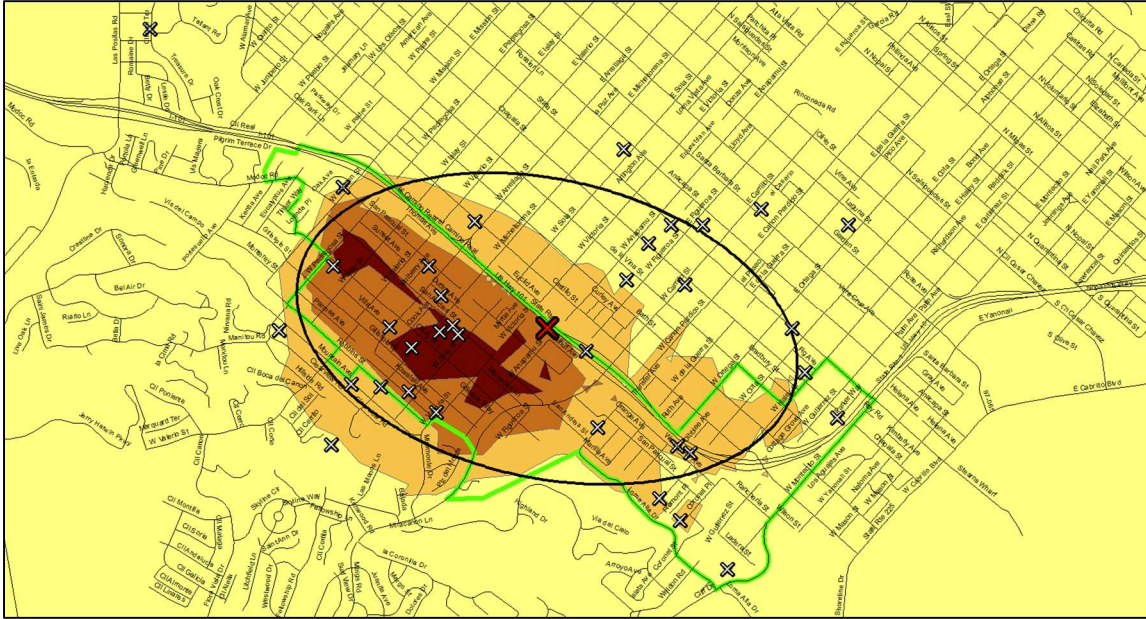


Figure 12. Cognitive agreement and centroids for District 3.

In contrast to Districts 1 and 3, the majority of residents in District 2 do not agree to the existence of a single community of interest largely contained by the district boundaries, but rather acknowledge the presence of two separate communities, one of which extends well outside the district. A southern 50%+ agreement region straddles Cliff Drive and extends from the summit of the Mesa down to the shoreline. The other 50%+ agreement region to the north occupies the Bel Air neighborhood north of West Valerio Street and spills into District 3 to take in a sizeable chunk of that district. This pattern suggests that District 2 residents identify more with subsections of the district rather than the area as a whole, perhaps owing to the large size of the district. Even so, the 40%+ agreement region encompasses almost all of the district (though with considerable spillover into District 3), so at some level there is an idea of a Mesa-wide district (Figure 13). The centroid locations give further credence to this dichotomy between a single large community and two smaller ones. While the mean centroid as well as five individual centroids (each representing a single polygon) are right between these smaller areas, indicating some belief in a single

community encompassing the Mesa, most are found in and around the smaller areas. (Quite a few are well outside the district in the downtown area, which pulls the standard deviational ellipse to the northeast, again possibly reflecting residents' places of work). Therefore it is clear that many residents of District 2 do not believe there is a single community of interest taking up the whole of the district but rather two separate ones occupying adjacent hills.

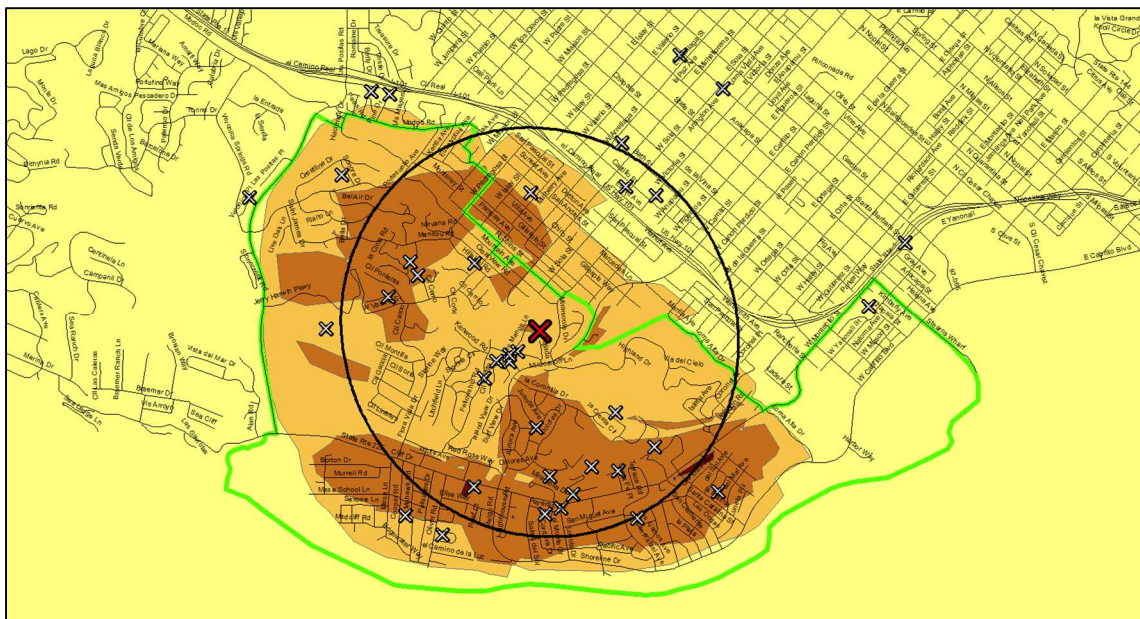


Figure 13. Cognitive agreement and centroids for District 2.

For each survey participant, I assess whether they drew one continuous polygon or more than one, the shape of those polygons, and whether the polygons were open or closed. Almost all respondents drew just one polygon, though three drew two, and three drew from four to six. Polygon shape is classified as either ovoid, rectangular, or feature-based, depending on whether it resembles an oval, it resembles a rectangle, or its edges closely followed features of the environment such as streets or the coastline (Figure 9). The ovoid and rectangular categories are mutually exclusive, but there are three instances in which a shape takes an ovoid form in one place and follows features elsewhere. Overall there are 61 completely or partially ovoid shapes, 47 completely or partially feature-based ones, and

three rectangular ones. Therefore most people did not draw lines that carefully followed particular features but rather made polygons that depicted a more roughly-defined region in their mind. The final aspect of shape I consider is whether the boundary defining a given polygon forms a closed circuit or is open on any side. I then judge whether the open polygons are truly open or intend to imply closed features. For example, several drew a polygon boundary that began at one point along the coast and ended at another point further down the shore; it is therefore likely that they intended the coastline to represent the remainder of the boundary. Of the 13 open polygons drawn, all but one clearly suggest having been intended as a closed one.

I next examine how well respondents felt the city's electoral districts represent their communities. Respondents indicated that the city's electoral districts represent their communities moderately well and nearly equally in all three districts. The average rating of how well the districts matched the communities (1 = "Very poorly" and 5 = "Very well") is 4.0, 3.6, and 3.5, for Districts 1-3 respectively. These are average responses hovering between "Neutral" and "Well." The standard deviation of these ratings is around 1.0 for all districts. The small differences between the districts prove not to be statistically significant ($F[2, 103] = 2.33, p > .05$). Also, residents of different districts drew community polygons whose size, shape, and location matched the city's electoral districts about equally well ($F[2, 103] = 1.79, p > .05$) (I describe the index used to quantify this spatial correspondence below). Factoring in the ethnicity of the respondent (Non-Hispanic or Hispanic) along with their district in a 2×3 MANOVA does not result in a significant interaction between ethnicity and district when it comes to the rating given or the spatial correspondence of the drawing ($F[2, 97] = 0.99, p > .05$). There is no significant main effect for ethnicity on rating

either ($F[3, 103] = 3.44, p > .05$). In short, I find no evidence that residents of one district believe their communities are better or more poorly represented by the city’s districts, nor that the shapes of the polygons they drew come any closer to reflecting that of their district.

Coding responses for the open-ended question about one’s definition of a community results in twelve categories (Table 6). More than three quarters of participants made reference to a group of people of some kind, so this is definitely an important aspect of the definition of a community. A slim majority of respondents also alluded to the interaction among people, which included any sense of “togetherness” or comradery. Explicit references to geography and residence were the other two categories found in over a third of people’s definitions. This coding task allows for the creation of a summary statement for what Santa Barbara residents believe a community to be, with decreasing confidence towards the end of it: *A group of people who interact with each other in close geographic proximity, living and supporting one another for their mutual benefit.* The boundary drawings and district ratings should in theory reflect these concepts.

Table 6. Categories for Phrases Used in Defining a Community (114 Respondents).

Name	Every mention of...	Percent Included
Group of People	A group of people	75.4%
Interaction	The interaction among people, including “together”	52.6%
Geography	Spatial/geographic area/proximity	40.4%
Residence	People residing	35.1%
Support	Giving and receiving of support among people	19.3%
Benefit	Activities benefiting/serving many people/greater good	18.4%
Unity	“Unity” or a related phrase	13.2%
Commonality	Shared traits or commonalities	12.3%
Neighborhood	“Neighborhood” or “neighbors”	10.5%
Economy	Occupations, jobs, or anything related to the economy	8.8%
Diversity	“Diversity” or a related phrase	2.6%
Culture	Cultural characteristics	1.8%

C. Correspondence Between Communities of Interest and Electoral Districts

Using demographic, land use, and survey data to identify two types of regions—thematic and cognitive communities of interest—allows for direct comparison between them as well as between a third type of region: the city council electoral districts. While I have already described how the thematic communities and electoral districts correspond, here I show how all three types of region compare with each other in a more quantitative fashion. I do so by overlaying one region on another in order to determine their overlap. Given these overlaps, I can then examine how similar the regions are to each other using a spatial similarity index that assesses the degree of overlap between the two, which depends on their relative locations, sizes, and (to some degree) shapes. Several such indices have been invented, each with its unique formula for computing spatial similarity (Frontiera, Larson, and Radke 2008). However, a number of these have difficulties with them that make them less attractive for use, such as taking a different form depending on the case or situation. For example, one measure uses one function if a region is completely contained by another and a different function if not. A simple and intuitive index with only one function in all cases is that developed by Hill (1990), which is:

$$\text{Spatial Similarity} = 2 \times O / (Q + D)$$

where Q and D are the areas of the two regions in question and O is the area of their overlap. Hill's index ranges from 0, meaning the regions are not similar at all because they do not overlap at all, to 1, where they are exactly the same location, size, and shape. (Note that the index does not much reflect shape similarity unless the overlap is very high.)

Comparing the thematic and cognitive communities of interest (the latter at the 40%+ agreement level) reveals a relatively high degree of spatial similarity between them

for all three districts, with each of the three pairs overlapping at 0.60 or better. While the similarity measures in Districts 1 and 3 hover around 0.60-0.70, that in District 2 greatly surpasses those at 0.87 (Table 7). The lower values for Districts 1 and 3 chiefly reflect the disparity in size between the thematic communities of interest and their cognitive counterparts; in District 1 the thematic largely engulfs the cognitive, and vice versa in District 3. District 1 residents appear to have a more limited conception of their community of interest relative to their thematic community; they may not be inclined to include those areas where fewer people live, and industrial and tourism use predominates. Residents of District 3 seem to be more willing to include outside thematic areas in their community, perhaps feeling that their thematic community is too narrow in scope, occupying just the land between the Mesa and freeway. In District 2, on the other hand, the two types of communities of interest correspond closely in their size and to a lesser degree their extent.

Table 7. Spatial Similarity Between Thematic and Cognitive COIs (Area in Square km).

Associated District	Area of Thematic (<i>T</i>)	Area of Cognitive (<i>C</i>)	Area of Overlap (<i>O</i>)	Hill 1990 Index: $2 \times O / (T + C)$
District 1	5.64	3.04	2.59	0.60
District 2	9.62	8.91	8.02	0.87
District 3	1.73	2.52	1.43	0.67

As for the correspondence between these communities of interest and the existing city council districts, the results show that the thematic ones are more similar to their associated districts than the cognitive ones are, with the overlap between electoral and thematic communities averaging about 0.15 greater than that between electoral and cognitive communities. Tables 8 and 9 present both comparisons. The spatial similarity ratings are quite high for the thematic communities, ranging from 0.83 to 0.93. In contrast, the cognitive communities reflect their associated districts less well, with a range of 0.65 to

0.82. Here District 2 once again has the highest value, while Districts 1 and 3 lag behind, especially in regard to the cognitive communities.

Table 8. Spatial Similarity Between City Council Districts and Thematic COIs.

District	Area of District (D)	Area of Thematic (T)	Area of Overlap (O)	Hill 1990 Index: $2 \times O / (D + T)$
1	4.96	5.64	4.49	0.85
2	10.16	9.62	9.19	0.93
3	2.21	1.73	1.63	0.83

Table 9. Spatial Similarity Between City Council Districts and Cognitive COIs.

District	Area of District (D)	Area of Cognitive (C)	Area of Overlap (O)	Hill 1990 Index: $2 \times O / (D + C)$
1	4.96	3.04	2.59	0.65
2	10.16	8.91	7.83	0.82
3	2.21	2.52	1.60	0.68

VI. Discussion

This study finds that even on a small, intra-city scale, thematically-defined communities of interest correspond with those that are cognitively-defined reasonably well, especially when one considers how much of the former is uninhabited or sparsely inhabited and thus less amenable to being included in people's cognitive regions. Clearly the thematically-defined communities of interest better reflect the electoral districts that the city created. One likely reason for this phenomenon is that the thematic communities of interest better reflect the comprehensive nature of electoral districts, the fact that districts must comprehensively cover the city. The thematic communities tend to cover more of the city than the cognitive ones in large part because of the different units that make up the two types. The cognitive community of interest in each district is the result of taking dozens of overlapping, mostly ovoid, drawn polygons and identifying an area covered by a certain

proportion of those polygons. In contrast, the thematic type is a product of clustering parcels according to their land use and the demographics of the Census tracts in which they are situated. While the drawn polygons overlap with each other and tend to leave out areas where very few or no people live or work, the parcels do not overlap but comprehensively cover all areas of the city regardless of the degree to which people use them (so many parcels are, and therefore much of the thematic community is, non-residential). In essence the parcels are administrative regions while the drawn polygons are cognitive regions; therefore, it makes sense that the type of community of interest consisting of small administrative regions—and so can itself be considered a large administrative region—corresponds better with the administrative electoral districts (Montello 2003).

This issue is best illustrated by the disparity in spatial similarity between the two types of communities of interest in District 1. While the thematic one compares with the district at 0.85 the cognitive one is only similar to its district at the 0.65 level. This is because the thematic one takes in all of the areas on the southern side of the freeway that are almost entirely devoted to industrial purposes or hotels for tourists, which apparently most district residents consider too far removed from their daily lives to include in their cognitive community of interest. Therefore it seems that thematic communities of interest do a better job of reflecting the extent of the electoral district since they both have comprehensive coverage. Granted, one could repeat this analysis by comparing only the inhabited parts of the district with solely residential parcels, but rarely are communities just defined by their residential components, as the California Constitution makes clear with its references to industrial and even agricultural areas. Cognitive communities of interest, on the other hand, do a better job of marking the spaces of highest agreement: the more people agree on a

given area being in their community of interest, the more justification that area has for being kept within the district boundaries as the core of whatever district is being created.

Another reason why the thematic communities of interest fit the electoral districts better than the cognitive communities of interest could be the different methods used to determine the two types. The cognitive community of interest was equated with a single, very large region of 40%+ agreement; the size of this one region came close to that of Districts 1 and 2 and even surpassed that of District 3. On the other hand, the thematic community of interest was formed by connecting small clusters of demographic and land-use variables in order to produce a larger region, according to whether a given cluster fell mostly within the boundaries of a certain district. The smaller, multifarious nature of these constituent clusters allowed for much more flexibility in crafting a community of interest that closely reflects its associated district, since there was rarely a cluster large enough to extend well beyond the district boundaries, and such clusters were outweighed by many more that were completely within the district. This is opposed to the large, unitary nature of the 40%+ agreement region, which could only be compared as it was. However, the relatively small size of the thematic clusters cannot completely explain the striking similarities observed in this study. Some clusters are actually quite large, like the WNF-OC/LD cluster in District 2 at 5.91 square kilometers, or even the H&K/MD cluster in District 3 at 0.77 square kilometers. Despite their large extent, these clusters still manage to fit remarkably well within district boundaries.

The fact that thematic communities of interest and their associated districts are so similar in location and extent, even though some clusters are still quite large, points to two potential explanations. It could be evidence either that the boundary drawers of the electoral

districts did an excellent job accounting for thematic attributes, or that the cluster boundaries by chance match those of the districts very well. The latter explanation makes the observed similarities an accident of the way the regions were formed. Do the boundaries just happen to adhere well because of a random set of arbitrary choices? One cannot deny that many somewhat if not very arbitrary decisions led to the thematic clusters in this study: by the Census with their tract boundaries, by Spielman and Singleton with their classification scheme, by the city with their density groupings, and by me putting it all together. Even so, the first explanation seems more convincing because a demographics company was hired by the city to draw up the city council districts. Their first priority was to ensure that the districts made sense demographically, especially in regard to the Hispanic concentrations. In sum, the reason for the greater similarity between the districts and thematic communities of interest is likely some combination of the comprehensive, small, plural nature of the latter's constituent units and the fact that the city paid more attention to demographics.

The problem with the cognitive community of interest as assessed in this study is that there is so little agreement about the location and extent of these entities. What explains this relative lack of agreement? Had I asked participants to draw a boundary around particular neighborhoods by name (such as the "Westside" or the "Mesa"), I might have obtained more agreement between them. Studies that asked respondents to define specific named neighborhoods like "Downtown" or "Koreatown" certainly suggest this (Montello et al. 2003; Bae 2015). Those neighborhoods have an identity that is firmly attached to a certain space. People differ somewhat on the details of the boundaries, but everyone concurs that there is a unique, distinct Downtown Santa Barbara centered on State Street and, likewise, a unique, distinct Koreatown straddling Wilshire Boulevard between the 10 and

101 Freeways. In contrast, “community of interest” is a much vaguer concept that greatly depends on two elements for each individual. First, it depends on which segments of one’s activity space a person decides to include. Does one incorporate just the area where he or she lives and interacts with neighbors, or those areas where he or she works and shops, or both together? Second, it depends on how large a person considers a community of interest to be. Does one include just those areas where he or she actually traverses during normal activities, or the surrounding areas as well (and how much of these outer areas)? This issue of the scale of the community of interest will be examined further below. For now, I conclude that cognitive communities of interest are real but rather idiosyncratic and open to much interpretation depending on the instructions designed to elicit them. I thus find relatively little agreement among people on their location, size, and shape.

Another interesting question regarding the cognitive community of interest concerns the question of why so many in District 2 identify more with subsections of the district rather than the whole. Why might people feel a greater sense of belonging with these smaller areas? Actually the majority agreement areas in this district are no smaller than those in the other districts; they are just smaller relative to the size of the district in which they are situated. Given the lower density of District 2 this means that residents conceived of communities of interest of similar size but lower density than those of residents elsewhere. Perhaps people resist the idea of a community of interest that is large in size even if it means incorporating fewer people; they consider space as much or more than they do population. It is difficult to conclude too much given the small number of districts in this study, but it would be an interesting idea to pursue in future research. It could also be that this pattern observed in District 2 is an effect of the terrain in that district, with the valley

separating the two hills acting as an effective cognitive barrier. Further investigation of districts with unique topography might elucidate this question.

Turning now to the way in which the survey was administered, some aspects of instrument design surely affected the results that the survey yielded. One choice I made with potential consequences was to provide respondents with a street map of the city with major streets labeled (Figure 5; a black and white version was given to survey participants). Giving this to them involved making a tradeoff between two competing objectives. The first concerned offering sufficient information for respondents to find familiar features and identify relevant places that they could include within their community of interest; the other sought to avoid forcing them into a network-based mindset in which they felt compelled to draw their lines along streets or highways. I concluded that ensuring that people drew their polygons in as informed a manner as possible trumped the concern about the network-based thinking. The fact that most people still drew polygons that did not follow any streets or features suggests that this was an appropriate decision.

Another choice made with important implications was to specify the scale of the community of interest as that of a city district or large neighborhood. I wanted to guide respondents into drawing a polygon with a similar size as that of their district in order to facilitate comparison between the two regions, and more fundamentally because city council districts are not drawn to conform either to a community of interest encompassing the entire city or beyond, or one limited to just a couple of street blocks. Still, several participants did not follow these instructions but rather drew citywide or block-focused polygons. Without this instruction, it seems likely that more participants would have drawn very large or very small communities of interest, thus potentially compromising comparisons even more. Still,

it is likely that there are upper and lower limits for what most people would reasonably call a community of interest; surely no one thinks of a whole state as such, nor do most think of just their street block that way. There is reason to believe, however, that the scale of a community of interest can be further honed between these bounds.

Evidence from the literature as well as the findings of this study point to a close approximation in scale between the community of interest and the neighborhood. The neighborhood, which might be defined as a geographically contiguous and socially homogenous categorization of space, is a concept that is rich and varied itself (Spielman and Logan 2013). The California Constitution suggests that the two concepts closely relate by juxtaposing “local community of interest” with “local neighborhood” in its list of entities whose geographic integrity should be respected (Article XXI, Section 2-d-4). The fact that Spielman and Singleton (2015a) interchange “neighborhood” and “community” has already been mentioned. Yet Monmonier (2001) pushes back on the idea that neighborhoods and communities of interest are synonymous, arguing that “communities of interest are almost always larger and more fragmented than one’s immediate neighborhood,” in large part due to improved transport and communication links (152-53). The cognitive communities of interest identified through this research support Monmonier’s notion, especially the fact that the one in District 2 incorporates several areas, such as Mesa, Bel Air Knolls, and the Upper Westside, which are identified by a local newspaper as distinct neighborhoods that are yet linked together as a larger “Mesa” (San Andres-Calleja 2016). These neighborhoods are also fragmented from each other by the hilly terrain. All these facts give credence to the idea that the community of interest can be considered a group of like-minded neighborhoods.

The issue of the scale is especially important because in California and elsewhere the community of interest criterion is employed for districting at the federal, state, and local government levels. If community of interest can be used as a basis for making a district as large as California's 8th congressional district (about 85,000 square kilometers) just as much as it can for making one as small as Santa Barbara's 3rd city council district (a little more than 2), can there still be a constant scale for the concept? The answer may well be yes, in large part because the scale of a community of interest need not be equivalent to that of the district itself. The California Constitution demonstrates this fact by placing "community of interest" in a list of other geographic entities that are generally smaller than a federal- or state-level district: cities, counties, and neighborhoods (Article XXI, Section 2-d-4). This stipulation mandates that efforts be taken to keep these entities whole, but that does not imply that there is only one entity for every district that should be kept wholly within that district. Rather, at least at these levels, there is a *collection* of cities, counties, neighborhoods, and communities of interest that falls within each district. A number of these entities will be located in the interior of the district and have no trouble fitting completely within the district boundaries, but many will lie at the periphery and require special care to not be split between different districts. So that is how the concept can still be relevant even with very large districts; such districts are made up of collections of communities of interest that are similar to each other.

Future research can investigate how people conceive of these collections in the context of larger-scale congressional and state legislative districts. A potential method to study this would be to have participants break up a given region into a certain number of subregions according to how many districts that region would merit at a particular level of

government: U.S. House of Representatives, State Senate, State Assembly, etc. I might have used this method for Santa Barbara's city council districts, instructing participants to break up the city into six subregions (there are six city council members) based on perceived communities of interest. Doing so would have allowed me to leave the scale undefined and give respondents freer rein to decide that for themselves. It would have also prevented respondents from drawing huge polygons encompassing the whole city, as they would have to stay within the city limits and segment that area into six smaller areas. However, such a study would have necessitated surveying across the city, as one could not conclude much about the city at large from the responses of just half the city's residents (those living in Districts 1, 2, and 3). Given the resources I had, such a survey was not feasible for this study. However, future studies might allow for such methods.

What remains to be determined is just how large an individual community of interest can be widely considered, not just a collection thereof. Even if one can confidently define it as a group of like-minded neighborhoods, people might differ on what number exactly a "group" entails. To determine the answer to this question, a future study would have to explicitly ask participants to define the extent of a single community of interest, and give them the freedom to make it as large as they want. Such a study would investigate the range of scales that people believe a community of interest can encompass, without necessarily a reference to any type of district. The value of leaving the scale unspecified lies in giving respondents more latitude to show what size a community of interest looks like to them. Perhaps a digital mapping platform allowing users to pan and zoom to whatever scale or extent they choose would give more meaningful results in this regard.

The possibility of non-response bias inevitably hovers over the results obtained by the survey. While the response rate was a decent 60%, still a good many declined to participate because they were about to leave the house, were in the middle of an activity, had no time to spare, or just preferred not to participate. Despite this fact I do not believe that non-response presents too great of a concern with my study because no particular group of people seemed to be too underrepresented in the survey. I obtained responses from a wide range of ages and appropriate amounts from the various ethnic groups considering their proportions in each district, with the possible exception of Hispanics in District 3. The slight excess of females over males may reflect a somewhat greater willingness to participate among the former, but I have no reason to think that beliefs about community of interest significantly differ with sex. Thus I do not find non-response bias to cast much doubt on the validity of my results.

Probably the greatest weakness of the survey design lies in its inability to assess variation within and not just between the communities of interest drawn by individuals. The fact that participants only had to draw a single line around what they believed to be their community of interest means that all space within that line—even just within—is considered 100% part of their community, and all space outside—even just outside—is considered 0% part of it. Yet, of course, people do not have such a monolithic conception of their cognitive region but rather recognize a gradation within it (Montello, Friedman, and Phillips 2014). Giving them the chance to represent variation within their region using some kind of cell rating task would provide greater clarity about not only which areas are most agreed upon but where the cognitive “cores” of each district are located, where the ratings are highest.

Such a time-intensive task would tax the patience of respondents and likely limit the amount of participants I could obtain, but this would be informative and worth the effort.

VII. Conclusion

By studying communities of interest both thematically and cognitively, this project shows that both definitions compare reasonably well with each other and have unique and important contributions to make to our understanding of what a community of interest actually is. These findings demonstrate that the concept of the community of interest does indeed gain precision in its definition by incorporating both some key thematic indices as well as the perspectives of individual residents. The thematic community of interest represents all the land uses that should be included when drawing district boundaries but are often discounted by the perspectives of individual residents. On the other hand, the cognitive community of interest depicts the area where most people agree upon and should represent the core or center of whatever district is being crafted. The fact that both types of communities cohere rather well with existing districts shows that it makes sense to consider communities of interest when drawing borders even at this small scale.

If state or city officials wish to give communities of interest proper consideration, they would serve their citizens well by identifying those areas where people most agree is in their community of interest. They could do so by hiring out one or more individuals to administer a survey to residents, using either the freeform polygon method as I did in this study or the segmentation into subregions method. This would likely provide greater feedback in quantity and quality than could be gained from a public forum, though more resources would be necessary. Once areas of strong agreement are identified, officials can make these areas the cores around which they form the rest of the district. They can then

utilize the thematic clustering method to ensure that the districts comprehensively cover the city. They would link such clusters according to how similar their attributes are to those of clusters at the core, after of course considering population, contiguity, and compactness. Such an exercise would be quite feasible for most authorities, and the benefits to citizens' sense of representation would very likely outweigh whatever costs might accrue. Hopefully more attention directed to both types of communities of interest will result in districts that are more representative and responsive to the needs and preferences of their citizens.

References

- Bae, C. J. 2015. Representations of an urban neighborhood: Residents' cognitive boundaries of Koreatown, Los Angeles. Master's thesis, University of California, Santa Barbara.
- Cain, B. E., K. Mac Donald, and M. McDonald. 2005. From equality to fairness: The path of political reform since *Baker v Carr*. In *Party lines: Competition, partisanship, and congressional redistricting*, ed. T. E. Mann and B. E. Cain, 6-30. Washington, DC: Brookings Institution.
- City of Santa Barbara. 2013. City of Santa Barbara general plan. Adopted 25 June 2013. <http://www.santabarbaraca.gov/civicax/filebank/blobdload.aspx?BlobID=16899>.
- City of Santa Barbara. 2016. District elections. Last updated 1 Feb. 2016. www.santabarbaraca.gov/gov/vote/district_elections.asp.
- City of Santa Barbara. 2016. Neighborhood construction & design activity. Last updated 1 Feb. 2016. <http://www.santabarbaraca.gov/services/construction/building/default.asp>.
- Constitution of California, Article XXI. http://www.leginfo.ca.gov/.const/.article_21.
- Courtney, J. C. 2008. From gerrymandering to independence: District boundary readjustments in Canada. In *Redistricting in comparative perspective*, ed. B. Grofman and L. Handley, 11-26. New York: Oxford University Press.
- Fan, C., W. Li, L. J. Wolf, and S. W. Myint. 2015. A spatiotemporal compactness pattern analysis of congressional districts to assess partisan gerrymandering: A case study with California and North Carolina. *Annals of the Association of American Geographers* 105 (4): 736-53.

- Frontiera, P., R. Larson, and J. Radke. 2008. A comparison of geometric approaches to assessing spatial similarity for GIR. *International Journal of Geographic Information Science* 22 (3): 337-60.
- Grofman, B. 1993. Would Vince Lombardi have been right if he had said: "When it comes to redistricting, race isn't everything, it's the *only* thing?" *Cardozo Law Review* 14: 1237-76.
- Handley, L. 2008. A comparative survey of structures and criteria for boundary delimitation. In *Redistricting in comparative perspective*, ed. B. Grofman and L. Handley, 265-84. New York: Oxford University Press.
- Hill, L. L. 1990. Access to geographic concepts in online bibliographic files: Effectiveness of current practices and the potential of a graphic interface. PhD thesis, University of Pittsburgh.
- Htun, M., and G. B. Powell. 2013. Between science and engineering: Political science, electoral rules, and democratic governance. In *Political science, electoral rules, and democratic governance: Report of the task force on electoral rules and democratic governance*, ed. M. Htun and G. B. Powell, 1-13. Washington, DC: American Political Science Association.
- Johnson, D. 2015. City of Santa Barbara: Drawing the lines. National Demographics Corporation. Published on 18 Mar. 2015.
<http://www.santabarbaraca.gov/civicax/filebank/blobdload.aspx?BlobID=50303>.
- Lynch, K. 1960. *The image of the city*. Cambridge, MA: The MIT Press.
- MacEachren, A. M. 1985. Compactness of geographic shape: Comparison and evaluation of measures. *Geografiska Annaler: Series B, Human Geography* 67 (1): 53-67.
- Mann, T. E. 2005. Redistricting reform: What is desirable? possible? In *Party lines: Competition, partisanship, and congressional redistricting*, ed. T. E. Mann and B. E. Cain, 92-114. Washington, DC: Brookings Institution.
- Medew, R. 2008. Redistribution in Australia: The Importance of One Vote, One Value. In *Redistricting in comparative perspective*, ed. B. Grofman and L. Handley, 97-105. New York: Oxford University Press.
- Monmonier, M. S. 2001. *Bushmanders and Bullwinkels: How politicians manipulate electronic maps and Census data to win elections*. Chicago: University of Chicago Press.
- Montello, D. R. 2003. Regions in geography: Process and content. In *Foundations of geographic information science*, ed. M. Duckham, M. F. Goodchild, and M. F. Worboys, 173-89. London: Taylor & Francis.

- Montello, D. R., M. F. Goodchild, J. Gottsegen, and P. Fohl. 2003. Where's downtown?: Behavioral methods for determining referents of vague spatial queries. *Spatial Cognition and Computation* 3 (2-3): 185-204.
- Montello, D. R., A. Friedman, and D. W. Phillips. 2014. Vague cognitive regions in geography and geographic information science. *International Journal of Geographic Information Science* 28 (9): 1802-20.
- National Demographics Corporation. 2015. Draft Plan 3B. Published on 26 Mar. 2015. <http://www.santabarbaraca.gov/civicax/filebank/blobdload.aspx?BlobID=50821>.
- National Demographics Corporation. 2015. Santa Barbara public forums summary report. Published on 23 Mar. 2015. www.santabarbaraca.gov/civicax/filebank/blobdload.aspx?BlobID=50344.
- Openshaw, S. 1983. *The modifiable areal unit problem*. Norwich, England: Geo Books.
- Phillips, A. 1995. *The politics of presence*. New York: Oxford University Press.
- Pitkin, H. 1967. *The concept of representation*. Berkeley and Los Angeles, CA: University of California Press.
- Potthoff, G. 2015. Santa Barbara agrees to settle district elections lawsuit. *Noozhawk*. Published on 15 Feb. 2015. www.noozhawk.com/article/santa_barbara_agrees_to_settle_district_elections_lawsuit.
- Powell, G. B. 2000. *Elections as instruments of democracy: Majoritarian and proportional visions*. New Haven, CT: Yale University Press.
- Powell, G. B. 2004. Political representation in comparative politics. *Annual Review of Political Science* 7: 273-96.
- Prescott, J. R. V. 1965. *The geography of frontiers and boundaries*. Chicago: Aldine Publishing Company.
- San Andres-Calleja, A. 2016. Basic facts about TMP. *The Mesa Paper* 9 (6). Published on June 2016.
- Spielman, S. E., and J. R. Logan. 2013. Using high-resolution population data to identify neighborhoods and establish their boundaries. *Annals of the Association of American Geographers* 103 (1): 67-84.

- Spielman, S. E., and A. Singleton. 2015a. Studying neighborhoods using uncertain data from the American Community Survey: A contextual approach. *Annals of the Association of American Geographers* 105 (5): 1003-25.
- Spielman, S.E., and A. Singleton. 2015b. Studying neighborhoods with uncertain census data. Github. Published on 22 July 2015.
https://github.com/geoss/acs_demographic_clusters.
- Vickrey, W. 1961. On the prevention of gerrymandering. *Political Science Quarterly* 76 (1): 105-10.
- Winburn, J. 2008. *The realities of redistricting: Following the rules and limiting gerrymandering in state legislative redistricting*. Lanham, MD: Lexington Books.
- Woodruff, A. 2012. Crowdsourced neighborhood boundaries, part one: Consensus. *Bostonography*. Published on 2 July 2012.
<http://bostonography.com/2012/crowdsourced-neighborhood-boundaries-part-one-consensus/>