

A Model of Farsighted Voting

Elizabeth Maggie Penn

October 30, 2007

Abstract

I present a new method of interpreting voter preferences in settings where policy remains in effect until replaced by new legislation. In such settings voters consider not only the utility they receive from a given policy today, but also the utility they will receive from policies likely to replace that policy in the future. The model can be used to both characterize long-term preferences and distributions over policy outcomes in situations where policy is ongoing and voters are farsighted.

“[W]e’re unprepared for the consequences of winning. Winning in court too soon could mean losing in the court of public opinion, in Congress and under the United States Constitution.”

—Matt Coles, Director of the ACLU Lesbian and Gay Rights Project, quoted in The New York Times on November 12, 2004. [14]

1 Introduction

Following the 2004 elections in which eleven states passed constitutional amendments banning gay marriage, gay rights groups began a concerted public effort to change their political strategy. The groups decided to temporarily retreat from what they had previously claimed to be their most important goal: winning the right for same-sex marriages. The reason for this move was simple. The groups realized that a challenge to such amendments in federal court was a risky proposition, regardless of whether the challenge was successful or not. An unfavorable court decision would establish a legal precedent that could be invoked in future cases. A favorable court decision could easily accelerate public and legislative support for a federal constitutional amendment banning same-sex marriage outright.

Political choice in the short term often involves long-term considerations because, as the above story demonstrates, decisions made today can greatly affect the types of decisions that are feasible tomorrow. Given the importance that individuals place on long-term outcomes, an essential step in understanding strategic voting behavior is understanding the future consequences of policy choice. These consequences are particularly relevant when policies remain in effect for an indefinite period of time, or are continuing. A continuing program, as opposed to a once-and-for-all program, is a policy that continues in effect until it is changed by new legislation.¹ When such a policy is enacted we can think of it as producing a path-

¹The terms “continuing program” and “once-and-for-all program” are used in Baron [3].

dependent stream of future legislation, with the status quo at any given time being the extant legislation at that time. Because of the long-term nature of these programs they are often of particular interest to legislators, activists and voters. Such programs include entitlements, social policy, and redistributive programs, with specific examples being minimum wage laws, social security benefits, the regulation of public health concerns such as air quality and automobile safety, and eligibility requirements for social welfare programs.

Decisions over continuing programs frequently affect the future choices that are available to a group; for example, it can be politically difficult to reinstate a tax once it has been repealed, to lower levels of entitlement spending, or to revoke a law concerning public safety. When legislators evaluate such policies they do so with an eye toward what the policy is likely to produce over time. There is therefore no reason to expect that legislative bargaining over a continuing program will be similar to legislative bargaining over a once-and-for-all program. While the behavior of a legislator may depend largely on his immediate preferences, it may also depend on how he sees today's choice as affecting future decisions. This tension between short and long-term interests is my primary focus.

In this paper I develop a model of bargaining over continuing programs in which individuals rank policies not only on the basis of the utility they yield today, but also with respect to the types of alternatives they will likely lead to in the future. The model differs from much of the bargaining literature in that the bargaining process does not end once a policy has been agreed upon, as in Baron and Ferejohn [4]. Instead, the chosen policy becomes the reversion point of the next round of bargaining, and remains in effect until it is replaced by a new alternative.² The model is used to both characterize voting behavior when individuals are farsighted and to characterize the types of policies that will emerge when programs are continuing.

²Baron [3], Kalandrakis [12], Riboni [20], and Bernheim et al. [7] also examine repeated bargaining games with endogenous reversion points. These papers will be discussed in Section 2.1.

The theory developed here captures several important features of politics that may not be captured in a one-shot model of bargaining. First, farsighted voters are not indifferent between different policies that provide them with the same level of utility. This is because the space of alternatives that defeat each policy, and that each policy defeats, are substantively different. This model demonstrates that in dynamic environments, the space of alternatives that can and cannot defeat a policy may have as much impact on individual decisionmaking as the substance of the policy itself. And second, farsighted voters will take the preferences of others into account when voting, not because of a behavioral assumption such as altruism or inequality aversion, but because the preferences of others will matter when passing future legislation. This implies that changing the preferences of a single voter will change the strategic behavior of every other voter in a predictable way. While both of these features hold true in any dynamic game, the goal of this paper is to not only characterize individual behavior on an equilibrium path, but also to characterize farsighted individual preferences in general. Most importantly, *this model calculates a general induced utility function for each player that captures his evaluation of every policy in terms of what it is likely to produce over time.*

The paper begins with several general existence results. I then apply the model to a number of standard legislative settings, including a one-dimensional spatial model, a divide-the-dollar game, and a two-dimensional spatial model in which players have either circular or elliptical preferences. I find that the theory is capable of making sharp predictions in each of these legislative settings, with the predicted outcomes frequently corresponding to policies that divide benefits fairly between like-minded individuals. Thus, while the model is intended specifically to analyze individual vote choice, it could also be interpreted as a model of endogenous party, or coalition, formation. The coalitions that tend to emerge consist of those individuals whose preferences are most similar, either in terms of spatial distance or intensity of preference. Furthermore, farsighted individuals tend to favor policies that yield equitable distributions of payoffs, and frequently

vote for certain normatively “fair” alternatives over their own ideal points. Relatedly, farsighted voters may shy away from implementing an inefficient Condorcet winner in favor of implementing a more efficient cycle of alternatives over time; thus farsighted voters may willingly concede utility in the short term for more beneficial policy outcomes later on.

The paper proceeds as follows: Section 2 describes the notation used and presents the model. Section 2.1 describes the path of play in more detail and discusses related literature. Section 2.2 presents the notion of a dynamically stable voting equilibrium, the equilibrium concept used in the analysis of the model. Section 3 presents several results about equilibrium existence that constitute the main theoretical contribution of the paper. When the policy space is finite I show that there always exists an equilibrium. When the policy space is infinite then there exists an equilibrium under certain conditions, and when the number of players is large the equilibrium is unique. Furthermore, I demonstrate that equilibrium behavior in this model is consistent with Bayesian Markov-perfect equilibrium behavior. Section 3.1 provides two analytic examples of the model in the setting of a finite policy space, and shows that Condorcet winners may not arise as policy outcomes when voters are farsighted. Section 4 discusses the specific applications of the model in greater detail and presents numerical results concerning these applications in continuous policy spaces. Section 5 concludes.

2 The Model

There is a collection of *voters* $N = \{1, 2, \dots, n\}$ and a compact set $X \subset \mathbb{R}^m$ of *alternatives*, or *policies*. X can be either finite or infinite. For each voter $i \in N$, preferences are represented by a real-valued *utility function*, $u_i : X \rightarrow \mathbb{R}$. When the set of policies is infinite I also assume that these utility functions are differentiable, and that their derivatives are uniformly bounded by some constant U .

A subset $C \subseteq N$ is called a *coalition* of voters. Coalitions that are large enough to enact a policy are called *decisive*, or *winning*. The collection of winning coalitions is W . It is simply assumed that this collection W is *monotonic* and *proper*. Monotonicity implies that if $C \in W$ and $C \subset C'$, then $C' \in W$, or that adding people to a decisive coalition yields another decisive coalition. Properness implies that if $C \in W$, then $N \setminus C \notin W$, or that any two decisive coalitions must have at least one individual in common.

2.1 Path of Play and Related Literature

The legislative process is modeled as a sequence of sessions in which votes on policy occur. In each round the status quo is determined by the bargaining outcome of the previous round. A policy to be pitted against the status quo arises exogenously in each round. This policy is drawn from a probability density Q . A vote then occurs between this policy and the status quo. Every voter receives a payoff from the winning policy and this policy then becomes the status quo of the next round of bargaining. This process is pictured in the figure below.

[FIGURE 1 HERE]

The probabilistic and exogenous nature of the proposal process distinguishes this model from Baron [3], Kalandrakis [12], Riboni [20], and Berhneim et al. [7], and warrants a short discussion. Both Baron and Kalandrakis analyze infinitely repeated bargaining games with endogenous reversion points. In these papers a player is selected at random to make a proposal in each round. The proposal is pitted against the status quo, with the winner becoming the status quo in the next round of bargaining. Baron looks at the case of a unidimensional policy space, and finds that outcomes converge to the ideal point of the median voter. Kalandrakis examines a three-player divide-the-dollar game and finds that a Markov perfect equilibrium of the game is characterized by a situation where the proposer in every round proposes the entire dollar

for himself, and this allocation is approved by a majority of players. Riboni considers a similar model in which policy changes are proposed by a strategic agenda setter (who is not a member of the voting body), and then voted upon by a committee. Bernheim et al. also consider a similar model in which legislators are recognized sequentially to make proposals and a finite number of proposals are made. Under weak conditions they find a result similar in spirit to Kalandrakis's, in which the final proposer receives the entire dollar in a divide-the-dollar setting.³

The focus of these papers differs from my paper in an important respect. While these papers are concerned primarily with predicting equilibrium policy outcomes and voter behavior on the path of play, the goal of this paper is to model voter preferences over *all* alternatives when voters evaluate policies in terms of what they are likely to produce over time. By separating voting strategies from the proposal process, results about farsighted preferences and behavior are easier to interpret. For example, in Penn [18] I endogenize the proposal process in this model by allowing players to make proposals themselves. The results I find are similar to those found by Kalandrakis and Bernheim et al. Thus, Kalandrakis's result that a legislative dictator emerges with certainty in every round is likely an effect of the particular endogenous proposal process assumed.⁴

The assumption of an exogenous and probabilistic process by which future proposals are generated reflects the notion that legislators are not certain of the policy proposals that will be brought to the floor in the

³Also worth noting are recent papers by Battaglini and Coate [5, 6] that examine the dynamics of taxation and public spending. In these papers, policy-making periods are linked by endogenous levels of public goods and public debt, respectively.

⁴Kalandrakis has different work [13] that specifically examines the importance of proposal rights in determining political power. He obtains the interesting result that *any* distribution of power can be obtained by simply manipulating proposal rights in a particular bargaining environment. The same is not true when manipulating other institutional features of the bargaining game such as voting rights.

future, but are aware of the current political climate and have priors over the distribution of future proposals. While “endogenous” is frequently interpreted as being more realistic than “exogenous” in any game theoretic model, the assumption of an exogenous agenda here makes the model easier to interpret, easier to apply to different situations, and, I argue, more realistic. This is because an agenda in this setting will depend on a number of complex factors (recognition rules, special interest and constituency pressures, the ordering of the legislative calendar, party control, etc.) that are beyond the scope of this model. For example, in game-theoretic models it is generally assumed that legislators are randomly recognized to propose alternatives to the floor via a particular recognition rule. This assumption yields a very specific kind of strategic agenda-setting process that is separate from the goal of this paper, which is simply to model voter preferences in an environment where policies are replaced by new policies over time.

Roberts [21] and Compte and Jehiel [9] study similar bargaining problems with randomly generated offers. Roberts considers a setting where policymaking is continuing and every alternative is equally likely to be brought to the floor. Compte and Jehiel assume the bargaining process ends when a certain number of players accept the random offer. They examine the effect of patience, the number of players and the majority requirement, and find that as the majority requirement increases more efficient outcomes are generated, but that it also takes longer to reach agreement. They motivate their random proposal process with the argument that it is rare for any individual or group to have full control over an agenda-setting process, and that even if a person did have full control over proposals, it would be difficult for him to perfectly target a collection of specific payoffs for the other players. They also find that simplifying the proposal process in such a way enables them to better analyze many empirical regularities that have not been supported by previous bargaining problems; for example, the fact that agreement is frequently more difficult to obtain when unanimous consent is required.

2.2 Dynamically stable voting equilibria

This section presents the model of farsighted voter preferences. Utility function $u_i(x)$ represents Player i 's one-shot payoff from policy x when x is enacted today. However, if Player i is farsighted then he not only cares about his payoff from having x enacted today, but also his future payoffs from the policies that will ultimately replace x .

Let Player i 's *value function* be denoted v_i , where $v_i : X \rightarrow \mathbb{R}$. This function represents the discounted sum of utility Player i can expect to receive from having policy x enacted today, given that a stream of policies will be enacted after x . Let $v = \{v_i\}_{i \in N}$ be the collection of all voter value functions, with \mathcal{V}^n being the space of all value functions (so that $v \in \mathcal{V}^n$). In equilibrium each voter will assign a “true” long-term value to every policy. This means that players will vote based on their equilibrium value functions and by voting this way they ultimately generate those same functions.

In equilibrium, value functions capture a consistency between beliefs and behavior; when a player behaves according to such a function, the value he assigns to a policy equals the true future expected value of that policy. When this holds for *all* players, then the vote strategies of players generate value functions that generate those same vote strategies. Thus, beliefs and behavior are entirely consistent with each other. The following equilibrium concept captures this idea.

Definition: A *dynamically stable voting equilibrium* is a collection of value functions, $v^* = \{v_i^*\}_{i \in N}$, such that for all $i \in N$ and $x \in X$,

$$v_i^*(x) = u_i(x) + \delta \sum_{y \in X} [v_i^*(y)p(v^*(x), v^*(y)) + v_i^*(x)(p(v^*(y), v^*(x)))]Q(y).$$

The case of an infinite X is defined analogously.⁵

⁵As the value function depends on the distribution from which proposals are drawn, or Q , $v_i^*(x)$ could also be written $v_i^*(x|Q)$.

The function $u_i(x)$ equals the utility player i receives from alternative x in one period. The probability that policy y will defeat status quo x is denoted $p(v(x), v(y))$. This function is simply assumed to be the probability that a winning coalition votes for y over x , given that every voter knows that policy selection will continue into the future. Thus voting decisions are made based on voters' value functions and not their utility functions (i.e. voters are farsighted). The specific functional form of $p(v(x), v(y))$ can be found in Appendix A; it is simply the probability of victory of y over x .

Q is the probability mass from which alternatives y to be pitted against the status quo are drawn. As discussed in Section 2.1, Q represents the fixed beliefs that voters have over the types of alternatives that will be brought to the floor in the future. These beliefs could be uniform over all alternatives (uninformative), or could be generated by fixed external pressures from political parties, special interests, constituencies, or simply the current political climate. While Q is assumed to be independent of the current status quo x , a similar model could be constructed where Q is conditioned on x . This construction would not dramatically change the analytic results of the model, but is omitted for ease of exposition.⁶ When the set X is infinite Q is instead a density. In this case, Q is assumed to have full support, and to be differentiable in y .

$\delta \in [0, 1)$ is a discount factor that represents players' time preferences. When δ is high, voters place greater relative weight on the future. For notational simplicity δ is assumed to be common for all players, but this assumption does not affect any of the analytic results.

One feature of this model that differentiates it from an endogenous proposal model is that the probability of victory of y over x , or $p(v(x), v(y))$, is averaged over all decisive coalitions. Alternatively, when proposals are endogenous the proposer does two things: he chooses an alternative to propose, and he identifies

⁶It would require a stronger assumption than Assumption 2 in order to prove existence when X is infinite (Proposition 2).

a particular decisive coalition to target his proposal to. Therefore, when proposals are endogenous policy outcomes will typically change at every round of voting unless the status quo is the proposer's ideal point. However, in this model, because winning coalitions are no longer uniquely determined but instead are averaged over, a consequence is that universalistic, or close to universalistic, majorities can arise given particular policy pairings. Thus, policy will remain at more desirable status quos for longer, and undesirable policies will be replaced quickly.

To summarize this section, the above model characterizes long-term individual preferences in a setting where policies are repeatedly being challenged and replaced by new alternatives. In a given round a transition from status quo x to new policy y is dependent upon two factors. First, policy y must be chosen to be pitted against status quo x from density Q . Second, voters must choose policy y over status quo x , and the probability that y defeats x when infinitely more rounds remain is denoted $p(v(x), v(y))$. The function $v_i(x)$ represents i 's discounted expected sum of utility when x is enacted today, given that infinitely more rounds of policy selection will occur. This equals $u_i(x)$, or i 's utility from x today, plus the discounted expected value of what x will ultimately lead to in the future.

3 Analytic results and examples

In the sections that follow, I will provide examples of dynamically stable voting equilibria in specific settings. The goal of this section is to provide more general results about the types of environments in which we can expect dynamically stable voting equilibria to exist and to be unique. Recall that an equilibrium in this setting is a collection of value functions such that, when individuals vote according to these functions, individual valuations of policies equal the true future expected values of those policies. Therefore, we are looking for fixed points in a mapping from one value function into another. Mathematically, this is different

than a more standard game theoretic setup in which we look for fixed points in action profiles. The difference stems from the fact that while value functions (which are necessarily cardinal) will yield action profiles, action profiles do not provide us with enough information to give us back cardinal valuations. Therefore, stronger conditions are required in order to obtain equilibrium existence in this setup than in a standard game theoretic setup.⁷

Propositions 1, 2, and 3 all focus on the problem of equilibrium existence. Although the voting behavior of players is purposefully left unspecified in this model, individual behavior is implicit in the definition of the function $p(v(x), v(y))$, or the probability of victory of y over x . While p is defined formally in Appendix A, let the probability that individual i (with value function v_i) votes for y over x be denoted $p_i(v_i(x), v_i(y))$. The three existence results all require that this probability be a continuous function for all individuals, or that individuals vote probabilistically.

Many authors have invoked the assumption of probabilistic behavior on the part of voters (see [16], [22], [1], and [10]), with the implication being that models of behavior are incapable of perfectly predicting vote choice, and that this reality should be incorporated into the calculations of voters. This model is consistent with the more game-theoretic formulation of probabilistic voting adopted by McKelvey and Patty [16]. As in their setup, and opposed to some more classic models of probabilistic voting, the expected payoff of casting a vote is represented correctly for each voter.⁸ In other words, individuals vote with the knowledge that other voters are also behaving probabilistically. As in other recent work on the topic, the assumption of probabilistic voting in this model is consistent with pure strategy equilibrium behavior in a Bayesian

⁷In particular, the difficulties with equilibrium existence here are mathematically similar to difficulties in obtaining existence in games with continuous action spaces.

⁸Relatedly, McKelvey and Patty also require the size of the electorate to be “large enough” in order to characterize their equilibrium.

framework in which individuals receive privately observed payoff shocks for each possible action, or vote. This formulation of probabilistic voting, and in particular, the effect of these random payoff disturbances on individual vote functions p_i , is presented in Appendix A.

The three existence results require different and nested sets of conditions on individual vote functions p_i .⁹ Proposition 1 simply proves that when the alternative space X is finite and individuals vote probabilistically, an equilibrium exists. Proposition 3 proves that if individual vote functions are differentiable and the derivatives are uniformly bounded by *any* constant, then a unique equilibrium will exist when the number of players is large enough, regardless of whether X is finite or infinite. Proposition 2 proves existence when X is infinite for any number of players, but requires that the derivative of p be bounded by a particular constant. In Appendix A this constant is defined both generally (Assumption 2) and more specifically for the case of a logit agent quantal response equilibrium. All proofs of the propositions can be found in Appendix B.

Proposition 1 *If X is finite, then there exists a dynamically stable voting equilibrium when p is continuous.*

Proposition 2 *If X is infinite, then there exists a dynamically stable voting equilibrium when the derivative of p is bounded by a particular constant.*

Proposition 3 *When n is large then there always exists a unique dynamically stable equilibrium when the derivatives of p_i are uniformly bounded, regardless of whether X is finite or infinite.*

To understand why additional requirements on transition probabilities p are needed to prove existence when X is infinite, note that the set of all functions over a finite alternative space is a vector space, while the set of all functions over an infinite and compact subset of \mathbb{R}^m is a function space. Compactness of the space

⁹Alternatively, as discussed in Appendix A, these assumptions can be thought of as requirements on the distribution of individuals' privately observed random payoff disturbances.

of value functions is generally needed in order to find a fixed point. While any closed and bounded subset of a finite-dimensional vector space is compact, closed and bounded sets of functions are rarely compact. Proposition 2 is proved by showing that the set of value functions is equicontinuous, and thus compact, when a certain restriction (Assumption 2 described above) on the derivative of transition probability p holds.

Proposition 3 proves existence differently, by showing that when the total number of players is sufficiently large we can construct an iterative definition of a dynamically stable voting equilibrium that is a contraction mapping. This definition (Equation 2 in Appendix B) is also used to perform the numerical estimations that follow in Section 4, and is discussed in greater detail in that section. Furthermore, this definition can be easily utilized to calculate farsighted evaluations of policy when only a finite number of rounds of policymaking will occur. The extension of this model to the setting of a finite number of rounds is discussed in Appendix B.

Last, it is important to note that under different specifications of the functions p_i a dynamically stable voting equilibrium is equivalent to other commonly used equilibrium concepts. For the final proposition and corollary of the paper I construct a game, Γ , in which equilibrium behavior is consistent with behavior generated by dynamically stable voting equilibria. In particular, I show that when individuals vote probabilistically, behavior in a dynamically stable voting equilibrium is equivalent to pure strategy Bayesian Markov-perfect Nash equilibrium behavior in the constructed game (Proposition 4). As a corollary it follows that when individuals vote deterministically, so that $p_i(v_i(x), v_i(y)) = 1$ if $v_i(y) \geq v_i(x)$ and zero otherwise, then at a dynamically stable voting equilibrium the collection of functions p_i constitute a Markov-perfect equilibrium. In both cases, the v^* vector represents the expected utility functions of the players, and strategies as specified by the functions p_i are consistent with the maximization of these expected utility functions. The proofs and the construction of Γ are relegated to Appendix B.

Proposition 4 *If individuals vote probabilistically and payoff disturbances are admissible, then at a dynamically stable voting equilibrium, v^* , the collection of functions p_i are consistent with behavior in a pure strategy Bayesian Markov-perfect Nash equilibrium of a game, Γ .*

3.1 A One-Dimensional Example: The Federal Marriage Amendment and Gay Rights

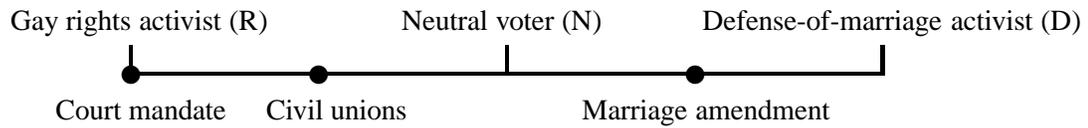
When the policy space is finite and small it is not difficult to solve for equilibria analytically. In this section and the next I will present simple analytic examples of dynamically stable voting equilibria. This first example depicts a one-dimensional spatial model, the second depicts a setting in which there is a majority preference cycle over a subset of the alternative space. In both examples I assume that discount factor $\delta = .9$ and that voting is deterministic, with

$$p_i(v_{it}(x), v_{it}(y)) = \begin{cases} 1 & \text{if } v_{it}(y) > v_{it}(x) \\ 0 & \text{if } v_{it}(y) < v_{it}(x) \\ \frac{1}{2} & \text{otherwise.} \end{cases}$$

To motivate this first example, consider the story presented in the introduction about the political strategies of gay rights groups after the 2004 elections. Suppose that there are three political actors in the model: a gay rights activist (R), a defense-of-marriage activist (D), and a neutral voter (N). Also, suppose that there are three possible political outcomes: a court-mandated overturn of state constitutional amendments banning gay marriage (“court mandate,” for short), legalization of civil unions and benefits for same-sex partners (“civil unions”), and a federal constitutional amendment banning gay marriage (“marriage amendment”).

Last, assume that the “marriage amendment” outcome is currently being strongly forwarded by special interests, and is three times more likely to arise as a policy proposal than the other two policy alternatives. Thus, $Q(\text{Marriage amendment}) = \frac{3}{5}$, and $Q(\text{Civil unions}) = Q(\text{Court mandate}) = \frac{1}{5}$.

The following figure depicts the hypothetical spatial location of the ideal points of the three players and the locations of the three policies.



The above figure generates the following two tables, which show the utility functions of the three players and the sum of expected utility each policy yields in the long term, at a dynamically stable voting equilibrium.

One-shot Utility

i	$u_i(\text{Court mandate})$	$u_i(\text{Civil unions})$	$u_i(\text{Marriage amend.})$
R	1	$\frac{3}{4}$	$\frac{1}{4}$
N	$\frac{1}{2}$	$\frac{3}{4}$	$\frac{3}{4}$
D	0	$\frac{1}{4}$	$\frac{3}{4}$

Farsighted (Equilibrium) Valuations

i	$v_i(\text{Court mandate})$	$v_i(\text{Civil unions})$	$v_i(\text{Marriage amend.})$
R	4.512	4.565	3.478
N	7.195	7.5	7.5
D	5.488	5.435	6.522

The following table summarizes the above information by depicting individuals' rankings over the alternatives when individuals are both myopic and farsighted. If Player i strictly prefers policy x to policy y , it is written $x \succ y$. If Player i is indifferent between the two, it is written $x \sim y$.

Individuals' Rankings of Alternatives

<i>i</i>	One-shot rankings	Farsighted rankings
Gay rights activist (R)	Court mandate \succ Civil unions \succ Marriage amendment	Civil unions \succ Court mandate \succ Marriage amendment
Neutral voter (N)	Civil unions \sim Marriage amendment \succ Court mandate	Civil unions \sim Marriage amendment \succ Court mandate
Defense-of-marriage activist (D)	Marriage amendment \succ Civil unions \succ Court mandate	Marriage amendment \succ Court mandate \succ Civil unions

As is consistent with a traditional spatial model, the neutral voter is always indifferent between legalized civil unions and a constitutional marriage amendment, and strictly prefers both of these policies to a court mandate legalizing gay marriage. However, the gay rights activist, with ideal point “court mandate,” will strictly prefer to implement “civil unions” rather than his own ideal point when he is farsighted. This is because when “marriage amendment” is more politically salient than the other two policies (i.e. is brought to the voters’ attention more often by density Q) the policy that makes the gay rights activist best off over time is not his ideal point, but the policy closest to his ideal point that can defeat a constitutional marriage amendment. Thus, it is in the activist’s best interest to concede some utility in the current round in order to prevent his least favorite policy from being quickly implemented. Finally, the defense-of-marriage activist strictly prefers “court mandate” to “civil unions” when he is farsighted, even though “civil unions” is closer to his ideal point. Loosely speaking, this is because at “court mandate” there is a 60 percent chance of transitioning to “marriage amendment,” the defense-of-marriage activist’s favorite policy, while at “civil unions” this chance drops to 30 percent.

The purpose of this example is to demonstrate several key features of this model. First, all three players

care both about the policy chosen today and about the types of policies that will replace it in the future. Second, farsightedness requires taking into account institutional particularities at any given time. In this example this is reflected in the exogenous Q term, which captures the fact that there is currently considerable pressure from interest groups to enact gay marriage bans at the federal level. And third, farsightedness also requires taking into account the preferences of other political actors. In the absence of a defense-of-marriage activist, the gay rights activist would have pursued a different political strategy. While this example is obviously highly stylized, it provides a clear picture of how this model works, and demonstrates that the predictions that this model yields are often quite intuitive.

A final point to note is that the “farsighted preferences” conceived of in this model are only farsighted with respect to the current political climate, or Q term. In this sense, the model predicts short-term behavior when individuals care about policy in the long run. In this example a strategic gay rights activist will not challenge state-level bans of gay marriage in the federal courts *given the current political climate*. Clearly the political climate will change over time, and these changes will necessarily change the predictions of the model.

3.2 A Condorcet Winner-Turned-Loser

In this next example there are three voters and four alternatives, $X = \{c, x, y, z\}$, with one alternative, c , being a Condorcet winner. The other three alternatives form a majority preference cycle, with $x \succ y$, $y \succ z$, and $z \succ x$. The Condorcet winner gives each voter a one-shot payoff of 2, while the expected value of an alternative in the majority preference cycle is 3. Thus, while every player is better off cycling through the alternatives x, y, z than remaining at Condorcet winner c , c always gives two players strictly greater utility in the short term than they would receive at any other policy. The following tables present players’ utility

functions and their farsighted valuations when Q is uniform over all alternatives (i.e. when every alternative is equally likely to be brought to the floor for a vote).

One-shot Utility					Farsighted (Equilibrium) Valuations				
i	$u_i(c)$	$u_i(x)$	$u_i(y)$	$u_i(z)$	i	$v_i(c)$	$v_i(x)$	$v_i(y)$	$v_i(z)$
1	2	8	1	0	1	28.7097	39.0463	30.109	20.8447
2	2	1	0	8	2	28.7097	30.109	20.8447	39.0463
3	2	0	8	1	3	28.7097	20.8447	39.0463	30.109

The following table depicts individuals' rankings over the alternatives when individuals are both shortsighted and farsighted. As in the previous example, a preference reversal occurs. This reversal changes Condorcet winner c into a Condorcet *loser* in farsighted valuations, or a policy that is majority-defeated by every other policy.¹⁰

Individuals' Rankings of Alternatives

i	One-shot rankings	Farsighted rankings
1	$x \succ c \succ y \succ z$	$x \succ y \succ c \succ z$
2	$z \succ c \succ x \succ y$	$z \succ x \succ c \succ y$
3	$y \succ c \succ z \succ x$	$y \succ z \succ c \succ x$

The logic behind why the preference reversal occurs takes two steps to reveal. First note that in the short run Player 1 prefers c to y , but in the long run this preference is reversed. However, c and y both lead to similar payoffs in the subsequent round for Player 1: since every alternative beats c when voters are farsighted, the

¹⁰Roberts [21] provides an interesting refinement of Condorcet winners to intertemporal settings that captures those policies that can also defeat cycles of alternatives. This example works precisely because the Condorcet winner here does not satisfy Roberts's stronger criterion of being a "generalized Condorcet winner."

expected payoff to Player 1 in a round following the implementation of c is $\frac{1}{4}(2 + 8 + 1 + 0) = \frac{11}{4}$. Player 1's expected payoff after y is implemented is the same: y is only defeated by x , and so Player 1's expected payoff in the round following implementation of y is $\frac{1}{4}(8) + \frac{3}{4}(1) = \frac{11}{4}$. However, c and y lead to different expected payoffs for Player 1 two rounds out. While y and c both lead to an expected payoff of $\frac{11}{4}$ in a subsequent round, y leads to an expected payoff of $\frac{57}{16}$ two rounds out, while c leads to an expected payoff of $\frac{47}{16}$ two rounds out.¹¹ Thus, while equilibrium evaluations in this framework capture expected payoffs for an infinite stream of future policies, the intuition behind why preference reversals occur can be seen in settings with a finite (and even small) number of periods.

This example demonstrates that, when voters are farsighted, outcomes will not necessarily coincide with many commonly known tournament solution concepts including the uncovered set, minimal covering set, tournament equilibrium set, Banks set, largest consistent set, and von Neumann-Morgenstern stable set, as all of these sets reduce to the core, if one exists.¹² In the spatial settings considered in the following section, outcomes do appear to coincide with elements of the von Neumann-Morgenstern stable set. However, this observation cannot be extended to a general preference environment. The relevant issue is that cardinality of preferences matters in this setting, whereas tournament solution concepts only require ordinal preferences. This same issue distinguishes this model from sophisticated voting (the standard definition of which is presented in [24]), which is defined solely with respect to ordinal preferences. Sophisticated voting will always

¹¹To see the logic of this, let y be implemented today, at $t = 0$. Tomorrow at $t = 1$ y will lead to either x (with probability $\frac{1}{4}$) or y (with probability $\frac{3}{4}$), for an expected payoff to Player 1 of $\frac{11}{4}$. At $t = 2$ there is a $\frac{1}{4}$ chance we will have been at x at $t = 1$, which will lead to x with probability $\frac{3}{4}$ and z with probability $\frac{1}{4}$, for an expected payoff of $\frac{24}{4}$. Similarly, there is a $\frac{3}{4}$ chance we will have been at y in the previous round, for an expected payoff at $t = 2$ of $\frac{11}{4}$. Thus the expected payoff at $t = 2$ when y was implemented at $t = 0$ is $\frac{1}{4}(\frac{24}{4}) + \frac{3}{4}(\frac{11}{4}) = \frac{57}{16}$.

¹²These six sets are defined in [17], [11], [23], [2], [8], and [25], respectively.

yield a Condorcet winner as the unique voting outcome, if a Condorcet winner exists. Here a Condorcet winner exists, but it is chosen with probability zero as an equilibrium policy outcome.

More generally, this example of a Condorcet winner-turned-loser provides a good basis of comparison between this model and the standard sophisticated voting setup that is frequently used to analyze forward-looking behavior in legislative settings. Sophisticated voting describes strategic voting behavior over a finite, predetermined sequence of alternatives. As in this model, sophisticated voters may seemingly exhibit “preference reversals,” in that they may vote in favor of alternatives that give them lower utility in order to beneficially affect the future path of play. However, under sophisticated voting, voters do not actually concede anything; individuals may vote against policies that they like, but only because they know that what they like cannot win. In this model, farsighted voters may take short term losses in order to do *better* in expectation than they could have in a one-shot game.¹³

4 Numerical examples in two-dimensional spaces

What follows is a look at several numerical estimations of this model in settings where the policy space is two-dimensional. The first setting is that of a three-player constant sum game and the second setting is that of a three-player spatial model where players have convex preferences. The graphs that follow depict both the equilibrium value functions of one of the players and the equilibrium distribution over observed

¹³This model also considers a different agenda framework than the amendment agendas considered under sophisticated voting. Not only is the proposal process in this model probabilistic, but agendas move “forward” in that the status quo is replaced at every round of voting. The standard sophisticated voting setup considers agendas as fixed orderings of alternatives, with policies sequentially eliminated through a planned series of pairwise votes. These agendas are commonly voted on “backward,” with a fixed status quo considered last.

outcomes.¹⁴ In all of the estimations it is assumed that voting is via majority rule and that players vote deterministically, as in the previous section. It is also assumed that every player has the same discount factor, $\delta = 0.9$, and that Q is uniform over the policy space. Q was chosen to be uniform simply as a baseline.

4.1 Three players divide a dollar

[FIGURE 2 HERE]

Figure 2 is a graph of Player 1's value function. The setting is a three-player divide-the-dollar game; the policy space equals the set of all divisions of the dollar between three people and a player's utility from a particular policy equals the amount of money he is allocated by that policy. The policy space is pictured, and Player 1's ideal point (the policy $x = (1, 0, 0)$) is at the top of the simplex. The bottom of the simplex denotes those policies that give Player 1 no portion of the dollar. The darkest areas correspond to the policies that yield the highest values for Player 1, and the lightest areas denote the policies that yield the lowest values.

¹⁴The estimations were run by discretizing the policy space into approximately nine hundred policies (for Example 1) or two hundred and sixty policies (Example 2) and then iterating the mapping defined in Equation 2 of Appendix B until it converged numerically to a dynamically stable voting equilibrium. The iterations were performed by letting $v_0 = \{u_i\}_{i \in N}$ and defining $v_{t+1} = g(v_t)$ for $t \geq 0$. Convergence was obtained in every example for a sup norm of .002. Once the equilibrium value function was found, the equilibrium distribution over outcomes was found by first drawing two policies from density Q , pitting them against each other (assuming that voters vote according to their equilibrium value functions), pitting the winner against a new policy drawn from Q , pitting *this* winner against a new policy drawn from Q , and so on. This process was repeated 200,000 times. The frequency with which each policy arose as an outcome generated the ultimate distribution over observed outcomes.

It is apparent that the policies that Player 1 values most are not near Player 1's ideal point, but rather those that divide the dollar about equally between himself and one other player, or $(\frac{1}{2}, \frac{1}{2}, 0)$ and $(\frac{1}{2}, 0, \frac{1}{2})$. The intuition for this is simple. Suppose that the status quo policy in a given round is Player 1's ideal point, $(1, 0, 0)$. Then whichever policy is chosen to be pitted against the status quo in the next round will win with near certainty, because every policy weakly defeats Player 1's ideal point. Conversely, the point $(\frac{1}{2}, \frac{1}{2}, 0)$, as an example, is more stable and less likely to be defeated by a new policy. This is why Player 1's *least* favorite policy is at $(0, \frac{1}{2}, \frac{1}{2})$. Not only does Player 1 get a payoff of zero from this policy, but it is also a relatively stable outcome, unlikely to be replaced quickly.

[FIGURE 3 HERE]

Figure 3 depicts the density over observed policy outcomes. The darkest areas correspond to the most frequently observed policies. In this example only a small subset of the total policy space is ever observed with positive probability. The observed policies appear to constitute a majority-rule core with respect to players' equilibrium value functions. Figure 2 demonstrates this—since the setting is symmetric, it is clear that each of Player 1's most-preferred policies is also the most-preferred policy of another player. This example demonstrates that the assumption of farsightedness gives us sharp predictions in this divide-the-dollar game. It predicts outcomes corresponding to the set of policies that divide the dollar evenly between all members of a minimal winning coalition. In this example, the likelihood of a policy defeating a status quo such as $(\frac{1}{2}, \frac{1}{2}, 0)$ is approximately 1%.

In social choice theory this set of predicted policies is referred to as the *von Neumann-Morgenstern stable set* (or simply *stable set*), and is defined by the nice property that no element in the set strictly defeats any other element in the set and that any policy outside the set is strictly defeated by an element of the set.¹⁵

¹⁵These two conditions are termed *internal* and *external stability*, respectively.

It is interesting to note that the predictions generated by this model when Q is uniform closely coincide with elements of the von Neumann-Morgenstern stable set in *all* of the spatial settings considered. However, as Section 3.2 demonstrates, this is not true in general.

It is also interesting to compare this example with a standard one-shot bargaining model in which players are recognized with equal probability to propose allocations of the dollar and if an allocation is approved by a majority of voters, it goes into effect. In the one-shot model proposals are made optimally with regard to future optimal behavior by players (in the event that a proposal is rejected). Baron and Ferejohn [4] demonstrate that the only stationary subgame-perfect equilibrium of the one-shot model is for the proposer to give himself $\frac{\delta(n-1)}{2n}$ of the dollar and to give $\frac{n-1}{2}$ other players $\frac{\delta}{n}$ each. In both models, only a bare majority of players receive a positive share of the dollar. In Baron and Ferejohn's model, members of the winning coalition receive different allocations of the dollar depending on whether or not they were the proposer. Here members of the winning coalition divide the dollar equally among themselves, as proposals are exogenous. Of course, ex ante outcomes in the one-shot model and this model are the same in expectation.

4.2 An asymmetric two-dimensional spatial model

This last series of pictures depicts a two-dimensional spatial model, where the ideal points of the three players are no longer symmetric but are located at $(0, \frac{1}{2})$, $(0, 0)$, and $(1, 0)$. The policy space is bounded by the lines connecting the ideal points of the three players.¹⁶ First, preferences are assumed to be circular so that players are indifferent between all policies equidistant from their ideal points. This implies that each issue dimension matters equally to each player. Then we will consider the case where two of the three players care more about one issue dimension than the other.

¹⁶When preferences are circular, this policy space corresponds to the Pareto set. When preferences are elliptical, as in the next example, the policy space subsumes the Pareto set.

Figure 4 depicts the spatial location of the ideal points of the three players and their indifference curves when preferences are circular. Figure 5 depicts the equilibrium value function of Player 1, whose ideal point is located at $(0, \frac{1}{2})$. Again, the darker areas correspond to the policies that Player 1 values most highly. Figure 6 depicts the frequency with which each policy is observed as an outcome, with the most frequent outcomes being darker in color than the less frequent ones.

[FIGURE 4 HERE]

[FIGURE 5 HERE]

[FIGURE 6 HERE]

In Figure 5 we can see that Player 1's highest-valued alternative is close to the point $(0, .25)$. Although not pictured, the equilibrium "highest-valued" policies of Players 2 and 3 are $(.03, .03)$ and $(.94, 0)$, respectively. Farsightedness induces Players 1 and 3 to prefer policies that may spark the formation of a coalition between themselves and Player 2, the most moderate player. Figure 6 shows that the most observed outcome is approximately $(0, .22)$, close to the alternative in the stable set corresponding to a coalition between Players 1 and 2, the two players whose ideal points are closest to each other. In this example the stable set consists of the points $\{(0, .19), (.28, .36), (.19, 0)\}$, approximately.

[FIGURE 7 HERE]

In the last example the preferences of Players 2 and 3 are now elliptical rather than circular, and are defined by the equation

$$u_i(x_1, x_2) = -\sqrt{(r_{i1} - x_1)^2 + 100(r_{i2} - x_2)^2},$$

where $r_i = (r_{i1}, r_{i2})$ is the ideal point of player i . Thus, Players 2 and 3 value the second (or y) dimension of the policy space ten times more highly than the first. The preferences of Player 1 have remained unchanged. Figure 7 shows the ideal points of the three players and their indifference curves. The dotted curve represents the contract curve of Players 1 and 3, and is the upper bound of the Pareto set.

[FIGURE 8 HERE]

[FIGURE 9 HERE]

Interestingly, even though Player 1's utility function is the same as in the previous example, his equilibrium value function is quite different than both his utility function and his value function in the previous example, when the preferences of the other two players were circular. Figure 8 shows that Player 1's most-preferred alternatives now lie close to the origin, the ideal point of Player 2. The reason for this is similar to the intuition behind the example given in Section 3.1. Because the indifference curves of Players 2 and 3 both favor policies that lie close to the x -axis, Player 1 knows that implementing a policy that appeals to him along the y -dimension is a lost cause. This is because the point $(0, \frac{1}{2})$, Player 1's ideal point, is the alternative in X that is *farthest* from the x -axis. Thus, he is willing to concede a great deal of utility along the second dimension of the policy space in order to collude with Player 2 along the first dimension.

The stable set in this example approximately equals $\{(.14, .03), (0, .01), (.09, 0)\}$. Figure 9 shows that there exists a single alternative, $(.09, .02)$, that arises with near certainty. This alternative is close to the alternative in the stable set corresponding to a coalition between Players 2 and 3. As in the previous example, this prediction corresponds to the most efficient element of the stable set; it is the element of the stable set that maximizes the sum of the players' utilities.

5 Conclusion

Societies frequently make decisions that will persist into the future. Indeed, it is not a stretch to argue that most policies that people care about are of this form, with examples being social policies, entitlements, budgets, and redistributive policies. This paper argues that there is no reason to expect that preferences over these “continuing” policies will be similar to preferences over “once-and-for-all” policies. This is because, when evaluating continuing policies, individuals consider not only their payoffs from the policies themselves, but also from what the policies will lead to in the future. I show that when preferences are considered in this way surprising outcomes can emerge in several standard legislative settings. For example, I demonstrate that a policy that is a Condorcet winner in a one-shot game is selected with probability zero as a policy outcome when voters are farsighted.

One implication of the model is that its predictions can give insight into the sorts of coalitions that may form in settings where policy is implemented over many rounds. While the existence of stable coalitions is undeniably central to political life, such coalitions can be difficult to understand from a theoretical perspective. This type of dynamic environment is perhaps one of the most natural in which to think of the formation of alliances, and this paper formalizes a common argument for why stability can arise and persist in the real world. In the theory presented here, individuals consider the trade off between the immediate value of a policy and the long-run stability of the coalition implementing that policy. Ultimately, this consideration leads to the recognition that policies that fairly divide benefits between members of a winning coalition leave individual players best off in the long run. The cooperation that emerges in this model does not rely on any threat of punishment other than the fact that current policies can be replaced by new alternatives.

The theory can also be interpreted as providing an explanation for why particular coalitions are more likely to form than others. Both the analytic and numerical examples demonstrate that farsighted voters

will frequently vote for policies that do not necessarily give them the highest one-shot payoff. The favorite policies of a farsighted voter will depend on a combination of his own preferences, the preferences of other voters, the voting power of other voters, and the likelihood with which certain policies will be brought to the floor. Thus, the model provides a nuanced characterization of voter considerations that encompasses many different elements of the institutional environment. It also provides a characterization of voting behavior that is estimable because the model yields distributional predictions.

While the model presented in this paper is purely formal, the theory is applicable to a variety of real-world legislative settings, as it utilizes only weak assumptions about the number of voters, their preferences, their respective voting weights, the majority requirement, and the policy space. However, the predictive power of the model will depend largely on the functional form of the proposal process Q . This process can be thought of as representing the likelihood with which particular policies will be considered by the group in the future. Estimating these likelihoods in real-world situations may provide insight into the prospective voting behavior of legislators, with the implication being that the perceived distribution of future policy considerations may be an omitted variable in some empirical models of legislative voting.

Appendix A: Assumptions on individual vote choice

Throughout, I assume that for all $x, y \in X$, $p(v(x), v(y))$, or the probability of transitioning from status quo x to policy y , given that x and y are put to a vote, can be written as the probability of victory of y over x :

$$p(v(x), v(y)) = \sum_{C \in W} \prod_{i \in C} p_i(v_i(x), v_i(y)) \prod_{i \notin C} (1 - p_i(v_i(x), v_i(y))) \quad (1)$$

where $p_i(v_i(x), v_i(y)) \in [0, 1]$ represents Player i 's probability of voting for y over x given value function v . It is assumed that p_i is independent of p_j for all $i, j \in N$, that $p_i(v_i(x), v_i(y)) + p_i(v_i(y), v_i(x)) = 1$,

and that p_i is increasing in $v_i(y) - v_i(x)$.

While the general model defined in Section 2 does not require any additional assumptions on the functions $p_i(v_i(x), v_i(y))$, imposing more structure on these functions enables us to obtain results about equilibrium existence (Propositions 1, 2, and 3), and about when equilibrium behavior is Markov-perfect (Proposition 4). In particular, the two natural specifications of individual vote choice I consider are deterministic behavior and probabilistic behavior. These two specifications are discussed and defined below. Deterministic voting simply assumes that individuals vote for y over x if the long-run payoff of having y implemented today is at least as high as the long-run payoff of having x implemented today.

Definition: Individuals vote *deterministically* if $p_i(v_i(x), v_i(y)) = 1$ if $v_i(y) \geq v_i(x)$ and zero otherwise.

Probabilistic voting assumes that, at each round of voting, each individual i receives an unobserved payoff disturbance θ_{ix} from casting a vote for policy $x \in X$. As is standard in models of probabilistic voting, these θ_{ix} terms are assumed to be independently and identically distributed across all policies x , voters i , and, implicitly, rounds of voting. Furthermore, the distribution of θ_{ix} has full support and a cumulative distribution function F that is twice continuously differentiable. If the payoff disturbances satisfy all of these properties, they are termed *admissible*.

Assuming this payoff structure implies the following definition of probabilistic voting, which will be shown in Proposition 4 to be consistent with an assumption that individuals play Markovian pure strategy Bayesian Nash equilibria.¹⁷

¹⁷See [15] and [16] for a more thorough presentation of probabilistic voting and agent quantal response equilibria.

Definition: Individuals vote F -probabilistically (or simply, *probabilistically*) if, if for all $i \in N$,

$$p_i(v_i(x), v_i(y)) = \Pr_F[v_i(y) + \theta_{iy} \geq v_i(x) + \theta_{ix}].$$

As discussed in Section 3, each of the following three existence results requires that individuals vote probabilistically. Proposition 3 also requires that for all i , $p_i(v_i(x), v_i(y))$ be differentiable in both of its arguments, and that these derivatives be uniformly bounded. And Proposition 2 additionally requires that these derivatives be uniformly bounded by a specific constant. These two additional assumptions are formalized and discussed below.

Assumption 1 $|\frac{\partial}{\partial v_j(x)} p_j(v_j(x), v_j(y))| \leq K$, for some $K \in \mathbb{R}_+$.

With respect to the assumption of probabilistic voting and Assumption 1, note that it is always possible to approximate a discontinuous function with such a continuous, differentiable one. These assumptions are substantively weak and, furthermore, they are not necessary in order to demonstrate that an equilibrium exists in specific settings. In the estimations and analytic examples presented, equilibria are shown to exist even when individuals vote deterministically. Assumption 2, however, imposes a real restriction on individual behavior because it limits how responsive voting decisions can be to payoffs. In the definition of this assumption, let $\bar{u} = \max_{x \in X, j \in N} u_j(x)$, and let $\underline{u} = \min_{x \in X, j \in N} u_j(x)$.

Assumption 2 $|\frac{\partial}{\partial v_j(x)} p_j(v_j(x), v_j(y))| \leq \frac{(1-\delta)^2}{\delta n |\bar{u} - \underline{u}|}$ for all $j \in N$.

First, note that the bound defined by this condition is very conservative. Existence may be obtained in far less restrictive environments. And second, while this assumption may seem strange, it can be interpreted in the context of standard models of probabilistic voting. In the standard example of a logit agent quantal response equilibrium in which $p_i(v_i(x), v_i(y)) = \frac{e^{\lambda v_i(y)}}{e^{\lambda v_i(x)} + e^{\lambda v_i(y)}}$ for some $\lambda \geq 0$, this assumption will

impose a restriction on λ for a fixed n and on n for a fixed λ .¹⁸ More specifically, in the case of majority rule and assuming that $u_i(x) \in [0, 1]$ for all $i \in N$ and $x \in X$, it will hold whenever

$$n\lambda \leq \frac{4(1-\delta)^2}{\delta}.$$

Since the right hand side of this equation is always positive, it follows that for any fixed number of players there will always exist a positive λ that guarantees equilibrium existence.

Appendix B: Analytic results

For the first three propositions we will define a function g that maps value functions into value functions, or $g : \mathcal{V}^n \rightarrow \mathcal{V}^n$ with $g = \{g_i\}_{i \in N}$ and $g_i : \mathcal{V}^n \rightarrow \mathcal{V}$. Specifically,

$$g_i(v(x)) = u_i(x) + \delta \int_{y \in X} v_i(y)p(v(x), v(y)) + v_i(x)(1 - p(v(x), v(y)))Q(y)dy, \quad (2)$$

with the case of a finite X defined similarly. It is useful to note that this function g can also be used to consider farsighted voting when there are only a finite number of periods of policymaking. Let $v_0 = \{u_i\}_{i \in N}$ and iteratively define $v_{t+1} = g(v_t)$, $t \geq 0$. Then $v_{it}(x)$ captures Player i 's valuation of policy x given that t rounds of policymaking will occur after x is implemented.

Proposition 1 If X is finite, then there exists a dynamically stable voting equilibrium when individuals vote probabilistically.

¹⁸A logit AQRE assumes that the θ_{ix} follow a type one extreme value distribution.

Proof: Since $\delta < 1$ and u_i is real-valued for all $i \in N$, the upper bound any individual's value function could take is $\frac{1}{1-\delta} \max_{x \in X} u_i(x)$, and the lower bound is zero. Thus, for every $v \in \prod_{i \in N} \mathbb{R}^X$, $v \in \prod_{i \in N} [0, \frac{1}{1-\delta} \max_{x \in X} u_i(x)]^X$, and so the set of value functions is bounded. Furthermore, the set of value functions is convex, since the convex combination of two bounded functions taking X to \mathbb{R} is itself bounded. Last, the set of value functions is closed, trivially. It follows that the set of value functions taking X into the real numbers \mathbb{R} is a nonempty, closed, bounded and convex subset of a finite-dimensional vector space, \mathbb{R}^X .

The mapping $g : \prod_{i \in N} \mathbb{R}^X \rightarrow \prod_{i \in N} \mathbb{R}^X$ (see Equation 2) is single-valued by definition, and is continuous by the continuity of every $p_i(v_i(x), v_i(y))$. By Brouwer's Fixed Point Theorem, there exists a $v \in \prod_{i \in N} \mathbb{R}^X$ such that $g(v) = v$. Thus, there exists a dynamically stable voting equilibrium. \square

When policy space X is infinite Assumption 2 is needed in order to guarantee existence, along with a definition and a lemma.

Definition: A set of real-valued functions $\mathcal{V}^* \subset \mathcal{V}$ is *equicontinuous* if for all $\epsilon > 0$, there exists a $\delta > 0$ such that

$$\rho(s, t) < \delta \text{ and } v_i \in \mathcal{V}^* \Rightarrow |v_i(s) - v_i(t)| < \epsilon.$$

To prove Proposition 2, we are concerned in particular with a set $\mathcal{B}_M^n \subset \mathcal{V}^n$ of vectors of differentiable functions taking X to \mathbb{R} whose derivatives are uniformly bounded by the constant M . This set is equicontinuous; let M be a bound for the derivatives of the functions in \mathcal{B}_M , and recall that for $v \in \mathcal{V}^n$, $\rho(v(s), v(t)) = \max_{i \in N} |v_i(s) - v_i(t)|$. For an $\alpha \in X$ and J equal to the dimensionality of the policy space, let $|\nabla v_i(\alpha)| = \max_{j \in J} |\frac{\partial v_i}{\partial \alpha_j}|$. Then, by an extension of the Mean Value Theorem, $\rho(s, t) < \delta$ implies that

$\rho(v(s), v(t)) = \max_i |\nabla v_i(\alpha)| \rho(s, t) \leq M\delta$, for some α on the line segment between s and t . Thus, given $\epsilon > 0$, the choice $\delta = \epsilon/(M + 1)$ demonstrates that \mathcal{B}_M , and thus \mathcal{B}_M^n , is equicontinuous.

Lemma 1 *If Assumption 2 holds, then the function g maps a closed, bounded, and equicontinuous subset of \mathcal{V}^n into itself.*

Proof: Define g as in Equation 2. Boundedness is attained because $\delta < 1$. Let \mathcal{B}_M^n be the set of vectors of differentiable functions whose derivatives are uniformly bounded by the constant M . The set \mathcal{B}_M^n is closed. I will show that there exists an $M \in \mathbb{R}_+$ such that for any $v \in \mathcal{V}^n$, if $v \in \mathcal{B}_M^n$, then $g(v) \in \mathcal{B}_M^n$. By Equation 2 we know that for all i ,

$$g(v_i(x)) = u_i(x) + \delta \int_{y \in X} v_i(y) p(v(x), v(y)) + v_i(x)(1 - p(v(x), v(y))) Q(y) dy$$

and thus,

$$\begin{aligned} |\nabla g(v_i(x))| &\leq |\nabla u_i(x)| + \delta |\nabla v_i(x)| \left(1 - \int_{y \in X} p(v(x), v(y)) Q(y) dy\right) \\ &\quad + \delta \int_{y \in X} (v_i(y) - v_i(x)) |\nabla p(v(x), v(y))| Q(y) dy, \end{aligned} \quad (3)$$

where, letting J equal the dimensionality of the policy space and letting f be any function of x , $|\nabla f(x)| = \max_{j \in J} \left| \frac{\partial f(x)}{\partial x_j} \right|$. From the definition of $p(v(x), v(y))$ we get

$$\begin{aligned} |\nabla p(v(x), v(y))| &\leq \\ &\sum_{C \in \mathcal{W}} \left[\left[\sum_{i \in C} |\nabla v_i(x)| \frac{\partial}{\partial v_i(x)} p_i(v_i(x), v_i(y)) \prod_{j \in C \setminus \{i\}} p_j(v_j(x), v_j(y)) \prod_{j \notin C} (1 - p_j(v_j(x), v_j(y))) \right] \right. \\ &\quad \left. - \left[\sum_{i \notin C} |\nabla v_i(x)| \frac{\partial}{\partial v_i(x)} p_i(v_i(x), v_i(y)) \prod_{j \in C} p_j(v_j(x), v_j(y)) \prod_{j \notin C \cup i} (1 - p_j(v_j(x), v_j(y))) \right] \right] \\ &= \sum_{i \in N} |\nabla v_i(x)| \frac{\partial}{\partial v_i(x)} p_i(v_i(x), v_i(y)) Z_i(\{p_j(v_j(x), v_j(y))\}_{j \in N \setminus \{i\}}) \end{aligned} \quad (4)$$

where, letting C_i^M equal the set of minimal winning coalitions that i is in,

$$Z_i(\{p_j(v_j(x), v_j(y))\}_{j \in N \setminus \{i\}}) = \sum_{C \in C_i^M} \prod_{j \in C \setminus \{i\}} p_j(v_j(x), v_j(y)) \prod_{j \notin C} (1 - p_j(v_j(x), v_j(y))). \quad (5)$$

$Z_i(\{p_j(v_j(x), v_j(y))\}_{j \in N \setminus \{i\}})$ represents the probability that Player i 's vote is pivotal given that all other players j vote according to the functions $p_j(v_j(x), v_j(y))$. Moving back to Equation 3, it follows that

$$\begin{aligned} |\nabla g(v_i(x))| &\leq |\nabla u_i(x)| \\ &+ \delta |\nabla v_i(x)| (1 - \int_{y \in X} p(v(x), v(y)) Q(y) dy) \\ &+ \delta \int_{y \in X} (v_i(y) - v_i(x)) \sum_{k \in N} |\nabla v_k(x)| \frac{\partial}{\partial v_k(x)} p_k(v_k(x), v_k(y)) Z_k(\{p_j(v_j(x), v_j(y))\}_{j \in N \setminus \{k\}}) Q(y) dy \end{aligned}$$

Let $U = \max_{j \in N, x \in X} |\nabla u_j(x)|$. U is assumed to be bounded. Let $B = \max_{j \in N, x, y \in X} |v_j(x) - v_j(y)|$, which is bounded by the assumption that $\delta < 1$ and X is compact. Let $K = \max_{j \in N} |\frac{\partial}{\partial v_j(x)} p_j(v_j(x), v_j(y))|$. Last, assume that $v(x) \in \mathcal{B}_M^n$, or that $M = \max_{j \in N, x \in X} |\nabla v_j(x)|$. Using the fact that Equation 5 is bounded between zero and one, we now get that

$$|\nabla g(v_i(x))| \leq B + \delta M + \delta * N * K * B * M.$$

It follows that if $v \in \mathcal{B}_M^n$, then $g(v) \in \mathcal{B}_M^n$ if

$$B + \delta M + \delta * N * K * B * M \leq M,$$

or

$$K \leq \frac{1 - \delta}{\delta * n * B}. \quad (6)$$

As $B = \frac{\bar{u}-u}{1-\delta}$, Equation 6 is exactly equal to Assumption 2. It follows that if Assumption 2 holds, then g maps a closed, bounded and equicontinuous subset of \mathcal{V}^n into itself. \square

Proposition 2 If X is infinite, then there exists a dynamically stable voting equilibrium when Assumption 2 holds.

Proof: The Heine-Borel Theorem in a function space tells us that a subset $\mathcal{V}^* \subset \mathcal{V}$ is compact if and only if it is closed, bounded, and equicontinuous.¹⁹ Lemma 1 proves that the set of value functions can be restricted to the compact set \mathcal{B}_M^n . Since the function $g : \mathcal{B}_M^n \rightarrow \mathcal{B}_M^n$ is continuous, we need only convexity of the set of value functions to prove that there exists an equilibrium value function.

Take the convex combination of any two value functions, $v, w \in \mathcal{B}_M^n$, so that for any $\gamma \in [0, 1]$, $\gamma v(x) + (1-\gamma)w(x) = z(x)$. Clearly z is continuous, since v and w are continuous. Furthermore, z is differentiable, and the derivative of z is bounded by the constant M . It follows that $z \in \mathcal{B}_M^n$, and that \mathcal{B}_M^n is convex. By Brouwer's Fixed Point Theorem, there exists a v such that $g(v) = v$. \square

Proposition 3 If Assumption 1 holds and n is large then there always exists a *unique* equilibrium, regardless of whether X is finite or infinite.

Proof: The proof is specifically for the case where X is infinite; the finite case can be proved similarly. For $w, z \in \mathcal{V}^n$, let $\rho(w_i, z_i) = \max_{x \in X} |w_i(x) - z_i(x)|$, and let $\rho(w, z) = \max_{i \in N} \rho(w_i, z_i)$. We must show that for any $w, z \in \mathcal{V}^n$, $\rho(g(w), g(z)) < \rho(w, z)$, or that g is a contraction mapping.

As in Equation 2, let $g_i : \mathcal{V}^n \rightarrow \mathcal{V}$. Thus, $g = (g_1, \dots, g_n)$. First consider the gradient vector ∇g_i . For all

¹⁹This theorem is a direct consequence of the Arzela-Ascoli Theorem. See [19, p. 217].

$x \in X$,

$$g_i(v(x)) = u_i(x) + \delta \int_{y \in X} v_i(y)p(v(x), v(y)) + v_i(x)(1 - p(v(x), v(y)))Q(y)dy.$$

Thus, the components of $\nabla g_i(v(x))$ can be defined using the partial derivatives

$$\frac{\partial g_i(v(x))}{\partial v_i(x)} = \delta \left[1 - \int_{y \in X} p(v(x), v(y))Q(y)dy \right] + \delta \int_{y \in X} (v_i(y) - v_i(x)) \frac{\partial p(v(x), v(y))}{\partial v_i(x)} Q(y)dy \quad (7)$$

and for all $j \in N \setminus \{i\}$,

$$\frac{\partial g_i(v(x))}{\partial v_j(x)} = \delta \int_{y \in X} (v_i(y) - v_i(x)) \frac{\partial p(v(x), v(y))}{\partial v_j(x)} Q(y)dy. \quad (8)$$

Using Equation 1 we get that for all $i \in N$,

$$\frac{\partial p(v(x), v(y))}{\partial v_i(x)} = \frac{\partial p_i(v_i(x), v_i(y))}{\partial v_i(x)} Z_i(\{p_j(v_j(x), v_j(y))\}_{j \in N \setminus \{i\}}) \quad (9)$$

where $Z_i(\cdot)$ is defined as in Equation 5. Recall that $Z_i(\{p_j(v_j(x), v_j(y))\}_{j \in N \setminus \{i\}})$ represents the probability that Player i 's vote is pivotal given that all other players j vote according to the functions $p_j(v_j(x), v_j(y))$. McKelvey and Patty ([16], Lemma 1) prove that when people vote probabilistically (i.e when for all $j \in N$, and all $x, y \in X$, $p_j(v_j(x), v_j(y)) \in (0, 1)$), all pivot probabilities $Z_i(\cdot) \rightarrow 0$ as n gets large.

Combining Equations 8 and 9, we get for all $j \in N \setminus \{i\}$

$$\frac{\partial g_i(v(x))}{\partial v_j(x)} = \delta \int_{y \in X} (v_i(y) - v_i(x)) \frac{\partial p_j(v_j(x), v_j(y))}{\partial v_j(x)} Z_j(\{p_k(v_k(x), v_k(y))\}_{k \in N \setminus \{j\}}) Q(y) dy.$$

By Assumption 1 we know that for all $j \in N$ and $x, y \in X$, $\frac{\partial p_j(v_j(x), v_j(y))}{\partial v_j(x)}$ is bounded by some constant.

We also know that the difference $|v_j(y) - v_j(x)|$ is bounded by a constant, since $\delta < 1$ and utility is bounded.

Since $Z_j(\cdot) \rightarrow 0$ as $n \rightarrow \infty$, it follows that for any $\epsilon > 0$ there exists an $M \in \mathbb{N}$ such that for all $n > M$,

$$\frac{\partial g_i(v(x))}{\partial v_j(x)} < \epsilon.$$

Using Equation 7, by the same logic it follows that for any $\epsilon > 0$ there exists an $M \in \mathbb{N}$ such that for all $n > M$,

$$\frac{\partial g_i(v(x))}{\partial v_i(x)} < \delta [1 - \int_{y \in X} p(v(x), v(y)) Q(y) dy] + \epsilon.$$

Define $|\nabla g(v)|$ such that

$$|\nabla g(v)| = \max_{\{i,j\} \in N} \left(\max_{x \in X} \left| \frac{\partial g_i(v(x))}{\partial v_j(x)} \right| \right).$$

Since $\delta [1 - \int_{y \in X} p(v(x), v(y)) Q(y) dy] \in (0, 1)$ for all $\delta < 1$, it follows that for n sufficiently large (i.e., ϵ sufficiently small), $|\nabla g(v)| < 1$.

By the Mean Value Theorem we know that

$$\rho(g(w), g(z)) \leq \rho(w, z) |\nabla g(v)|$$

for some v on the line segment between w and z . Since, for any $v \in \mathcal{V}^n$, $|\nabla g(v)| < 1$ for n sufficiently large, it follows that

$$\rho(g(w), g(z)) < \rho(w, z).$$

Thus, there exists an $M \in \mathbb{N}$ such that for all $n > M$, the function g is a contraction mapping. \square

The final proposition and corollary show that we can construct a game, Γ , such that behavior in a dynamically stable voting equilibrium is consistent with Bayesian Markov-perfect Nash equilibrium behavior in Γ .

Define Γ as follows:

- There is a collection of *players*, $N = \{0, \dots, n\}$, with Player 0 assumed to be nature and the remaining players being voters.
- There is a collection of *states*, $\mathcal{S} = X \times X$, with generic element s . At a given time t , $s^t = (x^t, y^t) \in \mathcal{S}$ can be interpreted as a *status quo policy* x^t and a *proposal*, y^t , to be pitted against the status quo.
- A player's *type* is denoted $\theta_i = (\theta_{ix}, \theta_{iy})$, with $\theta = \{\theta_i\}_{i=1}^n$. Let $\omega = (s, \theta)$ be a *type profile*. For each $i > 0$, let $\Theta_i = \mathbb{R}^2$ and $\Theta = \times_{i=1}^n \Theta_i$. Let F denote the twice-continuously differentiable cumulative distribution function of a probability distribution possessing full support on Θ_i and $\bar{F} \equiv F^n$ denote the cumulative distribution function of the resulting product measure on Θ .
- A *history* at time t is a sequence of type profiles and actions, $h^t = \{\omega^0, a^1, \omega^1, \dots, \omega^{t-1}, a^{t-1}, \omega^t\}$. Let \mathcal{H}^t denote the set of all possible histories at time t and \mathcal{H} denote the set of all possible histories. Note that the histories are from the voters' perspective, so that $h^0 = \{\omega^0\}$. Below, the true initial history, $h = \emptyset$, is used only for the consideration of Nature's determination of the initial type profile, ω^0 .
- At each time t each player $i > 0$ knows history h^t and has *action space* $A_i^t(h^t) = \{0, 1\}$. An action $a_i^t(h^t) = 1$ is a vote by player i for y^t over x^t . At $t = 0$, $A_0^0(\emptyset) = X \times X \times \Theta$. At $t > 0$ Player 0 has action space $A_0^t(\{h^{t-1}\}) = X \times \Theta$, where $a_0^t = (y^t, \theta^t)$. Let $a^t(h^t)$ be a profile of the actions taken at time t .

- Each player $i > 0$ is represented by an information structure, \mathcal{I}_i , which is a partition of \mathcal{H} such that any element of \mathcal{I}_i , I_i , contains histories of exactly one length. Denoting the subset of \mathcal{I}_i containing histories of any length t by \mathcal{I}_i^t and denoting the type of player i in period t following history h by θ_i^h , the information structure is assumed to satisfy the following condition. For any two histories h^t and \hat{h}^t , h^t and \hat{h}^t are in the same element of \mathcal{I}_i^t if and only if

1. $h^{t-1} = \hat{h}^{t-1}$,
2. $s^t = \hat{s}^t$, and
3. $\theta_i^h = \theta_i^{\hat{h}}$.

In words, players are assumed to observe all past actions and type profiles, the current state, and their own current type, but not other players' current types. (Note that the assumption that players observe each others' past types is unimportant, as I will be examining equilibrium strategies that do not depend upon any players' past types.)

- A *strategy* for player i is a mapping from information sets into the space of probability distributions over actions, denoted $\sigma_i^t : \mathcal{H}^t \rightarrow \Delta(A_i^t(h_t))$ for $i > 0$. Let \tilde{Q} be the probability measure generated by pdf Q . It is assumed that $\sigma_0^t(h^t) = \tilde{Q} \times \bar{F}$ at all $t > 0$, and that for $t = 0$, $\sigma_0^0(\emptyset) = \tilde{Q} \times \tilde{Q} \times \bar{F}$. σ^t denotes a strategy profile at time t .
- Payoffs for each player are defined by the value functions $v_i(h) = \sum_{t=0}^{\infty} \delta^t [u_i(x^t) + \theta_i^t a_i^t]$.
- For $t > 0$, ω^t depends on ω^{t-1} and a^{t-1} in the following way. $x^t = x^{t-1}$ if and only if for some $C \in W$ (the collection of winning coalitions), $a_i^{t-1} = 0$ for all $i \in C$. If not, then $x^t = y^{t-1}$. y^t is chosen by nature, as in the definition of a_0^t .

We are now ready to prove the last proposition, which is straightforward given the setup of the game. It requires the definition of probabilistic voting and payoff disturbances presented in Appendix A.

Proposition 4 If individuals vote probabilistically and payoff disturbances are admissible, then at a dynamically stable voting equilibrium, v^* , the collection of functions p_i are consistent with behavior in a pure strategy Bayesian Markov-perfect Nash equilibrium of Γ .

Proof: Markov perfection requires that if two histories h^t and \hat{h}^t have the same value of type profile ω^t , then $\sigma_i(h^t) = \sigma_i(\hat{h}^t)$ for all i . In considering Markov-perfect equilibria suppose that strategies σ_i are measurable with respect to $\omega_i \equiv (x, y, \theta_i)$. In other words, players condition only on the current state (x, y) and their current type, θ_i .

Let $\phi(x, y|\sigma)$ be the probability that y defeats x , given that voters vote according to strategies σ . Let $\phi_i(x, y|a_i, \theta_i, \sigma_{-i})$ be the probability that y defeats x conditional on type realization θ_i and i 's vote choice a_i . Note that the information structure \mathcal{I}_i and fact that the θ_i 's are i.i.d. in each time period imply that $\phi_i(x, y|a_i, \theta_i, \sigma_{-i}) = \phi_i(x, y|a_i, \theta'_i, \sigma_{-i})$ for all θ_i, θ'_i . In other words, θ_i only affects i 's action; it does not affect his beliefs about other players' actions.

Clearly, ϕ is measurable with respect to ω (i.e. ϕ is history-independent). Since we are considering a monotonic game, we know that $\phi_i(x, y|1, \theta_i, \sigma_{-i}) \geq \phi_i(x, y|0, \theta_i, \sigma_{-i})$. In other words, if Player i votes for y over x , then the likelihood that y defeats x is weakly greater than it would have been had Player i voted for x over y .

For a dynamically stable voting equilibrium, v^* , define $\sigma_i^{v^*}$ as follows:

$$\sigma_i^{v^*}(\omega_i) = \begin{cases} 0 & \text{if } v_i^*(x) + \theta_{ix} > v_i^*(y) + \theta_{iy}, \\ 1 & \text{otherwise.} \end{cases}$$

All that remains to be shown is that v_i^* is equal to the i 's true expected value given σ^{v^*} . To see this, write the probability that i votes for y over x given σ^{v^*} as: $p_i^\sigma(v_i^*(x), v_i^*(y)) = \int_{\mathbb{R}^2} \sigma_i^{v^*}(x, y, \theta_i) f(\theta_i) d\theta_i$.

It can be verified that $p_i^\sigma = p_i$, as utilized in the definition of v^* (through transition probability $p(v^*(x), v^*(y))$). Thus, v_i^* is Player i 's true expected payoff conditional on strategy profile σ^{v^*} . It follows that σ^{v^*} represents a sequentially rational (i.e., Bayes Nash equilibrium) profile of strategies for the game Γ . Since $\sigma_i^{v^*}$ is by definition Markovian with respect to $\omega_i \equiv (x, y, \theta_i)$, we can conclude that, for any dynamically stable voting equilibrium, v^* , the strategy profile σ^{v^*} is a Bayesian Markov-perfect Nash equilibrium. \square

The following corollary follows immediately from the fact that when individuals vote deterministically (as defined in Appendix A) then $\theta_{ix} = 0$ for all $i \in N$ and $x \in X$.

Corollary 1 If individuals vote deterministically then at a dynamically stable voting equilibrium, v^* , the collection of functions p_i are consistent with behavior in a Markov-perfect equilibrium of Γ .

References

- [1] Jeffery Banks and John Duggan. Probabilistic voting in the spatial model of elections: The theory of office-motivated candidates. In David Austen-Smith and John Duggan, editors, *Social Choice and Strategic Decisions*. Springer, New York, 2004.
- [2] Jeffrey S. Banks. Sophisticated voting outcomes and agenda control. *Social Choice and Welfare*, 1:295–306, 1985.
- [3] David Baron. A dynamic theory of collective goods programs. *American Political Science Review*, 90(2):316–330, 1996.

- [4] David Baron and John Ferejohn. Bargaining in legislatures. *American Political Science Review*, 83(4):1181–1206, 1989.
- [5] Marco Battaglini and Stephen Coate. Inefficiency in legislative policy-making: a dynamic analysis. *National Bureau of Economic Research, Working Paper #11495*, 2005.
- [6] Marco Battaglini and Stephen Coate. A dynamic theory of public spending, taxation and debt. *Mimeo, Department of Economics, Princeton University*, 2006.
- [7] B. Douglas Bernheim, Antonio Rangel, and Luis Rayo. The power of the last word in legislative policy making. *Econometrica*, 74(3):825–, 2006.
- [8] Michael Suk-Young Chwe. Farsighted coalitional stability. *Journal of Economic Theory*, 63:299–325, 1994.
- [9] Olivier Compte and Philippe Jehiel. Bargaining over randomly generated offers: a new perspective on multi-party bargaining. *Mimeo, C. E. R. A. S.-E. N. P. C., C. N. R. S.*, 2004.
- [10] Peter J. Coughlin. *Probabilistic Voting Theory*. Cambridge University Press, Cambridge, 1992.
- [11] Bhaskar Dutta. Covering sets and a new Condorcet choice correspondence. *Journal of Economic Theory*, 44:63–80, 1988.
- [12] Anastassos Kalandrakis. A three-player dynamic majoritarian bargaining game. *Journal of Economic Theory*, 116(2):294–322, 2004.
- [13] Anastassos Kalandrakis. Proposal rights and political power. *American Journal of Political Science*, 50(2), 2006.

- [14] Adam Liptak. “Caution In Court for Gay Rights Groups.” *New York Times*, New York, NY., November 12, 2004. Late edition—Final. Section A, Column 1, page 16.
- [15] Richard McKelvey and Thomas Palfrey. Quantal response equilibria in extensive form games. *Experimental Economics*, 1:9–41, 1998.
- [16] Richard D. McKelvey and John W. Patty. A theory of voting in large elections. *Games and Economic Behavior*, 57:155–180, 2006.
- [17] Nicholas Miller. A new solution set for tournaments and majority voting: Further graph-theoretic approaches to the theory of voting. *American Journal of Political Science*, 24:68–96, 1980.
- [18] Elizabeth Maggie Penn. A model of farsighted voting and agenda formation. In *Cooperation and Social Choice: How Foresight can Induce Fairness*. Ph.D. Dissertation, California Institute of Technology, 2003.
- [19] Charles Chapman Pugh. *Real Mathematical Analysis*. Springer, 2002.
- [20] Alessandro Riboni. Committees as substitutes for commitment. *Mimeo, Département des Sciences Économiques, Université de Montréal*, 2005.
- [21] Kevin Roberts. Condorcet cycles? A model of intertemporal voting. *Social Choice and Welfare*, 29:383–404, 2007.
- [22] Norman Schofield. Local political equilibria. In David Austen-Smith and John Duggan, editors, *Social Choice and Strategic Decisions*. Springer, New York, 2004.
- [23] Thomas Schwartz. Cyclic tournaments and cooperative majority voting. *Social Choice and Welfare*, 7:19–29, 1990.

- [24] Kenneth Shepsle and Barry Weingast. Uncovered sets and sophisticated voting outcomes with implications for agenda institutions. *American Journal of Political Science*, 28(1):49–74, 1984.
- [25] John von Neumann and Oskar Morgenstern. *Theory of Games and Economic Behavior*. Princeton University Press, Princeton, NJ, 2nd edition, 1947.

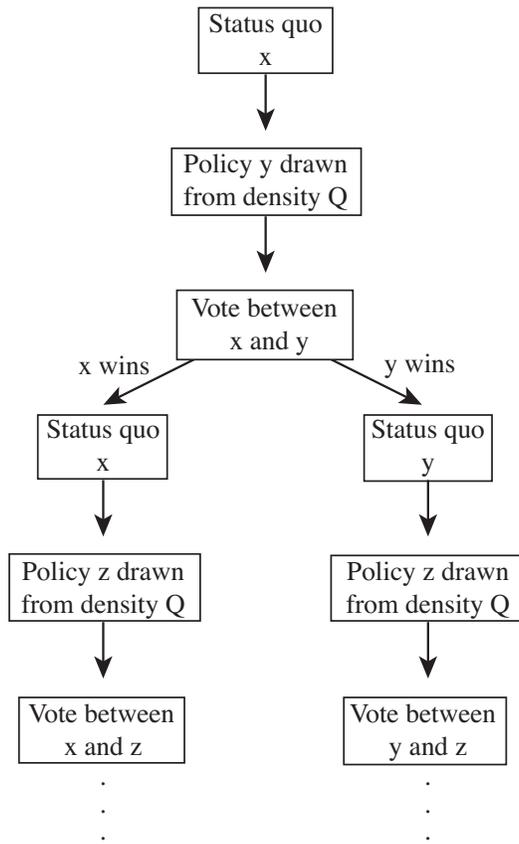


Figure 1: Path of play

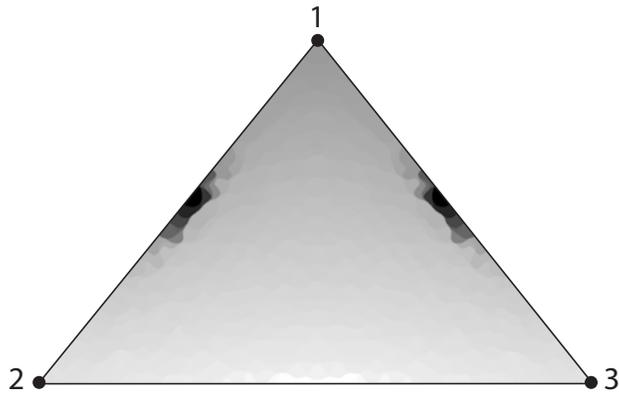


Figure 2: Player 1's value function.

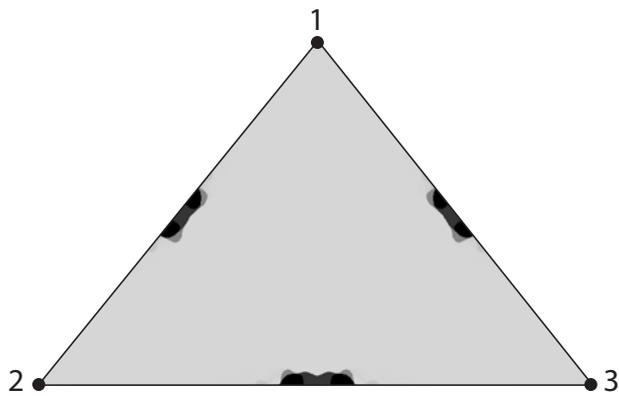


Figure 3: Density over outcomes.

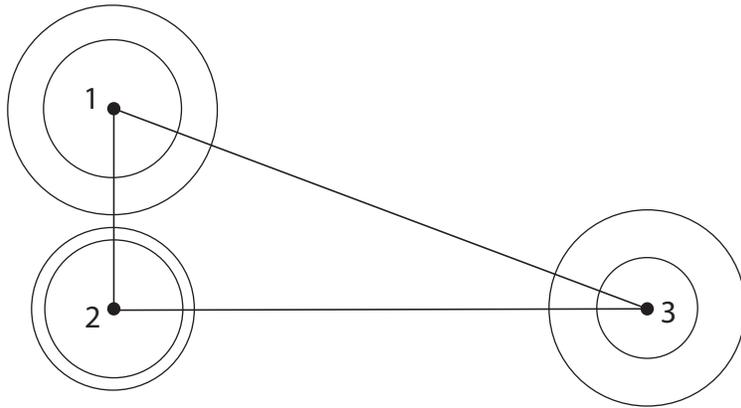


Figure 4: Two-dimensional spatial model with circular preferences.

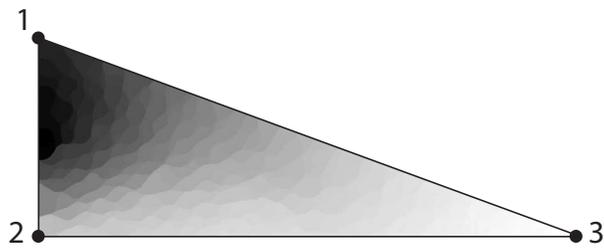


Figure 5: Player 1's value function with circular preferences.

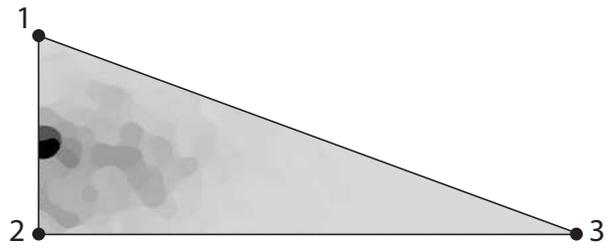


Figure 6: Density over outcomes with circular preferences.

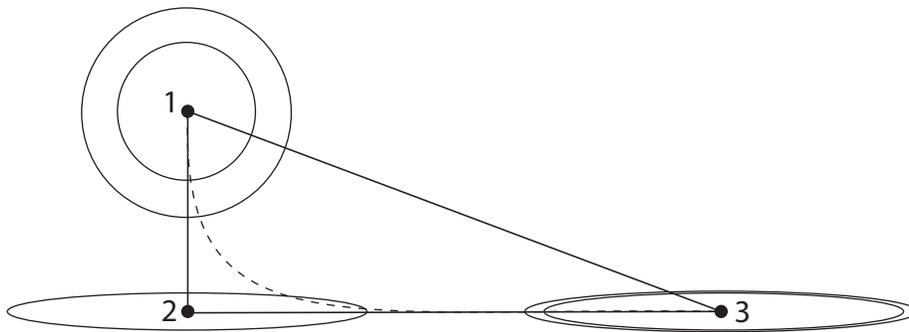


Figure 7: Two-dimensional spatial model where Players 2 and 3 have elliptical preferences.

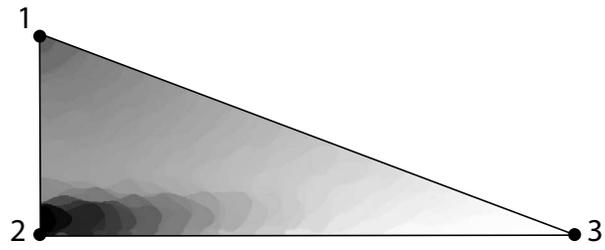


Figure 8: Player 1's value function when Players 2 and 3 have elliptical preferences.

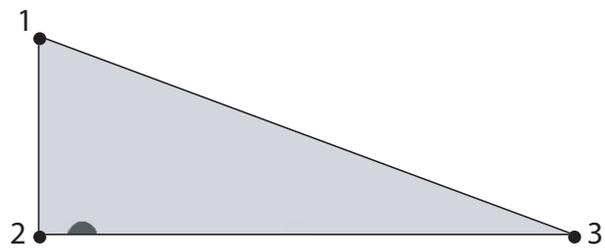


Figure 9: Density over outcomes when Players 2 and 3 have elliptical preferences.