DIVISION OF THE HUMANITIES AND SOCIAL SCIENCES

CALIFORNIA INSTITUTE OF TECHNOLOGY

PASADENA, CALIFORNIA 91125

A COMPARISON OF POLITICAL INSTITUTIONS IN A TIEBOUT MODEL

Scott E. Page
California Institute of Technology

Ken Kollman
University of Michigan

John H. Miller
Carnegie Mellon University

SOCIAL SCIENCE WORKING PAPER 926

May 1995

A COMPARISON OF POLITICAL INSTITUTIONS IN A TIEBOUT MODEL

Scott E. Page Ken Kollman John H. Miller

Abstract

In this paper, we construct a computational model of Tiebout competition. We show that the notion that Tiebout competition, as a result of enforcing efficiency, renders institutional arrangements unimportant does not preclude the possibility that political institutions may differ in their ability to sort citizens. In particular, institutions which perform poorly given a single location, may perform better when there are multiple locations because they allow for improved sorting. We demonstrate that insights from simulated annealing, a discrete nonlinear search algorithm, may explain this improvement.

Keywords: Tiebout Competition, Computational Models, Political Institutions.

Communications should be sent to Scott E. Page, 228-77 Caltech, Pasadena, CA 91125, email scotte@mrfloods.caltech.edu, phone 818 395-4216

A COMPARISON OF POLITICAL INSTITUTIONS IN A TIEBOUT MODEL

Scott E. Page* Ken Kollman† John H. Miller‡

1 Introduction

The performance of a political or economic institution depends critically upon its ability to structure micro–level incentives to be in agreement with macro–level goals (Schelling, 1978). In many important situations, micro–level incentives are consistent with multiple equilibria (Axelrod and Bennet, 1993; DeVany, 1994), and one role of institutions may be to steer agents towards the best configuration, or at least bias outcomes towards better configurations. In this paper, we compare the performance of various political institutions in a model of Tiebout competition. We examine whether some political institutions induce better sorting of citizens among competing jurisdictions, leading to configurations of higher aggregate utility.[1]

In a Tiebout model, local jurisdictions compete for citizens using bundles of public goods. Citizens then sort themselves among jurisdictions according to their preferences. Tiebout's (1956) original formulation challenged Samuelson's (1954) conjecture that public goods could not be allocated efficiently. The original "Tiebout hypothesis" has since been extended to include additional propositions. Prominent among them is that Tiebout competition, as a result of enforcing efficiency, renders local politics unimportant: a political institution able to attract and retain citizens cannot waste resources, i.e., it must be efficient (Hoyt, 1990). This argument does not preclude the possibility that political institutions may differ in their ability to sort citizens according to preferences, which is the focus of this paper.

*Division of Humanities and Social Sciences, California Institute of Technology 228-77, Pasadena, CA 91125.

†Department of Political Science and Center for Political Studies, University of Michigan, Ann Arbor, MI 48109.

‡Department of Social and Decision Sciences, Carnegie Mellon University

[1]Given that Tiebout models assume sorting with respect to preferences over public goods and services, the perspective of this paper is that configurations with greater homogeneity of preferences within districts create higher utility, which would not be the case if sorting occurs along racial or cultural dimensions then (Schelling 1978).

To investigate this question, we construct a computational Tiebout model. Under a variety of conditions and parameters in the model, significant differences arise in the performance of distinct political institutions. In the model, citizens confront a finite number of local public issues, which may be considered as either local public projects or local policies. Citizens have linearly separable preferences over these local public issues, and they choose to locate in the jurisdiction offering the most attractive array of decisions on the issues. Within this skeletal framework we compare the abilities of various political institutions to aggregate preferences and to sort citizens. Among the institutions considered are referenda, two–party competition, and multiple party competition under a variety of voting rules. In most respects, the findings agree with standard Tiebout models. We find, for instance, that both aggregate utility and the amount of movement between jurisdictions increase with the number of jurisdictions. We also find, rather unexpectedly, that institutions which perform poorly in a model with one jurisdiction perform well when there are multiple jurisdictions with migrating citizens. Political institutions not yielding a unique policy prediction, but whose policies sometimes cycle, appear to induce more sorting than institutions which yield unique policy predictions. As a result, institutions with policy cycles might generate better configurations of citizens' locations. We offer an explanation for this phenomenon based on insights from the literature on nonlinear search algorithms. The analysis suggests that in order for "political instability" to improve outcomes in a multi-jurisdiction environment, the degree of instability must be positively correlated with the heterogeneity of citizens' preferences. Later in the paper we construct measures of instability and heterogeneity of preferences, and find that three of the institutions we consider, two–party competition, Borda rule, and proportional representation, exhibit positive correlations between the two measures.

The remainder of the paper consists of five sections. In section 2, we briefly survey the empirical and theoretical literature on Tiebout's theory, and relate past research to the research presented here. The next two sections contain a description of the model and an informal explanation of the tradeoff between local political stability and global efficiency. In section 5, we describe our computational findings, emphasizing the relative performances of various political institutions. The final section contains directions for future work.

2 The Tiebout Hypothesis

This section summarizes several important papers since the original Tiebout formulation and indicates how this paper departs from previous research. Tiebout's hypothesis has been investigated in depth both theoretically and empirically.[2] The theoretical literature has uncovered flaws in Tiebout's argument, and many theorists, notably Bewley (1981), have clarified and extended Tiebout's original model. Bewley's characterization of a Tiebout Equilibrium (TE) consists of five conditions. First, each agent maximizes her utility subject to a budget constraint. Second, each agent inhabits the region she

[2]See Dowding, John, and Biggs (1994) for a recent survey.

most prefers. Third, each firm maximizes its profits. Fourth, regional governments have balanced budgets. Fifth, each regional government's choice of tax schedule and bundle of public goods is consistent with its preferences. Bewley considers two types of governments: *democratic governments* which maximize the welfare of their citizens and *entrepreneurial governments*, which have objectives, such as maximizing their local population, that are independent of the preferences of their citizens. For the purposes of this paper, the former type of government is more relevant. Bewley shows through a series of examples that regardless of whether the local public goods' costs are independent of the size of the local population or proportional to the size of the local population, there exist economies with TE which are not Pareto optimal. In other words, the Tiebout hypothesis does not hold. The existence of examples with suboptimal TE by itself is not sufficient to dispose with Tiebout's theory. More troubling is Bewley's result that the Tiebout hypothesis can be expected to hold only when the number of jurisdictions is at least as large as the number of citizens—a rather unlikely situation. Subsequent to Bewley's analysis, theorists have attempted to resurrect the Tiebout hypothesis, by relaxing his assumptions or including frictions in the relocation decisions. Among the assumptions of Bewley's simplified model are that agents have perfect, costless mobility, and that there are no external effects between agents and regions. Recently, economists have begun to pay attention to the importance of these external relationships (Durlauf, 1994), which may be relevant to the findings presented in this paper.[3] Nevertheless, we know of no formal theoretical research comparing the performance of qualitatively different political institutions in a Tiebout setting.

The empirical literature on Tiebout focuses on three questions: Does competition among jurisdictions lead to efficiency? Do bundles of local public goods enter into jurisdictional choice? How do various political institutions influence the provision of public goods? Tests of the first question belong to a larger literature on the size of government, commonly referred to as the *Leviathan Literature.* If, following Tullock (1965) and Niskanen, (1971), politicians are budget maximizing bureaucrats, then in the absence of competition, government budgets would grow too large. Attempts to test this theory and, implicitly, the Tiebout hypothesis, have taken many forms. Typically, these empirical models regress the size of government on the number of competing local jurisdictions. A negative relationship would support the Tiebout hypothesis that competition leads to greater efficiency. Unfortunately, empirical findings are inconclusive at best and appear to depend upon whether the local jurisdictions are single or multi–purpose (Zax, 1989).[4] Tests that bundles of public goods enter into jurisdictional choice fall into two categories: macro–level and micro–level studies. In a number of cases using aggregate data, fiscal differences among jurisdictions have been shown to have a significant effect on migration (Reschovsky, 1979).[5] These studies find that fiscal factors play a larger role in pushing

[3]A related literature examines the growth and development of cities (Arthur, 1989; Krugman, 1991). The work of Jacobs (1984) has been a driving force behind these models.

[4]Zax shows that the lowest levels of tax appeared to be paid when there is high fragmentation among multi–purpose jurisdictions, low fragmentation of single purpose districts, and decentralized governments. We use the word *appear* because the widespread existence of special districts to exploit returns to scale makes clean tests difficult.

[5]Cebula (1974) has shown that welfare payments had a significant effect on interstate migration in

people out of districts than in pulling people into new districts, which is not surprising given the informational asymmetry (Fox, Herzog, Schbottman, 1989). Other studies rely on micro–level survey data to examine the link between public services and goods, and jurisdictional choice. The findings in the micro–level studies are less conclusive, although school quality and tax rates appear as significant in several studies (Sharp, 1984; Percy and Hawkins, 1992; Percy, 1993). Finally, welfare policies across states in the U.S. do seem to influence migration patterns of the poor (Peterson and Rom, 1989).

The third and final question is most relevant for our purposes. There have been few empirical studies of the effects of political institutions on the provision of public goods. A notable exception is a paper by Romer and Rosenthal (1979) on school spending referenda. They find that the reversion level (what policy will be if the referenda fails) influences the outcomes of referenda. As a result of the meager literature in this area, we have little guidance from past work on parameter values. Thus, in what follows we are careful to test our model under a wide variety of conditions.

Our model departs from previous theoretical approaches in several important respects. First, with a few exceptions, our primary interest, comparing the performance of political institutions, has been neglected to a large extent in the Tiebout literature. A typical Tiebout model takes the political institution, usually majority rule, as exogenous and constant. Here we vary institutions and measure performance, an approach consistent with the literature on mechanism design. Second, aside from an analysis of an extended example, we do not compare equilibria. This decision stems partly from a desire to avoid oversimplifying an already stark model by making the restrictive assumptions necessary to guarantee the existence of equilibria. Also, our theoretical approach emphasizes the dynamic nature of population migration, political responsiveness, and the production of public goods. People may continue relocating in response to changing policies, while local policies react, in turn, to changes in constituencies. The system of policy choices by governments and jurisdictional choices by citizens can be viewed as a complex adaptive system in which movements and policies are determined by the preference of citizens and the political institution. Some institutions may create complex systems which settle quickly into equilibria while others may create systems which never equilibrate.[6] Third, our assumptions of how organizations make decisions differ from much of the previous literature in political science. Rather than model political parties as fully informed and optimizing, we rely on an adaptive party model in which parties gather information by taking polls and make decisions on the local public issues using heuristics (Kollman, Miller, and Page, 1992) . The adaptive party model will be described in more detail in the next section.

the U.S. from 1965–1970.

[6]One may also argue that local populations never completely stabilize in reality. Arguably, these movements may be part of a drift towards an equilibrium, but if the equilibrium point changes before the population can locate the previous equilibrium, restricting attention to equilibrium phenomena may be misleading.

3 The Model

Presentation of the model is divided into three parts. We begin by characterizing preferences. We next describe an adaptive party model and conclude with descriptions of the political institutions under consideration.

3.1 Preferences

We assume that there are a agents, k jurisdictions, and n binary variables, which we call *local public issues* (LPIs) upon which each local government is obliged to make a decision.

Def'n: The set of *agents* $A = \{1, .., a\}$.

Def'n: The set of *jurisdictions* $K = \{1, .., k\}$.

Def'n: The set of *local public issues* $N = \{1, .., n\}$.

A *configuration* is a mapping of agents to jurisdictions.

Def'n: A *configuration* $\Gamma : A \rightarrow K$.

To distinguish between issues and decisions made on issues, we refer to the latter as *local public decisions* (LPDs) and to a set of decisions on all n issues as a *local public decision array* (LPDA).

Def'n: A *local public decision* at location j on issue ℓ is $d_{j\ell} \in \{N, Y\}$.

Def'n: A *local public decision array* at location j is $d_j \in \{N, Y\}^n$.

The set of all local public decision arrays will be called a *global public decision array* (GPDA). As mentioned, a LPI may concern funding a public project such as a beach,[7]

[7]In this case, a LPD of Y might denote that the beaches are staffed and cleaned for the summer.

or alternatively, a LPI may pertain to a local policy on an issue, for example, public smoking laws. We assume that agents have linearly separable preferences on LPIs and that their per unit value for each LPI lies in the interval $[-\frac{4}{n}, \frac{4}{n}]$ distributed uniformly.

Def'n: Agent i's *per unit utility* on issue ℓ, $a_{i\ell} \in [-\frac{4}{n}, \frac{4}{n}]$ distributed uniformly.

An agent's utility from an LPDA equals the sum of her values on the LPDs.

Def'n: Agent i's *utility from LPDA* d_j, $u_i(d_j) = \sum_{\ell=1}^{n} a_{i\ell} \cdot \delta(d_{j\ell})$, where $\delta(Y) = 1$ and $\delta(N) = 0$.

A straightforward calculation verifies that the expected value to an agent of an arbitrary bundle of projects equals zero and the expected value to an agent of her optimal bundle equals one.

3.2 Adaptive Parties

Comparing the performance of political institutions in a formal model poses several difficulties, not least among them is determining electoral outcomes when theories do not arrive at unique predictions. To overcome this problem, we borrow Kollman, Miller, and Page's (1992) adaptive party model in which parties adapt platforms using search heuristics applied to polling information. We then average over many sample elections in order to make comparisons across institutions.

Our modeling of parties as incrementally adaptive and reliant on incomplete information contrasts with the rational choice approach in which parties either have complete information or act as Bayesians. ¿From our perspective, no single model of party behavior is likely to be accurate. Computational modeling provides us the flexibility to consider various models of party behavior and to learn which findings are particular to the assumed party behavior and which are generic. Practical considerations impose restrictions on the amount and type of information available to parties, on parties' computational abilities, and on the maximum allowable policy change in any one election. For example, we assume that in order to maintain both credibility and a coherent organization, parties can make only minor alterations in their LPDA in any one time period. We further assume that the parties rely exclusively on polling data as opposed to having direct knowledge of the utility functions of agents.

Computational constraints limit the actual number of behavioral rules we consider. In the computational experiments performed for this paper, we test three types of adaptive search rules for parties. These are random search, genetic algorithms (Holland, 1975), and directed search techniques. Our findings to date suggest that few implications depend significantly on the heuristic used by parties, provided the parties are neither omniscient nor dim witted.

The findings presented in section 5 of this paper include only a *multi-step hill climbing* heuristic, which can be described as follows: prior to the first election, a random initial platform LPDA is assigned to each party. In subsequent elections, parties begin with their LPDA from the previous election. Parties alternate adapting their LPDAs in an attempt to increase their vote totals. Suppose that a party's initial LPDA is given by y^*. The algorithm has two steps. First, a new LPDA y_1 is randomly created in a neighborhood of y^*.[8] If the new LPDA yields a higher vote total, y_1 becomes the party's LPDA. This process continues for a few iterations and then another party adapts. After each party has adapted several times, an election is held.

3.3 The Institutions

We model four political institutions: *democratic referenda, two–party competition, proportional representation,* and *Borda rule.* These labels only imply that we have created abstract models of the corresponding real world institutions. In order to formulate tractable representations of actual institutions, we have resorted to strong assumptions, and our approximations are crude at best. We have sought, however, to capture as many relevant aspects of each institution as possible without overly complicating the analysis.

3.3.1 Democratic Referenda (DR)

We model democratic referenda as majority rule on each LPI.[9] The assumption that there are no external effects between projects implies that sincere voting is a dominant strategy for all agents. The outcome of democratic referenda is the *median LPDA* at each jurisdiction, which is defined as follows:

Def'n: The *median LPDA at j given* Γ, $d^m_j(\Gamma)$:

$$d^m_{j\ell} = Y \quad \text{if } \mid \{i : \Gamma(i) = j \text{ and } a_{i\ell} > 0\} \mid \; > \; \mid \{i : \Gamma(i) = j \text{ and } a_{i\ell} < 0\} \mid$$

$$d^m_{j\ell} = N \quad \text{otherwise.}$$

[8]In our formulation, a neighborhood of a LPDA y consists of all LPDAs whose Hamming distance from y is less than or equal to three. Formally $N(y) = \{\hat{y} : \; d(y, \hat{y}) \leq 3\,\}$.

[9]In the case of a tie, we assume that the LPD equals N.

The LPDA d_j^m maximizes utility at jurisdiction j given Γ if and only if on every issue the mean agent value and the median agent value have identical signs. Generally speaking, democratic referenda locates a policy of relatively high aggregate utility given a configuration. In addition, democratic referenda is relatively stable: the policy prediction is unique and an individual agent migrating into or out of a jurisdiction rarely changes the median LPDA. Suppose that a single agent moving into a jurisdiction has a positive value on issue one. The only way that she can change the LPD on issue one is if, prior to her moving, an equal number of agents had positive and negative per unit values on that issue.

3.3.2 Two–Party Competition

We model two–party competition as competition among political parties advocating LPDAs. Each agent votes for the party proposing the LPDA which yields her higher utility. Two–Party competition is not as stable as democratic referenda, which produces a unique outcome equal to the median voter's preference on each LPI. Even with the linearly separable preferences considered here, there need not be a Condorcet winner. Policy predictions cannot be guaranteed to be a unique without severe restrictions on preferences (Plott, 1967). This has led theorists to try to formalize the equilibrium set for two–party competition. The top-cycle set (McKelvey, 1976), the uncovered set (McKelvey, 1986), and the minmax set (Kramer, 1977), are among the proposed solution sets. The absence of a unique equilibrium in two–party competition for large classes of preferences calls into question the effectiveness of democratic decision making (Bates, 1990).

Def'n: A LPDA d_j belongs to the *top cycle set at jurisdiction* j if for any platform $\bar{d}_j$ which defeats d_j, there exists a sequence of LPDAs $\{\hat{d}_i\}_{i=1}^m$ such that d_j defeats $\hat{d}_1$, $\hat{d}_i$ defeats $\hat{d}_{i+1}$ for all i and $\hat{d}_m$ defeats $\bar{d}_j$.

Claim 3.1 states that our assumptions about the utility functions of agents imply that the top cycle set contains the median platform d_j^m at each jurisdiction.

Claim 3.1 *At each jurisdiction, d_j^m belongs to the top cycle set.*

pf: Let $\mid d_j - \hat{d}_j \mid = \mid \{\ell : d_{j\ell} \neq \hat{d}_{j\ell}\} \mid$ equal the Hamming distance between two LPDAs. Choose d_j in the top cycle set at jurisdiction j. Let $p = \mid d_j - d_j^m \mid$. If $p = 0$, then the proof is complete. Choose $\bar{d}_i$ for $i = 1$ to p such that $\mid d_j^m - \bar{d}_i \mid = i$, $\mid \bar{d}_i - \bar{d}_{i+1} \mid = 1$ for all i, and $\bar{d}_p = d_j$. A simple calculation shows that d_j^m defeats $\bar{d}_1$ and that $\bar{d}_i$ defeats $\bar{d}_{i+1}$ for i less than p, which completes the proof.

In the next section, we will rely on the fact that the median LPDA array belongs to the top cycle set. It need not belong to the *uncovered set*, a refinement of the top cycle set.

Def'n: A LPDA d_j belongs to the *uncovered set at jurisdiction j* if there does not exist a LPDA d such that d defeats d_j and also defeats any LPDA defeated by d_j.

In the following example, the median LPDA lies outside the *uncovered set.*

Example: Suppose that there are three agents, a, b, and c, and three LPIs. Define preferences as follows:

Preferences			
Agent	LPI 1	LPI 2	LPI 3
a	+1	-0.4	-0.4
b	-0.4	+1	-0.4
c	-0.4	-0.4	+1

The median LPDA given these preferences is NNN. However, NNN does not belong to the uncovered set. The LPDA YYY defeats NNN and also defeats YNN, NYN, and NNY, the three LPDAs defeated by NNN.

This example demonstrates that predicting the LPDA resulting from two–party competition depends critically upon the solution concept. The top cycle set encompasses the entire space, so it has no predictive value. The uncovered set can be shown to equal every LPDA *except* the median LPDA, which has an obvious focal advantage. Fortunately, for our purposes, predicting an exact LPDA is less important than knowing bounds on the set from which it will be chosen. As previously discussed, we rely on an adaptive party model in which parties with incomplete information choose LPDAs using heuristics, so none of these solution concepts corresponds exactly to the LPDAs our adaptive parties will implement. However, the general characteristics of each should roughly correspond to what our adaptive parties find. Consider two scenarios: In the first, agents' preferences are sufficiently homogeneous so that there is a Condorcet winner. In the second, not only is there no Condorcet winner, but the top cycle set and the uncovered set encompass the entire space, or nearly the entire space as in the above example. Evidence from computations described in section 5 suggests that in the first scenario, one or both of the two adaptive parties will locate the Condorcet winner, while in the second scenario, the two parties will probably alternate in office. More importantly, they will advocate distinct LPDAs.

3.3.3 Proportional Representation (PR)

In our proportional representation (PR) system, each agent votes for one of several parties and each party receives a number of seats approximately proportional to their vote total. We assume no distortion between the percentage of the vote that a party receives and the percentage of seats it obtains in the legislature.[10] Unlike two–party competition, no party need receive a majority of the vote. When there is no majority party, there must be a second stage process of deciding policies. In our formulation, once a party is allocated seats, it votes sincerely in a series of referenda among the parties on the LPIs. The weight of a party's vote is equal to its percentage of the popular vote. So, if a party received 18% of the vote advocating the LPDA YNY, then that party would vote Y, N, and Y on the first, second, and third LPIs respectively. These three votes each would be assigned a weight of 0.18. The LPD on an issue equals Y if the total weight of the parties advocating Y exceeds 0.5 and N otherwise.[11]

In elections with more than two parties, an agent may have an incentive to vote strategically. Suppose that an agent preferred party a, advocating LPDA NNN to party b, advocating YYY. Suppose also that on the first LPI, that the agent prefers the LPD Y, and that the percentage of the vote won by parties advocating Y on the first LPI is approximately 0.5. If the percentage of the vote won by parties advocating N on the other two LPIs is significantly greater than 0.5, then the agent may wish to vote for party b. Such examples notwithstanding, we assume that agents vote sincerely. We take up the issue of strategic voting in the discussion at the end of the paper.

There are two additional issues which we should address: the number of parties and parties' incentives. In many PR systems, a minimum percentage of votes is required in order to win any seats. This acts as a constraint on the number of parties. The formation of parties and the merging of existing parties are difficult to model. Therefore, the number of parties in our model is an exogenous parameter. We experiment with three, four, and seven parties. With respect to incentives, we assume that parties are only concerned with maximizing the percentage of the vote they receive. We leave for future work the situation in which parties have policy preferences.

[10] There is a substantial literature on distortions between vote totals and legislative seat allocations. See Taagepera and Shugart (1989).

[11] These assumptions do not preclude a party receiving a small percentage of the vote from having substantial power.

3.3.4 Borda Rule (BR)

Under Borda rule, each agent ranks the parties in order of preference. An agent allocating votes to m parties, contributes $m-1$ votes to her favorite candidate, $m-2$ votes to her next favorite candidate, and decreases by one the number of votes to each succeeding candidate in order of preference so that the least preferred candidate receives no votes.

Borda rule is similar to proportional representation in that the transitions from votes to seats and from seats to LPDAs require second stage decision processes. As in our model of PR, we assume that the proportion of seats is identical to the proportion of votes, and that LPDs are determined by weighted democratic referenda among the parties with weights determined by the number of seats. We also assume, as before, that parties are only interested in maximizing their percentage of the vote.

4 Political Instability

In this section, we discuss how political instability may permit greater sorting of citizens and therefore be beneficial in a model of Tiebout competition. The argument hinges on the relationship between the level of political instability within jurisdictions and the degree of homogeneity of preferences at each jurisdiction. At this point we should clarify that at present we do not have formal proofs for many of the ideas we put forth. What follows is an informal explanation of our results. The discussion borrows insights from recent research on condensed matter physics. Later in this paper, we present findings from computational experiments that strongly support our explanation. Mathematical results may be possible, but our impressions to date are that any such results would be difficult to obtain even in a much simplified model. We believe, however, that the lack of sufficient mathematical tools should not preclude analysis. In fact, our computational model provides insights into how such proofs might proceed.

4.1 Two Examples

Two examples begin the discussion. The first example demonstrates how noise may improve sorting in a Tiebout setting. The second example shows how the Tiebout equilibria with respect to two–party competition may be preferred to the Tiebout equilibria with respect to democratic referenda in a multiple jurisdiction model.

4.1.1 Example: The Advantage of Noise

In this simple example, there are two jurisdictions α and β, ten agents, and one local public issue. Suppose that local public decisions are made using democratic referenda. To simplify notation, we identify agents by their preferred decision on the LPI in lower case letters. If agent 1 has a positive value on the LPI, she is denoted by y. Let α contain five agents $\{y, y, y, n, n\}$ and β also contain five agents $\{y, y, y, n, n\}$. The LPD in each jurisdiction equals Y. If agents do not consider the impact of their movements on policy, then no agent would want to relocate.

Suppose that we introduce noise and relocate the two agents who both prefer n, from β to α. Now, the n's are in a majority in α. As a result, the three agents in α who prefer y will then move to β. In this new configuration, each agent lives in a jurisdiction in which the LPD agrees with her preferred decision. Moreover, in this new configuration, if we were to randomly relocate any two agents, they would have no impact on the LPDs in either location, and they would immediately return to their original jurisdictions. Aggregate utility has increased. The point of this example is to show that there may be multiple equilibrium configurations, some of which are more stable then others. Noise may be able to improve aggregate utility by forcing a transition from one configuration to another one.

4.1.2 Example: Improved Sorting

A simple example demonstrates the tradeoff between local political stability and global efficiency. Prior to describing this example, we must clarify what we mean when we say that a configuration, Γ, is a *Tiebout Equilibrium with respect to an institution.* Given an institution, we first need a rule for the set of GPDAs which can result from each configuration. For democratic referenda, this rule consists of the median global public decision array. For two–party competition it might consist of all GPDA such that each LPDA belongs to either the top cycle set or the uncovered set in its jurisdiction. In the example below, the uncovered set and the top cycle set are identical for all relevant configurations of agents, so to simplify matters we will use the top cycle set as the rule for generating GPDAs. Thus, for the purposes of this example, a configuration Γ is a *Tiebout Equilibrium with respect to two–party competition* if for any GPDA, d, d_j lies in the top cycle set for all jurisdictions j given Γ, and no agent wants to relocate.

In the example, there are two jurisdictions: α and β, eight agents: a, b, c, d, e, f, g, and h, and three LPIs: 1, 2, and 3. Define preferences as follows:

Preferences			
Agent	LPI 1	LPI 2	LPI 3
a	+1	+1	+1
b	+1	-1	+0.5
c	+1	+0.5	-1
d	+1	-1	-1
e	+1	-1	-1
f	-1	-0.5	-1
g	-1	+0.5	-1
h	-1	+0.5	-1

Assume the following configuration of agents to jurisdictions: α contains agents a, b, and c, and β contains agents d, e, f, g and h. If democratic referenda is the political institution, then the LPDA in α is YYY and in β the LPDA is NNN. It is easy to show that no agent wants to relocate. Therefore, the configuration of agents and the GPDA form a Tiebout Equilibrium with respect to democratic referenda. A simple calculation shows that the aggregate utility equals 4. We now show that two–party competition does not support this configuration of agents to jurisdictions as a Tiebout equilibrium. In two–party competition the LPDA YYY is not a Condorcet winner in α. The LPDA YNN defeats it by a vote of two to one.[12]

One point of this example is to show that comparing the performance of institutions at a single jurisdiction in isolation may oversimplify the analysis, so prior to analyzing stability, we consider jurisdiction α in isolation. Democratic referenda would be preferred to two-party competition on utilitarian grounds. The sum of the utilities to agents a, b, and c from the LPDA YYY equals 4 while the sums of their utilities of the other LPDAs in the top cycle set YNN, YNY, and YYN equal 3, 3.5, and 3.5 respectively. If we assume for the moment that the LPDA resulting from two–party competition is randomly chosen from this set, then the lack of a unique outcome from two–party competition yields lower aggregate utility.

With multiple jurisdictions, the lack of a unique outcome from two–party competition may induce sorting and increase aggregate utility. Including jurisdiction β in the analysis, we find that two–party competition does not support this configuration of agents as an equilibrium. More importantly, this configuration lies in the basin of attraction of an equilibrium configuration which has a higher aggregate utility.[13] First, notice that if the agents in α elect a candidate advocating a LPDA from the set {YYY, YYN, YNY}, they will not create a strict incentive for any agents from β to move to α. If, however, a candidate wins election in jurisdiction α advocating the platform YNN, then agents d and e strictly prefer jurisdiction α to jurisdiction β, and they will relocate. Jurisdiction α would then contain agents a, b, c, d, and e and have a Condorcet winner, the LPDA YNN. Jurisdiction β would now contain agents f, g, and h and would have a new Condorcet winner, NYN. It is easy to show that this configuration of agents is a Tiebout Equilibrium with respect to two–party competition. By Claim 3.1 it is also a Tiebout Equilibrium with respect to democratic referenda. More importantly, the aggregate utility of this configuration equals 5.5 which exceeds the aggregate utility of 4 from the previous configuration

[12]Incidentally, YNN is also not a Condorcet winner. It can be defeated by either YYN or YNY. These four LPD form the top cycle set, which in this case also equals the uncovered set.

[13]By basin of attraction we mean the set of connected configurations that lead to higher aggregate utility than the status quo.

To summarize, two–party competition may result in the agents in jurisdiction α electing a candidate advocating a LPDA, YNN, which lowers their utility, but which encourages agents to move into α. The new configuration of agents is a Tiebout Equilibrium with respect to both two–party competition and democratic referenda, and yields a higher aggregate utility.

4.2 Insights from Simulated Annealing

To develop intuition for our results, we apply insights from the literature on *simulated annealing*, a nonlinear optimization algorithm. Simulated annealing is a sequential search algorithm applied to a real valued function, f. To apply simulated annealing, one must first create a neighborhood structure. Each point, x, in f's domain belongs to a neighborhood, $N(x)$, which contains at least one point different from x. Simulated annealing works much like local hill climbing: at each step a new point in the neighborhood of the status quo point is tested and becomes the status quo point if it has a higher value under f. Where simulated annealing differs from hill climbing is that it also accepts points which have lower values with some probability. This probability depends upon the difference in the function values, $\Delta(x, \hat{x}) = f(x) - f(\hat{x})$, and a *temperature*, $T(t)$, which depends upon the time spent searching, t. Formally, if x is the status quo point and $\hat{x}$ is the lower valued tested point, the probability of acceptance is written:

$$p(x, \hat{x}, t) = e^{\frac{-\Delta(x,\hat{x})}{T(t)}}$$

Notice that if the difference in function values, $\Delta(x, \hat{x})$, is large relative to the temperature, $T(t)$, then the probability of acceptance is low. The temperature can be interpreted as the degree of leniency. A high temperature allows for almost any new point to be accepted. A low temperature allows for a new point to be accepted only if the difference in function values is small. The temperature decreases to zero according to an *annealing schedule*. When the temperature nears zero, search converges to a local optimum with respect to the neighborhood structure attained.

A substantial body of theory exists to explain the performance of simulated annealing. Hajek (1984) has shown that given any function there exists an annealing schedule such that the simulated annealing procedure converges to the global optimum.[14] To some extent these results are misleading. The proofs, whether relying on Markov chain theory or real analysis, require that every point be evaluated. This begs the question: why perform simulated annealing instead of exhaustive search? The primary reason is performance. In practice, simulated annealing has been very effective at locating good solutions with relatively few searches, even though it may often fail to locate global optima.

[14] Similar results can be found in a special issue of *Algorithmica* (1991) Vol. 6, dedicated to simulated annealing.

A promising line of inquiry into the performance of simulated annealing considers the structure of local optima for the function being optimized. Given a neighborhood structure and a search algorithm, each local optimum has a basin of attraction. The larger a local optimum's basin of attraction, the more likely search ends at that local optimum. Simulated annealing performs better on functions with positive correlations between the values of local optima and the sizes of their basins of attraction. Local optima with small basins of attraction, and relatively low values, are more likely to be rendered unstable by the noise from the annealing schedule.[15] In effect, the mistakes due to the temperature smooth over the smaller local optima, but are less likely to disturb local optima with large basins of attraction. Eventually, the optima with large basins of attraction become inescapable as the temperature tends to zero. Thus, the outcome of simulated annealing results in a distribution across those local optima with larger basins and relatively larger optima.

In relating these insights to the performance of political institutions in a Tiebout model, we find strong, though not exact, connections. For explanatory purposes, we restrict the discussion to two political institutions: democratic referenda and two–party competition. In the previous section, we showed that democratic referenda yields a unique LPDA in each jurisdiction and that two–party competition offers less stability. We can interpret this instability as making mistakes in that the LPDAs selected typically have lower utility than the median. We can consider a global public decision array (GPDA) together with the configuration of agents it generates as a point in the domain. A neighborhood of a point could consist of all GPDAs which lie within a fixed distance of one another (according to some metric) together with their corresponding re-configurations of agents.

In order for a configuration of agents to be a Tiebout Equilibrium with respect to democratic referenda, no agent should want to relocate given the LPDA in her home jurisdiction. We know from Claim 3.1 that at each jurisdiction the top cycle set contains the median LPDA chosen by democratic referenda. If adaptive parties choose LPDAs from a subset of the top cycle set (which includes the median platform), then a configuration is less likely to be a Tiebout Equilibrium with respect to two–party competition than with respect to democratic referenda.[16] The fact that there are fewer equilibrium configurations with respect to two–party competition in no way implies that they have higher average utility. To complete the argument, we must return to the structure of local optima.

Simulated annealing locates good local optima because the mistakes bias search away from local optima with relatively low values and into local optima with relatively high values. This is more likely to occur when optima with low (high) values have small (large) basins of attraction. In such spaces, simulated annealing performs "as if" it can

[15]Basins of attraction are determined by the neighborhood structure. The minimal basin of attraction consists only of those points in the neighborhood of a local optima. A large basin of attraction might consist of all neighbors of all neighbors of all neighbors of the neighborhood of a local optima.

[16]Recall that in order to be a Tiebout Equilibrium with respect to an institution, no agent should want to relocate for *any* assignment of LPDAs from the sets of LPDAs in agreement with the institution.

recognize whether a local optima's value was relatively high or low, and escapes lower valued local optima. We propose that two–party competition performs similarly in a Tiebout model: *two–party competition makes more (fewer) mistakes in relatively low (high) valued configurations enabling it to act 'as if' it recognizes the value of a local optima and to escape it (remain in it) if it has a low (high) value.*

Two characteristics of Tiebout competition generate this fortuitous bias in error making. First, aggregate utility for a configuration of agents is positively correlated to the homogeneity of preferences at each jurisdiction. Second, more homogeneous preferences result in more stable outcomes from two–party competition.[17] If, for example, an LPDA is a Condorcet winner, we would expect adaptive parties to locate it. Combining these two effects, if agents in a configuration are not very homogeneous at each jurisdiction, their aggregate utility will be low and two–party competition may tend to roam about the space of LPDAs rendering the configuration unstable. If the agents in a configuration are homogeneous at each jurisdiction, their aggregate utility will be high and two–party competition will select from a small set of LPDAs. This second configuration is more likely to be stable. Thus, among the stable configurations with respect to democratic referenda, those stable with respect to two–party competition should be biased towards configurations with higher aggregate utility.

5 The Computational Model

In this section, we describe findings from computational experiments on Tiebout competition. Computational findings can be particular to the parameter values selected, so we have endeavored to test many sets of parameter values. Results reported below appear to be robust to reasonable variations. Findings presented below using one–thousand agents should not be interpreted to mean that the computations involving two–hundred or ten–thousand agents result in contradictory findings. The opposite is true, in fact. The particular parameter values chosen: one–thousand agents, eleven issues, and one, three, seven and eleven jurisdictions, are drawn from within a much larger set of parameters for which the results appear qualitatively similar.[18]

This section begins with a description of the sequence of events and an explanation of the adaptive search rules used by parties. We then summarize findings from the one jurisdiction model and multiple jurisdiction models. Findings are included from experiments designed to test the explanation offered in section 4 on simulated annealing. We also relate preliminary findings from two extensions of the basic model. The first extension assumes a cost of relocation and the second assumes that agents exit the economy with a small probability and are replaced by randomly created agents.

[17] An aspect of this second characteristic has been addressed formally by McKelvey (1986) who showed that the size of the uncovered set decreases with the level of symmetry of preferences.

[18] We have tried to balance the benefits of increasing the number of agents and costs in computation time of increasing the number of jurisdictions. In the future, we hope to complete more computations and to run more formal tests of robustness.

5.1 Sequence of Events

The computational model begins with a procedure in which agents' preferences are created and in which agents are assigned randomly to jurisdictions, with each agent being assigned to each jurisdiction with equal probability. We next begin a series of relocation decisions by agents and local public decisions by the political institutions until agents have had ten opportunities to relocate. For some political institutions, an equilibrium will have been attained at this point. Democratic referenda almost always settles into an equilibrium within four relocations for the parameter values used. The other institutions show no significant increases in aggregate utility after the first ten relocations.

Following standard Tiebout models we assume that an agent moves to the jurisdiction providing her the highest utility.[19] The choice of an LPDA by a political institutions can be relatively complicated for the institutions other than democratic referenda. Two–party competition begins with two randomly created LPDAs each representing a party. The parties alternate adapting their LPDAs with the hopes of winning election. In the hill–climbing algorithm, which was used in the findings summarized below, a party was allowed five sets of eight platform adaptations. In each adaptation, the party tests a LPDA which differs from its status quo LPDA on at most three of the eleven local public issues. If the new LPDA receives more votes, then it becomes the new status quo. At the completion of the eight LPDA adaptations, another party adapts. After each party has adapted five times, an election is held and the final LPDA is determined.

For the multiple party models, the respective parties also alternate updating. Each party updates its random initial LPDA with the intent of maximizing its vote total. So that the comparisons would be meaningful, we allowed the same number of adaptations per party as in two–party competition. Thus, in the findings shown below, each had five opportunities to conduct eight platform adaptations using the exact same hill–climbing algorithm used in two–party competition. Although the algorithms for parties adapting in response to Borda rule and proportional representation are identical, adaptive party behavior differs under the two institutions because of how voters' preferences translate into vote totals. An adaptation which may result in more votes according to proportional representation, which only measures how many voters most prefer a party, may result in fewer votes than under Borda rule.

5.2 Single Jurisdiction Findings

For the most part, the findings in the single jurisdiction model are not surprising. For aggregate utility, democratic referenda performs best, followed by Borda rule, two–party competition, and proportional representation, in that order.[20] Democratic referenda

[19] If two or more jurisdictions provide identical utility to an agent, the agent chooses the jurisdiction with the lower index unless her current jurisdiction provides maximal utility in which case she remains there.

[20] In one case Borda outperforms democratic referenda.

produces an outcome of the median on each LPI, which produces nearly maximal utility given our assumptions on the distribution of preferences. The one unexpected finding is that Borda rule outperforms two–party competition and proportional representation regardless of the number of parties. In all three institutions involving parties, the resulting LPDAs are significantly better than randomly generated LPDAs. The characteristics of each institution differ considerably. In two–party competition, parties quickly adapt LPDAs of high aggregate utility. Once the parties have located a good region of LPDAs, they tend to wander within that region, alternating which party wins. Under both Borda rule and proportional representation, a party may benefit from appealing to a faction of voters and adapting an LPDA with low aggregate utility. This tendency to adapt less representative LPDAs would appear to undermine the performance of both institutions. Yet, given that the LPDA implemented usually does not belong to a single party but, instead, is decided upon by weighted majority rule, the unrepresentative LPDAs of the individual parties may combine to form a LPDA of high aggregate utility.

The above discussion partially explains how Borda rule and proportional representation can result in LPDAs of relatively high aggregate utility. The question remains as to why Borda rule outperforms proportional representation in the single jurisdiction model. One reason may be that Borda rule promotes greater stability than proportional representation. If a party makes a minor alteration in its LPDA under proportional representation, it may create a large enough shift in votes to change the LPDA resulting from weighted voting among the parties. Under Borda rule, a minor alteration in a party's LPDA may change its ranking for many people, but the effect on relative vote totals will not be very large. An additional reason Borda rule results in better outcomes is that under Borda rule parties care about how they are ranked, not just how many agents prefer them most, dampening their incentives to adapt LPDAs which appeal to a small fraction of the agents.

Table 1 contains numerical findings from fifty trials of the single jurisdiction model with one–thousand agents and eleven LPIs. We have normalized utilities so that the expected maximal utility to an agent equals one-hundred and the expected utility to an agent of a randomly selected platform equals zero.[21]

[21]In the computations generating the findings shown, the adaptive parties used a multi–step hill–climbing algorithm. Similar results obtain when parties used a random search algorithm or a genetic algorithm.

Table 1

Single Jurisdiction: Utility

Institution	Agg. Utility	(s.d.)
Democratic Ref.	2.49	(1.60)
Two–party Comp.	1.84	(1.74)
Borda Rule (3 parties)	1.92	(1.62)
Borda Rule (4 parties)	2.45	(1.78)
Borda Rule (7 parties)	2.59	(1.67)
Prop. Rep. (3 parties)	1.02	(1.73)
Prop. Rep. (4 parties)	1.05	(1.82)
Prop. Rep. (7 parties)	1.29	(1.69)

Both Borda rule and proportional representation improve significantly as the number of parties increases from three to seven.[22] With seven parties, Borda rule yields the highest aggregate utility of any political institution. Intuition for this finding is straightforward. As the number of parties increases, each agent has a party which advocates a LPDA near their utility maximizing LPDA. In the limit, as the number of parties equals the number of agents, these types of institutions become equivalent to democratic referenda.

5.3 Multiple jurisdictions

With multiple jurisdictions, agents can sort according to their preferences. Tiebout's theory predicts that aggregate utility will increase with the number of jurisdictions. We present findings from models with three, seven and eleven jurisdictions, and for all four institutions considered, aggregate utility increases with the number of jurisdictions. The increases from the one jurisdiction model to the three jurisdiction model are most dramatic. Also, the *performance of the political institutions nearly reverses.* Democratic referenda, which performed best in the one jurisdiction model, now yields the lowest aggregate utility and proportional representation now performs second best. The disparity in performance appears to increase with the number of jurisdictions.

[22]Note that here, and elsewhere, the terms in parentheses in the table are the standard errors of the distribution, not of the mean.

Table 2

Multiple Jurisdictions: Utility

Institution	3 Loc's	(s.d.)	7 Loc's	(s.d)	11 Loc's	(s.d.)
Democratic Ref.	34.40	(2.31)	48.51	(2.04)	55.64	(2.05)
Two–party Comp.	34.74	(1.83)	49.35	(1.74)	56.64	(1.96)
Borda Rule (4 parties)	35.63	(1.54)	52.05	(1.67)	59.12	(1.55)
Prop. Rep. (4 parties)	35.53	(1.53)	51.41	(1.67)	58.51	(1.69)

Earlier in the paper we described how, with multiple jurisdictions, two–party competition might outperform democratic referenda because it would induce more sorting of agents. Configurations of agents which had relatively low aggregate utility would be less likely to be stable with respect to two–party competition. A small policy change might lead to a small migration, which in turn might lead to yet another policy change and still more migration, until eventually a stable configuration has been located. One approach to capturing the extent of sorting is to keep track of the total number of agent relocations for each political institution. Table 3 shows that two–party competition, as well as Borda rule and proportional representation, results in significantly more relocations than democratic referenda with three, seven, and eleven jurisdictions.

Table 3

Number of Agent Relocations

Institution	3 Loc's	(s.d.)	7 Loc's	(s.d)	11 Loc's	(s.d.)
Democratic Ref.	864.16	(121.90)	863.2	(73.55)	887.3	(37.22)
Two–party Comp.	915.68	(103.71)	1162.5	(61.03)	1293.7	(50.43)
Borda Rule (4 parties)	1816.18	(247.95)	2525.82	(272.80)	2753.2	(216.26)
Prop. Rep. (4 parties)	1237.48	(191.52)	1623.6	(220.88)	1762.8	(209.48)

The findings in Table 3 agree with our proposed explanation. As an additional test of whether the utility differences stem from improved sorting, we measure the average Hamming distance between the LPDAs across jurisdictions.[23] To make the findings

[23]We also measured the variance of agents' preferences in each jurisdiction and found an identical ranking of the institutions. Democratic referenda results in the most variance, followed by two–party competition, then proportional representation. Borda rule results in the least variance in preferences, which might be interpreted as the most homogeneous preferences at each jurisdiction.

more understandable, we normalize the Hamming distances, reporting the *percentage of maximum possible Hamming distance.* In this way, normalized Hamming distance lies between zero and one hundred, and with eleven LPIs, if two LPDAs disagree on exactly 3 issues then their normalized Hamming distance equals three–elevenths, or 28%.

Table 4

Normalized Hamming Distance Between LPDAs

Institution	3 Loc's	(s.d.)	7 Loc's	(s.d)	11 Loc's	(s.d)
Democratic Ref.	61.58	(5.92)	51.64	(2.28)	50.76	(1.65)
Two–party Comp.	65.21	(2.89)	53.96	(1.64)	52.26	(0.90)
Borda Rule (4 parties)	66.18	(2.06)	56.59	(0.80)	53.98	(0.49)
Prop. Rep. (4 parties)	66.67	(2.00)	56.16	(0.99)	53.78	(0.53)

In all but one case, Borda Rule results in the most heterogeneous LPDAs followed by proportional representation, two–party competition, and democratic referenda respectively. This ranking agrees with the utility rankings and the number of relocation rankings. More relocations and greater LPDA heterogeneity imply better sorting and, therefore, higher aggregate utility.

Average LPDA heterogeneity decreases as the number of jurisdictions increases. This was expected. Two LPDAs may have a heterogeneity equal to one, but there does not exist a triple of LPDAs such that all subsets of size two have heterogeneity equal to one, so the maximal possible average heterogeneity decreases. In general, as the number of jurisdictions increases, the LPDAs become more crowded in the space of possible LPDAs reducing their average heterogeneity.

Much of our argument as to why the institutions involving party competition outperform democratic referenda hinges upon the assumption that outcomes from these institutions are less stable, i.e., the LPDAs often change when agents preferences have a high degree of heterogeneity. We now test whether this in fact occurs. To do so we create populations of agents of varying degrees of heterogeneity. We then measure the number of changes in LPDAs over time and verify whether the number of LPDA changes increases with the heterogeneity of preferences.

To create populations of agents with specified levels of heterogeneity, each agent's ideal point was made a convex combination of a *base preference* and an *individual preference.* Both an agent's individual preference and the common base preference were drawn from the same distribution as agents' preferences in the basic model. We can write an agent's preferences as

$$(1-\theta) \cdot base \; + \; \theta \cdot individual.$$

We refer to θ as the *degree of heterogeneity.* In the findings shown, we examine a one jurisdiction model with two hundred and fifty agents. The restriction to one jurisdiction guarantees that none of the changes in LPDAs are attributable to the shifting population. We vary θ between 0.5 and 1.0. Values of θ lower than this create agents whose preferences are too homogeneous: almost all have the same preferred platform. The findings presented below are from fifty series of ten elections. *Flips* equals the number of LPDs which changed in each election.

Table 5

Flips as a Function of Heterogeneity

θ	Dem Ref		Pro. Rep.		Borda		Two–Party	
0.50	0.00	(0.00)	0.69	(0.38)	0.81	(0.27)	0.24	(0.15)
0.55	0.00	(0.00)	0.80	(0.33)	0.90	(0.25)	0.21	(0.16)
0.60	0.00	(0.00)	0.80	(0.37)	0.98	(0.29)	0.23	(0.24)
0.65	0.00	(0.00)	0.77	(0.40)	1.01	(0.29)	0.26	(0.24)
0.70	0.00	(0.00)	0.88	(0.48)	1.03	(0.31)	0.31	(0.37)
0.75	0.00	(0.00)	1.01	(0.53)	1.13	(0.35)	0.38	(0.44)
0.80	0.00	(0.00)	1.26	(0.55)	1.13	(0.36)	0.53	(0.49)
0.85	0.00	(0.00)	1.94	(0.69)	1.23	(0.45)	0.68	(0.73)
0.90	0.00	(0.00)	2.84	(0.71)	1.51	(0.51)	1.37	(1.11)
0.95	0.00	(0.00)	3.78	(0.62)	2.14	(0.68)	2.65	(1.08)
1.00	0.00	(0.00)	4.02	(0.59)	2.48	(0.76)	3.36	(1.12)

Democratic referenda shows no flips for all degrees of heterogeneity. This occurs because with a fixed population, democratic referenda results in a unique outcome, and in this scenario no agents move into or out of the jurisdiction. For the other three institutions we find, as expected, that the average number of flips increases with the degree of heterogeneity. These findings strongly support our assumption that as preferences in a jurisdiction become more homogeneous the propensity for LPDAs to wander decreases.

The final test required to substantiate our argument is whether annealing improves the performance of democratic referenda. To test annealing we introduce a probability with which agents chose to live in a suboptimal jurisdiction. We lower that probability over time to create an annealing effect. The average utility increased in almost every case. For example, in an eleven jurisdiction model, when the initial probability of moving to a suboptimal jurisdiction was five percent and where this probability decreases by five percent each period for ten periods, aggregate utility increases from 55.35 to 55.61, which is significant at a 95% level. These findings suggest that noise alone may be enough to improve outcomes. The findings above, which suggest that the error rate is correlated with the aggregate utility of the configurations, lend support to the notion that noise, or small mistakes, may be beneficial.

5.4 Relocation Costs

Relocation costs reduce the number of agent relocations, and as a result reduce the amount of sorting. The inclusion of relocation costs can affect the absolute and relative performance of the political institutions we consider. At the extreme, if relocation costs are so high that no agent would ever want to relocate, then a multiple jurisdiction model would be equivalent to multiple single jurisdiction models. Institutional performance in the one jurisdiction model and the multiple jurisdiction model would therefore be equivalent. The relevant question is whether small relocation costs have proportional effects or large effects. We find that that the effects are roughly proportional. Small increases in relocation costs have correspondingly small effects on the number of agent relocations and on the ability of agents to sort.

Tables 6 and 7 presents findings from a seven jurisdiction model in which the cost of relocation equals one divided by the number of LPIs, $\frac{1}{11}$, and in which the cost of relocation equals two divided by the number of LPIs, $\frac{2}{11}$, respectively.

Table 6

Relocations Costs = 1/(number of LPIs)

Institution	Utility	(s.d.)	Reloc's	(s.d.)	LPDA Het.	(s.d.)
Democratic Ref.	47.54	(2.04)	762.90	(43.86)	50.84	(2.36)
Two–party Comp.	49.13	(1.87)	1007.60	(54.40)	54.10	(1.35)
Borda Rule	51.02	(1.68)	1828.20	(186.72)	56.23	(0.77)
Prop. Rep.	50.65	(1.73)	1237.56	(131.97)	56.23	(0.71)

Table 7

Relocations Costs = 2/(number of LPIs)

Institution	Utility	(s.d.)	Reloc's	(s.d.)	LPDA Het.	(s.d.)
Democratic Ref.	46.54	(1.99)	681.88	(29.91)	51.05	(2.37)
Two–party Comp.	48.02	(2.03)	862.66	(48.15)	54.13	(1.57)
Borda Rule	49.73	(1.97)	1358.38	(105.97)	55.97	(1.10)
Prop. Rep.	49.70	(1.87)	1031.22	(91.68)	56.07	(0.97)

The ranking of the institutions according to utility and number of movements is identical in the two cases. Moreover, the number of relocations and aggregate utility decreases. Differences in heterogeneity are insignificant.

5.5 Probabilistic Replacement of Agents

We now examine a model in which agents leave the global economy with a small probability and are replaced by agents with randomly drawn preferences. In Table 8, we present findings from a seven jurisdiction model in which each agent is replaced with probability equal to 0.02. Rather than run the model for just ten iterations, we experimented with longer time horizons. The findings below are taken from a model with only forty relocations as we found that increases beyond forty or fifty did not appear to have significant effects.

Table 8

Probability of Replacing Agents = 0.02

Institution	Utility	(s.d.)	Reloc's	(s.d.)	LPDA Het.	(s.d.)
Democratic Referenda	48.32	(1.56)	1332.0	(90.94)	52.71	(2.16)
Two–party Competition	48.82	(1.28)	1641.7	(96.47)	54.39	(1.31)
Borda Rule	51.62	(1.38)	5321.0	(399.11)	56.52	(0.66)
Prop. Rep.	51.11	(1.23)	2841.0	(348.40)	56.73	(0.56)

The relative rankings of the institutions are the same as in the seven jurisdiction model without replacement. In the probabilistic replacement scenario, the two multiple party institutions result in massive numbers of relocations relative to the other two institutions. Not unexpectedly, the aggregate utility and level of LPDA heterogeneity is much higher for these two institutions as well.

6 Discussion

This paper makes the claim that in a model of Tiebout competition some institutions perform better than others in that they induce more sorting and generate on average higher aggregate utility. Ironically, these institutions often perform relatively poorly in a single jurisdiction model. The reason for this performance reversal is that the multiple location model is a complex environment with many possible stable configurations. A minor mistake, which is harmful in the single jurisdiction model, may be beneficial in the multiple location model in that it can dislodge the system off a relatively bad local optima. In our model of Tiebout competition, instability in policies causes people to leave jurisdictions they do not like and to join jurisdictions they do like. Occasional wrong moves by people can lead to an avalanche of moves by others, leading to better outcomes (better sorting) in the end.

In the multiple jurisdiction model, aggregate utility, the number of relocations, and the heterogeneity of LPDAs were highly correlated. Institutions which generate high aggregate utility induce more relocations and result in more heterogeneous LPDAs. We also find that the number of policy flips increases with the degree of heterogeneity of preferences for the electoral institutions considered. Thus, our findings agree with the annealing argument put forth in section 4 as to why institutions which perform poorly with only one jurisdiction might perform well with multiple jurisdictions.

There are many strong assumptions in the current model. The parties lack policy preferences, and voters are not strategic. If we allow parties with policy preferences, then outcomes may even be less representative then under the current model. The relevant issue is whether the number of policy flips will be positively correlated with the degree of heterogeneity of preferences. We see no reason why this would not be the case. Strategic voting, on the other hand, may dampen this effect. With strategic voters, outcomes from Borda rule and proportional representation may become more representative, reducing the number of policy flips which occur when agents have heterogeneous preferences.

The general picture that emerges from our work is quite intriguing. Tiebout models are just one example of a broad class of phenomena that must "sort" agents among alternative configurations (other examples include models of coalition formation and organization). Under decentralized sorting mechanisms, these systems can get trapped in suboptimal configurations. If, however, there are means by which these poor configurations can be annealed in an appropriate manner, then the global system can escape these configurations and achieve superior outcomes. We find that certain political institutions act as natural annealing mechanisms, and this insight provides an important link for understanding the dynamical behavior and ultimate performance of such institutions.

A previous version of this paper was presented at the annual meetings of the American Political Science Association, New York City, 1994. A portion of this work was funded by NSF grants SBR 94–09602, SBR 94–11025, and SBR 94– 10948

References

Arthur, W. B. 1989. "Competing Technologies, Increasing Returns, and Lock–in by Historical Events" *Economic Journal* 99 pp 116- 131.

R. Axelrod, and D. S. Bennett, 1993. "A Landscape Theory of Aggregation," British Journal of Political Science, 23, pp 211–233.

Bates, Robert H. 1990. "Macropolitical Economy in the Field of Development." in *Perspectives on Positive Political Economy*, ed. James Alt and Kenneth Shepsle. New York. Cambridge University Press.

Bewley, T.F. 1981. "A Critique of Tiebout's Theory of Local Public Expenditures," *Econometrica* 50 1183-1205.

Cebula, R.J. 1974 "Interstate Migration and the Tiebout Hypothesis: An analysis according to race, sex, and age." *Journal of the American Statistical Association* 69 pp 876-879.

DeVany, A. 1994. "Hard Cores, Soft Cores, and The Emergence of Self-Organized Coalitions", UCI Working Paper.

Dowding, K. John, P. and Biggs, S. 1994 "Tiebout: A Survey of the Political Empirical Literature." *Urban Studies* 31 pp 767-797.

Durlauf, S. 1992 "Multiple Equilibria and Persistence in Aggregate Fluctuations," *AEA Papers and Proceedings* pp 70-74.

Fox, W.F. Herzog, H.W. Schbottman, A.M., 1989. "Metropolitan Fiscal Structure and Migration." *Journal of Regional Science* 29 pp 523-536.

Holland, John H. 1975. *Adaptation in Natural and Artificial Systems.* Ann Arbor, University of Michigan Press.

Holland, John H. and John H. Miller. 1991. "Artificial Adaptive Agents and Economic Theory." *American Economic Review, Papers and Proceedings.* 81:365–70.

Hoyt, W.H. 1990 "Local Government Inefficiency and the Tiebout Hypothesis: Does Competition among Municipalities Limit Local Government Inefficiency." *Southern Economic Journal* 481-496.

Jacobs, Jane. 1984. *Cities and the Wealth of Nations.* Random House, New York.

Kollman, Ken, John H. Miller, and Scott E. Page. 1992. "Adaptive Parties in Spatial Elections." *American Political Science Review.* 86:929–37.

Kramer, Gerald. 1977. "A Dynamical Model of Political Equilibrium." *Journal of Economic Theory* 15:310-334.

Krugman, P. 1991 "Increasing Returns and Economic Geography," *Journal of Political Economy* 99 pp 483-499.

McKelvey, Richard D. 1976. "Intransitivities in Multidimensional Voting Models and Some Implications for Agenda Control." *Journal of Economic Theory.* 12:472–482.

McKelvey, Richard D. 1986. "Covering, Dominance, and the Institution-Free Properties of Social Choice." *American Journal of Political Science* 30:283-314.

Niskanen, W.A. 1971 *Bureaucracy and Representative Government* Chicago, Aldine–Atherton.

Percy, S.L. 1993. "Revisiting Tiebout: moving rationales, exiting behavior, and governmental responses to metropolitan mobility." Presented at the Midwest Political Science Association, Chicago, IL.

Percy, S.L. and Hawkins, B.W. 1992. "Further tests of individual –level propositions from the Tiebout model." *Journal of Politics* 54 pp 1149-1157.

Plott, Charles. 1967. "A Notion of Equilibrium and Its Possibility under Majority Rule." *American Economic Review* 79:787-806.

Reschovsky, A. 1979. "Residential choice and the local public sector: An alternative test of the Tiebout hypothesis." *Journal of Urban Economics*,16 pp 501-520.

Romer, P. and Rosenthal H. 1979. "Bureaucrats Versus Voters: On the Political Economy of Resource Allocation by Direct Democracy." *Quarterly Journal of Economics* pp 563-587

Samuelson, P. 1954 "The Pure Theory of Public Expenditure." *Review of Economics and Statistics*, 36, pp 387-389.

Schelling, T. 1978 *Micromotives and Macrobehavior*, Norton, New York.

Sharp, E.B. 1984. "Exit, voice and loyalty in the context of local government problems," *Western Political Quarterly* 37 pp 67- 83.

Taagepera, Rein, and Shugart, Matthew. 1989. *Seats and Votes.* New Haven, Yale University Press.

Tiebout, C.M. 1956 "A pure theory of local expenditures" *Journal of Political Economy* 64 pp 416-424.

Tullock, G. 1965 *The Politics of Bureaucracy* Washington DC: Public Affairs Press.

Zax, J. 1989 "Is there a Leviathan in your neighbourhood?" *American Economic Review* 79 pp 271-283.