

# Local Diversity and Polycentric Democracy

RYAN MULDOON\*

## ABSTRACT

*Polycentric democracy holds great promise as a means to deal with the core democratic challenge of intractable disagreements that stem from our diversity. Instead of simply having majorities win, polycentric orders aim to find a way of letting everyone get what they want. Polycentric solutions are often presented as cost-free, as only those who choose to join a club or move to a new locale are affected. This paper examines the ways in which quite serious costs to third parties can arise in polycentric orders, and then offers some possible ways to recover the benefits of polycentricity while mitigating some of these costs.*

Especially at the local level, democratic orders have a number of ways of handling the diverse interests and needs of citizens. There is, of course, deliberation and majority-rules voting, which allows people to determine some rule for all citizens. But more interestingly, especially when we consider provisioning local public goods or services, we gain extra options that have serious prima facie appeal: we can either spatially split the electorate (via foot voting between existing jurisdictions or creating a new jurisdiction), or we can use private governance to allow for club goods provision that does not require support of the general electorate. These latter two options are in several ways superior to deliberation and majority vote for accommodating a more diverse citizenry. After all, we can do a far better job of ensuring that citizens are getting closer to the sort of social compact they are interested in, because they can more easily get more of the costly social goods that they want, and fewer of the ones that they do not. However, I argue that these mechanisms come with serious costs, especially when we consider the longer run dynamics of these approaches. These approaches have helped fuel stark racial and socio-economic segregation and large differences in quality of governance. Most worrisome, these costs are frequently shouldered by those outside of private governance contracts and those who are being moved away from. In particular, I am interested not only in material harms, but in the democratic harms that those outside of the extra polycentric arrangements bear, in the form of reduced democratic voice, and a weaker real option of foot voting. The aim of this paper is to explore the extent to which we can recover the benefits of polycentric democracy while taking these externalities into account.

## I. THE CHALLENGE OF DIVERSITY

At the core of any democratic system is a need to harness the benefits of a diverse populace while managing the disagreements that arise from that same diversity. Different

---

\* Ryan Muldoon is an Associate Professor of Philosophy and Director of the Philosophy, Politics and Economics Program at the University at Buffalo. His work focuses on the challenges and opportunities of diverse societies.

citizens have different interests, different wants, and different needs. While much of this diversity can be handled by having a robust private sphere, in which individual preferences or choices do not need to be turned into a matter of public policy, there will remain a fairly broad public sphere in which we will need to find ways of managing our differences. The basic challenge, of course, is that we must make a single choice, despite our differences, about the rules that we are all to follow. Whatever rules we select might leave much to individual choice, but regardless of how extensive the set of rules we choose for ourselves, it remains the case that we must select some rules for ourselves. Although it would be nice if our differences were always worked out via a robust debate that led to consensus, much more typically we find ourselves at odds with each other, and for any particular political choice that we make as citizens, some portion of the population will wish we made another one. This is a core function of any liberal democratic order—offering an account of why people are bound by rules they would have not chosen (and may have indeed not selected when given the option).

To make this worse, insofar as we have political coalitions around clusters of issues, it is perfectly possible that some minority group will not merely lose on *a* vote, but on a whole *series* of votes. Not getting your way every now and again is of course part of the democratic bargain, but citizens may well feel that *systematically* losing is something else entirely. If part of democratic participation is having some sense of authorship over the rules to which one is bound, that notion of authorship becomes increasingly tenuous the more one sits on the losing end of votes. Especially in first-past-the-post, winner-take-all democratic orders, if some minority group holds 30% of the votes, it might well mean that rather than getting their way 30% of the time, they simply never get their way. This may mean that for minority groups of various sorts—especially those that aren't able to make common cause with a broader political coalition—their democratic voice is essentially brought down to zero, because they will not be able to exercise any real authorship over the rules that bind them.

It could be that a single set of rules for everyone might fail to take not just mere preferences seriously, but that it fails to accommodate different plans of life. What's at stake is not merely having the set of rules a citizen would most enjoy, but having a set of rules that makes space for, or even enables, a particular form of life that is perhaps a bit different from the mainstream. An Amish plan of life requires a (somewhat) different set of rules,<sup>1</sup> as does a Quaker life,<sup>2</sup> or living on a kibbutz. Less dramatically, the rules that might make sense in small-scale farming community would simply fail to address crucial issues in Manhattan life, and vice versa. If we are to take diversity seriously, and respect all citizens on equal terms, such equal respect may entail allowing for different rules that cover different situations.

## II. THE PROMISE OF POLYCENTRICITY

Polycentricity offers a potential way forward for the problems of persistent, stubborn disagreements. It is wonderfully simple: instead of resolving a dispute by means of a majority vote, resolve a dispute by separating political authority. This can happen in two ways: either geographically, by means of having non-overlapping geographic jurisdictions that can have

---

<sup>1</sup> Thomas W. Foster, *Separation and Survival in Amish Society*, 17 SOCIO. FOCUS 1 (1984).

<sup>2</sup> PINK DANDELION, *THE QUAKERS, A VERY SHORT INTRODUCTION* (2008).

different sets of rules governing them, or non-geographically, by means of having different, optional, sets of rules covering different people in the same geographic area by means of using private governance arrangements. In the first case, where people can “vote with their feet” by moving between jurisdictions, jurisdictions are put into competition with one another for citizens. Citizens simply look to identify the jurisdiction whose rules and local public goods most matches their personal preferences or their particular needs and then settle in that jurisdiction, where they will tend to have fewer disagreements with their (new) fellow voters. In the second case, citizens generally stay within their existing municipality, but if most voters are uninterested in provisioning some public good, those who can form a club and provide the good for themselves as a club good. Crucially, both approaches aim to resolve disputes not by requiring that the minority submit to the will of the majority, but by finding room to accommodate both by allowing multiple sets of rules to apply at once. The core idea is that we can agree to disagree about the best set of rules (or even if there is such a thing) and instead focus on the idea that we can find communities of interest that are spaces of democratic consensus, or at least convergence, within a broader space of democratic disagreement. Let’s examine each case to see why these are options that hold real democratic appeal.

Tiebout developed a highly influential model of “foot voting,” framed in terms of a solution to a problem of the inefficiency of public expenditures.<sup>3</sup> Tiebout argued that while a national model of public expenditures would always suffer from preference mismatch, local public expenditures would not have to have this problem because a number of nearby municipalities could be treated as if they were offering competing goods bundles in a spatial economy.<sup>4</sup> So, with a suitably mobile population, voters will move to the municipalities that offer them the bundle that most appeals to them.<sup>5</sup> Municipalities with too few citizens for their given social contract will try and entice new citizens, and municipalities with too many citizens will try and put up roadblocks to new entry.<sup>6</sup>

This kind of model is straightforwardly appealing. Why fight with your neighbors over the construction of a municipal golf course when you could just move to a town that already has one or to a place where instead people prefer tennis? If you do not like paying high tax rates, move where taxes and services are low. If you like excellent public schools and do not mind paying for them, move to a high-tax district with strong schools. This model manages disagreement by citizens simply moving away from the people who disagree with them and moving toward the people with whom they agree. Unsurprisingly, there is plenty of evidence that housing choice is in part driven by what is now called “Tiebout sorting”— people go to where the amenities they want are and steer clear of places that do not offer the policy mix that they prefer.

We can easily see, especially in the confines of a model, how we can make sense of recovering the value of political authorship in diverse societies. People can move to the places that have rules that they would like to be bound by, and can then positively endorse those rules in future political activities. Presumably, as issues that come up that might be at least a little

---

<sup>3</sup> Charles M. Tiebout, *A Pure Theory of Local Expenditures*, 64 J. POL. ECON. 416 (1956).

<sup>4</sup> *Id.*

<sup>5</sup> *Id.*

<sup>6</sup> *Id.*

divisive within their more politically coherent communities, those disagreements can be resolved in part because it is easy to draw from a deep well of comity and mutual understanding. If those new disagreements prove to be more foundational, there is always another round of sorting available. Voice and loyalty are available to resolve political disputes, but exit is a powerful option, regardless of whether it is exercised first or as a last resort. Tiebout sorting just relies on the availability of a menu of municipal options that citizens can choose between and sufficient mobility for people to go where they want.

The second sort of polycentricity owes much to Buchanan. Buchanan argued contra Samuelson that rather than there being a sharp division between private goods and purely public goods or activities, we can understand a middle-ground option—clubs, and their attendant club goods.<sup>7</sup> Although I think we can understand some versions of non-overlapping geographic jurisdictions with the use of the Buchanan model, I am more interested in focusing on the model's use for understanding clubs within a single geographic area. The core idea is the following: instead of a stark divide between private goods that one acquires via normal market mechanisms and public goods that can only be provisioned via state action, there is a whole range of goods that can be collectively procured by groups of people.<sup>8</sup> For instance, while it is possible to purchase a pool for one's backyard, often people cannot do so, either because it is too expensive, or because they lack the requisite space. One option would be to lobby the municipal government to construct a public pool; however, this option would entail raising everyone's taxes, and perhaps not enough people in the town value pools over libraries, filled potholes, or cash. Here, we would run into the standard problem of democratic disagreement. Tiebout sorting suggests the solution would be to move to another municipality that already has ample numbers of public pools. But perhaps for various reasons people don't want to move. So, a final option presents itself: find a group of like-minded individuals and agree to jointly supply and maintain the pool. That is, form a pool club. The members of the pool club get a good deal—they all get access to a pool, while paying substantially less than they would have to provision it themselves. Likewise, the other members of their municipality get a good deal—they do not have to do or pay anything and can just continue with the existing set of public goods that they already pay for. Although those outside the club cannot access the pool, the reason for a private pool club was that the general public didn't want a pool enough to pay for it, anyway.

On this sort of model, there is a “base” set of rules that everyone follows and a “base” set of public goods that everyone could agree on, but beyond those, people can choose to take on extra rules or pay extra to supply themselves with “local public goods” or “club goods.” The most exciting and interesting version of this model can have an arbitrarily large number of clubs, some in direct competition with each other, others only indirectly, in any given location that would allow citizens to pick and choose between a whole variety of possible rules and goods. Instead of a municipality itself offering an extensive set of services, all bundled and paid for with a mandatory tax, the municipality itself might only do whatever the local version of a “minimum state” might be, and then if it were to offer any other services, they would be optional and in competition with other possible providers.

---

<sup>7</sup> James M. Buchanan, *An Economic Theory of Clubs*, 32 *ECONOMICA* 1 (1965).

<sup>8</sup> *Id.*

The promise of this sort of approach is that we do not have to spatially sort to resolve our disagreements. Instead, we just pick between clubs that offer different club goods, and different attendant membership fees or sets of rules that we would be held by as club members. We resolve disagreements by forming communities of interest and allowing them to coexist with each other.

On both basic models of polycentricity, we find something rather appealing: disagreements can be resolved without resorting to any coercion. Disagreeing parties simply agree to disagree and do what they can to find like-minded people to jointly provide the goods or rules that they favor. Political minorities regain political authorship in virtue of their ability to organize amongst themselves to provide the rules they want to live by. The majority does not get in their way, as the majority is free to do as it pleases, as well.

### III. THE PERILS OF POLYCENTRICITY

In the standard presentations of either Tiebout sorting or Buchanan-style clubs, there is little that is not to like. These mechanisms of diffusing the problem of disagreement are powerful because they work by increasing citizen satisfaction. More people get more of what they want, and fewer people are on the hook for paying for things they don't support. However, in this section, I want to suggest that there are some extra challenges with polycentric approaches that need to be taken seriously—some of the nice consequences of the models flow from the idealizing assumptions of the models themselves. When we consider real-world conditions, there are problems that we need to contend with.

Broadly speaking, there are a few main categories of concern. We will start with the Tiebout model, which relies on assumptions that frequently fail to obtain. Assumption 1 states that citizens are fully mobile.<sup>9</sup> Assumption 3 is that there are a large number of municipalities to choose from.<sup>10</sup> Assumption 4 is that there are no frictions due to employment.<sup>11</sup> Assumption 5 is that there are no externalities or diseconomies between communities.<sup>12</sup> Let's examine each of these in turn.

The first assumption—the mobility assumption—seems at best decreasingly true. In the United States at least, rates of domestic migration have fallen from about 20% of the population moving in a given year through the 1960's to about 10% now.<sup>13</sup> These mobility rates reflect mostly intra-county mobility: only about 3.5% of the population moved beyond county borders.<sup>14</sup> So far, this might just mean that most people like where they live, and so there is no particular need for between-municipality comparisons. However, U.S. Census data suggests that only

---

<sup>9</sup> Tiebout, *supra* note 3, at 419.

<sup>10</sup> *Id.*

<sup>11</sup> *Id.*

<sup>12</sup> *Id.*

<sup>13</sup> PETER J. MATEYKA, U.S. CENSUS BUREAU, *DESIRE TO MOVE AND RESIDENTIAL MOBILITY: 2010–2011* (2015).

<sup>14</sup> *Id.* Since the median county size in the US is about 622 square miles, this suggests a fairly limited range of mobility.

about 20% of those who indicate a desire to move do so.<sup>15</sup> Of the people who desire to move, it is mostly younger renters who actually move.<sup>16</sup> Homeowners rarely move. This is somewhat unsurprising, because homeownership adds transaction costs to a move that are not present when just contending with rental agreements. Further complicating the picture, the least common stated reason for moving is a dissatisfaction with public services.<sup>17</sup> The most common stated reason is a desire to be in a nicer house.<sup>18</sup> This all combines to suggest that the mobility assumption is quite a bit stronger than what we find empirically. If Tiebout sorting is happening, it is mostly happening amongst the young, who move at greater rates. But even there, most moves are driven by the desire for a nicer house, and secondarily a nicer neighborhood. Local safety and public services are the least common stated reasons for a move.<sup>19</sup> Insofar as this is true, then the majority of the moves that we see are not really evidence for Tiebout sorting. Those moves are, broadly speaking, people upgrading their housing as they earn more money. The other disconcerting aspect of this is that only one in five people who want to move do move, which suggests differential abilities to move.<sup>20</sup> If this were the case, then using Tiebout sorting to resolve political disputes is only serving the relatively privileged, possibly at the expense of those who are unable to be as mobile.

Suppose, though, that we are already pretty well-sorted and that young adults mostly move (within their counties) to municipalities that suit their interests. This is fine, as this is mostly what the Tiebout model is about—movement within a local region between competing municipalities. Even there, however, there is at least a difficult match to the empirical data, in part because this is very differently realized in different parts of the United States. There is a fair amount of variation in how many local governments are available within a single county: it varies from just one to a high of 150. The median number is around ten. So, sorting amongst municipalities is more possible in certain parts of the country. Counties that contain the most municipalities are either in large, sparsely populated states or very densely populated states. In the sparse case, moving between municipalities implies a pretty distant move and thus, may not really match what Tiebout had in mind. This leaves us with very densely populated counties as good matches for Tiebout sorting. So, the size of the menu is quite different across the country, and, as mentioned before, the rate of inter-state moves is very low. This at least narrows the applicability of the model for the average citizen.

So far, the main drawback of Tiebout sorting is that it likely does not happen all that much, and perhaps there is a bias for whom it happens. If anything, that would suggest we should use it more, given its many attractive features. Two likely barriers to movement here are owning a home, which adds significant costs to moving, and employment prospects. As we noted before, Tiebout assumes away any discontinuities in employment prospects. This is a very reasonable assumption for an analytic model, but this is indeed a prime reason that people choose to move or stay put. People flocking to San Francisco are likely not doing so because they agree with its housing policies but because there is a huge boom of technology jobs in the San Francisco area.

---

<sup>15</sup> *Id.*

<sup>16</sup> *Id.*

<sup>17</sup> *Id.*

<sup>18</sup> *Id.*

<sup>19</sup> *Id.*

<sup>20</sup> *Id.*

Likewise, the population explosion recently experienced in North Dakota was not because of some revolution in new municipal codes, but instead because of a jobs boom brought on by fracking technology in the oil and gas industry. Depopulation in the Rust Belt is perhaps partially explained by municipal failures, but a much more important explanation is the broad economic trends of deindustrialization, automation, and increased global trade that caused several manufacturing centers to shed tens of thousands of jobs. Unfortunately, the transaction costs involved in homeownership are often tied up in employment prospects because a loss of a major employer can both motivate people to leave and simultaneously make it difficult to do so if housing prices plummet alongside the local job market. Likewise, a boom in jobs may create shortages in housing, even if municipalities might otherwise want more citizens. We could just say here that the results of the Tiebout model should motivate us to develop policies to overcome these hurdles.

However, the last contentious stated assumption—that there are no externalities or diseconomies between communities—is more worrisome. After all, there are many possible sources of negative externalities between communities. A simple one is that, if we assume that training for skilled professions is broadly exogenous to any given local municipality's political economy, then public schools are in direct competition with each other for teachers. This would be fine if different municipalities had close to the same capacity to pay—then willingness to pay would just be a measure of how much each area valued education, and there could be reasonable variance across towns. But if wealthier towns simply face smaller burdens in paying higher teacher salaries, then higher quality teachers will quite rationally go to where they can get paid the most for their services. Since teacher quality has such a large effect on lifetime earnings, poorer areas are going to have a harder time improving the social mobility of the next generation. We can tell similar kinds of stories for other local officials, police, emergency services, and administrators. Likewise, if towns are divided economically in this way, we should expect that businesses will typically locate in wealthier districts where people have more money to spend, providing better job opportunities for locals.

The worries here all have a common form: when Tiebout sorting is driven by *ability* to pay, rather than *willingness* to pay, we are likely to encounter several negative feedback loops that will not just make those communities more unequal but will make it more likely that worse-off communities will lack the capacity to do better for themselves because the engines of economic growth or social mobility will tend to locate in more prosperous areas. This negative feedback loop can show up in more punitive tax policy as well; poor areas tend to have higher tax rates to make up for a smaller tax base, which makes it less likely that wealthier people will move there or for new businesses to start. It is worth noting here that taxes may be higher even while providing a narrower range of services: it is just that having any kind of school, a police department, a fire department, and other basic services have high fixed costs. So, it might not be a matter of profligate social spending or failure of belt-tightening, but simply the fiscal reality of providing basic services at all.

Importantly, none of these dynamics have to be the result of people intentionally trying to harm the poor; it is just the result of sorting, plus some negative spillovers. Note here that the objection is not that the default expectation ought to be that the wealthy subsidize the poor, and it is wrong when they fail to do so. Instead, the objection is that the wealthy *outbid* the poor on

those resources that promote social and economic mobility. So, not only do residents sort themselves, but so do the labor market and educational opportunities. This would be less problematic if the supply of teachers or firms were fully endogenous to the local political economy. In that case, the wealthier municipalities would simply be creating more teachers to put to work in educating their kids. It is hard to argue that it is wrong to choose to invest more in your kids, whether that is with time or money. But instead, we face a more-or-less fixed pool of teachers relative to the political choices of local municipalities, and so the allocation of that fixed pool of teachers is a zero-sum problem between local school systems. We see a similar story with (national) firms choosing a new location to operate. Tax policy is a plausible policy lever here, but places that already have a scant tax base may not be able to offer tax incentives or compete on service provision. Municipalities are too small of a political unit to be fully responsible for their political fortunes—too many things just are not in their control. And as a result, many things that on a national scale might rightly be understood as endogenous to policy choices are on a municipal scale best understood as exogenous choices. Likewise, where there may be an opportunity for positive-sum interactions in, say, training teachers, at the national level, the market for teachers at the local level is zero-sum.

A variant of this basic model is people sorting *out* of some municipality. Rather than simply not choosing to live somewhere in the first place, we can imagine what happens when people choose to leave a place that does not suit them. On the one hand, this is great! People should be able to find social arrangements that fit their broader plans of life. But, as Tiebout notes in his assumption six, patterns of community services have an implied optimal community size.<sup>21</sup> He thinks of this in terms of the carrying capacity of a local beach, but perhaps more important is the inverse: there is a certain threshold of people below which a town will fail. If tax revenues get too small, the town will not be able to support basic services. If there are not enough people, businesses cannot stay open. As businesses close, the tax problems get worse, and that further encourages a downward spiral of disinvestment. One could say simply that this was a social arrangement that proved to be undesirable. But, note here that there could still be a fair number of people in the town as it fails. If all of those people more or less liked that social arrangement best, then through no fault of their own, they lost their favored social arrangement at the same time as they lost their capacity to move to a better situation. What is worse, we should expect that as we approach such a scenario, the people who have the most resources or flexibility will be the quickest to exit. The people left behind are put in an ever-worsening situation, all because they are unable to leave when they were better off before the town's collapse. As in the previous case, it could well be that the precipitating event for this downward spiral is something that is outside of municipal control. So, municipal failure could well come about not because of a bad set of policies, but because of bigger trends. And people who (quite rationally) move away might, in fact, contribute to the hardships of the people who remain. Once again, this does not depend on a claim that the rich ought to subsidize the poor but just the idea that there are threshold problems in the provision of most goods, as Tiebout recognized.

These cases, I think, indicate some serious worries with the Tiebout sort of model, not simply because some people are made better or worse off—that always happens one way or another—but because how people are made better or worse off is outside of their political control. It is not that they made a bad bet on a social arrangement, it is that their social

---

<sup>21</sup> Tiebout, *supra* note 3, at 419.



arrangement was broadly dictated to them by the choices of other people who are outside of their jurisdiction. They are left to respond somehow to forces that are not subject to their political authority. This is possibly worse than the problem of being a permanent minority. At least when you are a permanent minority, you can attempt to exercise voice and get people to care on your behalf, in part because there is at least some minimal sense in which you are all in it together. But in this Tiebout scenario, you are not really in the same political community as the people whose choices undermined your ability to carry out your plan of life. Any assistance from them would be understood as supererogatory behavior, not something owed because of the imposition of harms.

Perhaps it is just that the Tiebout model has some problems, but we can get what we want by focusing more on clubs. So, let us turn to the scenario in which people are not picking between political jurisdictions, but are instead staying within a single jurisdiction and are opting in to different clubs that can impose new rules on them or offer them new club goods. This should help us avoid the problem of being burdened by choices happening outside of one's political jurisdiction.

The Buchanan model is, like the Tiebout model, wonderfully straightforward. And at a glance, it is quite appealing. If Alice and Bob want a golf course, but Carol does not, Alice and Bob can form a club to provide themselves with the golf course without making Carol pay for it. This result is great. Where I want to put a bit of pressure is thinking about what happens once some innocuous clubs are in place.

Buchanan rightly noted that a theory of clubs is really “a theory of optimal exclusion.”<sup>22</sup> For clubs to be successful, they need to be able to provide benefits to their members at the exclusion of non-members. Otherwise, problems of free-riding arise, and the motivation to join the club shrinks to the point where the club can no longer sustain its club goods. But there are two problems with exclusion that I want to put some pressure on. First is the idea that exclusion itself is one of the goods that some clubs provide, rather than a means to secure club goods. Second is the idea that exclusion, combined with greater organization, can mean that clubs can more easily exert their will on non-members in a way that burdens them.

The first idea rests on an empirical regularity: many clubs do not increase their membership up to the carrying capacity of the club goods that they provide. Golf clubs hold their membership down, VIP areas often have lots of extra space, etc. This is a bit strange when looking at the basic incentives of the Buchanan model, as club members should aim to encourage more members to decrease the costs of provisioning the club good, up until the point of the marginal new member causing a decline in the quality or availability of the good. After all, a swimming pool is not fun if there are too many people in it to swim yourself, and fishing in an overfished pond is boring at best, frustrating at worst. Clubs have good reasons to limit their size based on the club goods they provide. But, in many cases, we see the prevention of new members much sooner than this. In other work, I model this as clubs offering their members exclusivity itself as a consumable good. Even if we suppose that this exclusivity is solely a matter of numbers and not, say, demographically based, new entrants to a community might be blocked despite a willingness to pay for services. If the “rump” of people who are left out are not

---

<sup>22</sup> Buchanan, *supra* note 7, at 13.

numerous enough or well-capitalized enough to start their own club, they are stuck without the amenities they would like, despite their availability and a willingness to pay. It is once again relevant to note here that the concern is not that the wealthy owe subsidies to the poor—it is that people who have the willingness and ability to pay for some goods are still denied access by gatekeepers.

One might choose to respond by stipulating that this is not something we should concern ourselves with. After all, liberal societies generally recognize rights of association, and part of that right is the ability to choose who you associate with and who you do not. So, if a golf club wants to be exclusive, let them. I am in general quite sympathetic to this view. However, this can be made more complicated in those cases where the loss is not merely the ability play golf, but the potential civic loss if those clubs occupy an important civic niche. Golf clubs and other recreational clubs frequently can serve as crucial networking centers and important sources of social capital. While associational rights are an important part of any liberal democratic account, we somehow need to balance those against ensuring that citizens have the real value of their civic rights, such that access to office is open to all on fair terms. The conceptual barrier between the public sphere and the private sphere is more permeable than we normally appreciate, precisely because clubs and other associations end up playing an important role in both. So, much like Buchanan’s insight that we ought to see goods not as a strict binary of public or private, so too we might want to see associations as taking on more or less “public-ness” depending on their broader civic role. So, while we might not be concerned if a local juggling club decided that a criterion of membership was the ability to juggle five balls, thus keeping their membership lower than their space and equipment could accommodate, we would have reason to worry if a Rotary club excluded Muslims, given the frequency with which Rotary clubs serve important civic roles in local political life.

Indeed, if we look at the (quite recent) history of associations exercising their rights to exclude, we find that exclusion was not driven purely by the carrying capacity of the club goods being provisioned, but frequently it was driven by demographic considerations. Many labor unions were exclusively for whites. Homeowner’s associations broadly began as a means of contractually ensuring that a new neighborhood would never have black residents.<sup>23</sup> There still are racially exclusive golf clubs. Churches in the United States are immensely racially segregated. Residential segregation has fueled quite extensive segregation in public schools as well.<sup>24</sup> While I clearly have the right to pick who my friends are, it is far less clear that I get to pick my neighbors as well. It is still less clear that I have any sort of right that allows me to block arbitrary categories of people from entering a public school system.

A variant on the exclusion worry stems not from intentional exclusion, but compositional details of groups. For instance, in my hometown, a leisure services division of the town government facilitated recreational sports and served as a sort of market-maker for recreational courses, like poetry writing or yoga. Where I went to college, recreational sports were mostly “beer leagues” —coordinated by bars. Where my wife grew up, recreational sports were “church league” —coordinated by local churches and their associated church groups. These are just a few

---

<sup>23</sup> RICHARD ROTHSTEIN, *THE COLOR OF LAW: A FORGOTTEN HISTORY OF HOW OUR GOVERNMENT SEGREGATED AMERICA* (2017).

<sup>24</sup> *Id.*

of the possible ways for a community to organize recreational sports. But very likely, each is going to be seen as more or less welcoming to different sorts of people. Atheists or religious minorities might not feel comfortable on a church league softball team, even if the people involved see themselves as welcoming. Likewise, beer league teams might not be great for recovering alcoholics, or people who have a dim moral view of alcohol consumption. It is plausible that say, undocumented immigrants, or people who are behind on their taxes, would be worried about the extra state contact implicit in participating in a town-led league. If there are lots of alternatives, this may not matter much. In some instances, though, network effects may make it significantly harder to launch alternatives once one organization exists. This is especially true if part of the value of joining the organization is the improved access to its other members. If social life or access to civic institutions pass through particular clubs and not others, then the presence or perception of exclusion starts to matter a great deal more.

The second sort of worry worth contending with is what happens dynamically when we have several clubs. While I think this is a somewhat general problem, I will focus on one of the more popular forms of a club: homeowners' associations (HOAs). HOAs are, in themselves, pretty innocuous. If you would like to live in a neighborhood with stricter rules and have a contractual guarantee that your neighbors will mow their lawns regularly and keep flowers out front, an HOA is just fine. Where HOAs get more complex is when they start providing themselves with more services or own their own infrastructure. Then, they become a private government, whose contracts are enforced by the state but whose procedures and decisions do not have to be constrained by the same rules that municipal governments would be operating under. On the one hand, this provides us with more potential experiments in governance arrangements, which might be a benefit. But on the other, it empowers HOAs to exclude people in ways that municipalities are blocked from doing.

Most interestingly, when we consider private governance relations, we find that members of these HOAs tend to provide themselves with services that are *duplicative* of what their municipal government provides and then frequently work to gain exemptions from taxes that would otherwise pay for those public services. HOAs, by their very existence, result in neighborhoods which are better organized politically than others without such compacts.<sup>25</sup> This means that they will be on average more effective at organizing to support their interests in municipal politics. An HOA with a pool might decide that the town shouldn't spend its money on public pools and thus, will be more effective than their political opponents in organizing and getting out votes. HOAs, insofar as they see themselves as mini-municipalities within a larger one, frequently decide that they would rather raise their HOA fees to assure access to some exclusive good than pay the lower cost of municipal taxes for the same—but no longer exclusive—good. And since they certainly do not want to pay twice, they will provide themselves with a club good and block others from getting a public good. This sort of withdrawal is harmful not because others do not get to get the benefit of their tax dollars, but because the private government has reasons to undermine the public government and the political organization to do so. In other words, rather than Alice and Bob getting their own golf club because Carol did not want it, Alice, Bob and Carol all want a golf course, but Alice and Bob want a golf course without Carol on it. Therefore, they undermine Carol's capacity to

---

<sup>25</sup> ROBERT JAY DILGER, NEIGHBORHOOD POLITICS: RESIDENTIAL COMMUNITY ASSOCIATIONS IN AMERICAN GOVERNANCE (1992).

collectively organize, while supplying themselves with what they want. They are then free to exclude Carol because, after all, it's a private club, and Carol is not a member.

#### IV. A WAY FORWARD?

Polycentricity, despite its faults, is a promising approach to deal with real diversity. Different people really do value things differently, and not everyone wants the same things or wants the same plan of life. It is important to find solutions to better satisfy a broad range of interests and plans of life amongst a large citizenry, as that would better secure the real value of their rights. Allowing for more club goods, and for a more diverse set of rules across different clubs or municipalities, is an appealing option. That said, the primary worry is that, at their core, sorting and clubs are just as much about excluding others as they are about accommodating a variety of interests.

If we are to preserve the very real benefits of these tools for managing disagreement, there must be some mechanism to maintain the viability of the idea of equal citizenship. Equal respect does not need to mean (and I think, likely ought not mean) that we are all bound by the same social contract, especially in those cases when it means some people will systematically fail to win the rules that they would most like to live under. In a community of equals, equality should manifest itself by offering all people the same real opportunities to explore their favored plan of life. This may require different sorts of social contracts for different sorts of lives. So long as people have real choices between these myriad options, the existence of different rules promotes a fairer, more just political community that takes seriously the differing wants and needs of a diverse population.

However, as we have seen in polycentric orders, people can still be burdened by the decisions of others, and these burdens can be harder to remedy if they are made in other jurisdictions. Sorting into different clubs or different municipalities does not mean that there are no spillover effects. Likewise, within a jurisdiction, there can be problems when exclusion becomes the goal, rather than a tool to help people get what they want when others do not want the same thing.

For a polycentric order to respect the value of equality, some steps are needed to bring political realities into closer alignment with the assumptions of those models. Perhaps most importantly, we need a way to secure a positive right to movement. While people may face no formal impediments to moving, they may not have the real value of such rights if they are underwater on their mortgage, excluded from other communities, or lacking sufficient information to engage in a search. That only 20% of people who would like to move do in fact move should be a red flag for advocates of Tiebout sorting.<sup>26</sup> One way of helping to remedy this problem would be to provide assistance for moving aimed at those with less mobility. This may take the form of cash assistance or may also look more like an assisted matching program, helping people find communities (and potential employers) that are a good fit for their values,

---

<sup>26</sup> MATEYKA, *supra* note 12.

interests, and skills. An assurance of assistance with a move, in whatever form it might take, may help resolve many of the issues with Tiebout sorting. Notably, however, this does not speak to the problems of zero-sum interactions between municipalities and may make the problem of municipal collapse even more acute. The benefit would be that more vulnerable citizens would face less of a risk of being stuck in a town's downward spiral. But those same citizens might mourn the loss of their community.

Community collapse is a difficult problem to solve, and it is not clear that there are good solutions available. Municipalities with shrinking fiscal capacity could be plausibly supported by the state, but while public support may keep the lights on for basic services, it will not change the basic dynamic that encourages disinvestment and emigration. It is possible that the state could facilitate similarly teetering communities to physically combine, to gather enough people in one place in order to allow for a stable community. But doing so would require a group of people to give up their homes and their community for a new one that's perhaps similar, but not the same. Many people would reject such a proposal.

If we turn our attention to clubs in the same municipality, we may want to draw on similar sorts of policy responses. If we are looking to combat "semi-public" clubs becoming exclusionary in ways that inhibit the realization of equal citizenship, we may have two complementary options. Towns, or perhaps counties, could more effectively serve as market-makers or clearinghouses of possible clubs as a way of making exclusionary clubs less problematic. If one golf club welcomes all comers, it matters less if another club is exclusionary. Facilitating the establishment of new clubs can also help bring new ideas to the forefront, and thereby allow for more social dynamism. For example, exclusionary golf clubs could be made irrelevant by inclusionary axe-throwing clubs or climbing clubs. A market-making role for municipalities or counties would reduce search costs for citizens and may help avoid any one given club from becoming too entrenched or central to civic life.

In those cases where it might be difficult to maintain multiple competing clubs, we might think that there could be some kind of inclusionary presumption in the club. Insofar as these clubs possess a legal corporate form, we could imagine a corporate form that assures membership in exchange for favorable tax treatment or some other benefit. We can think of this as clubs selling their right to exclude for some other club benefit. We can then have a clearer sense of which organizations view themselves as having a civic role and which consider themselves to be fully private.

Finally, we may wish to disallow clubs' (most notably HOAs') tax exemptions for "duplicative" services. This eliminates one of the core incentives for an unhealthy longer-run dynamic of replacing public goods with private goods. Clubs are a wonderful tool to facilitate inclusion of divergent and diverse interests, not purely exclusionary interests. If the state wishes to grant them economic benefits, like tax exemptions, those benefits ought to facilitate this broader public interest. If not, members of private clubs can make the choice of whether or not they want to pay for the full cost of their desire for exclusion.

These measures, when taken together, should help secure the role that polycentric arrangements have in promoting quite valuable methods of resolving democratic disputes, while

mitigating the very real costs that they can impose on others. While there are likely several ways of trying to balance competing liberal interests here, these proposals are aimed at protecting robust rights of association, while also promoting an account of political equality that respects a diverse citizenry.