Policy Competition in Real-Time[†]

Faruk Gul

and

Wolfgang Pesendorfer

Princeton University

February 2016

Abstract

We formulate a dynamic model of party competition with a one-dimensional policy space. Policy choices at different times are linked because parties cannot change their policies abruptly; instead, policy adjustment happens gradually. Parties are uncertain of the median's policy preferences at the time they choose their policies. Our results relate the steady state equilibrium of the dynamic game to parties' beliefs about voter preferences. In particular, we show that for symmetric games, the steady-state outcome is the local equilibrium of the corresponding static Wittman game.

[†] We thank Martin Osborne and John Duggan for their very helpful comments on an earlier draft.

1. Introduction

In this paper, we provide a dynamic version of Wittmans's (1983) model of policy competition that incorporates a constraint on parties' adjustment rates of their policy positions. Our work is motivated, in part, by evidence that Adams and Somer-Topcu (2009) provide for the *Party Dynamics Hypothesis* (PDH); that is, the hypothesis that political parties respond to rival parties' policy changes. In their empirical test of PDH, Adams and Somer-Topcu distinguish between this strategic response and a party's own response to the change in public opinion that instigated its rival's policy shift and identify the former with PDH. Our model not only provides a framework for analyzing PDH, but also for understanding how the anticipated response of a party affects its rivals initial reaction to the public opinion, or more broadly, to changes in voter preferences.

The Downs-Hotelling model of two-party competition yields a counterfactual result: in the unique equilibrium, the two parties offer identical policy positions that equal the median's preferred choice. To bring models of spatial competition inline with the empirical evidence, Wittman (1983) modifies the Downs-Hotelling model in two ways.[1] He assumes that parties are uncertain about the location of the median voter when they choose their policies and that parties care about policy. We refer to a static candidate competition model with these features as a *Wittman game*. In a Wittman game, the uncertainty about the location of the median voter lessens a party's incentive to move closer to the opponent's policy. The resulting equilibrium features two distinct policy positions that depend on the parties' preferences and on the underlying uncertainty about voter preferences.

To model the constraint on parties' rates of policy adjustment, we formulate our dynamic model in continuous time and assume that parties must choose *policy trajectories*. Our central hypothesis is that parties cannot change policies abruptly but instead must adjust them gradually. The time derivative of a policy trajectory measures the speed of change, and our requirement of "no abrupt changes" translates into an upper bound on this derivative. We consider two versions of our dynamic model: in the first version, elections occur at the same frequency as policy adjustments; in the second version, elections are infrequent relative to policy adjustments.

[1] Wittman (1983) builds on the probabilistic voting models of Hinich and Ordeshook (1969), Hinich, Ledyard and Ordeshook (1972,1973), and others. See, Duggan (2014) for a survey.

There is evidence suggesting that political parties' ideological positions change little between adjacent elections (Dalton and McAllister (2014)). One reason for this stability may be that parties occupy equilibrium positions that need little adjustment between elections. A second reason may be that parties simply cannot adjust their ideological position very abruptly. Communicating changes in a party's ideological position to activists and voters is costly and time consuming. Furthermore, a party may lose many of its activists and volunteers after breaking with its past ideological position, rendering it uncompetitive were it to attempt such a shift. Finally, voters may find large and abrupt ideological changes untrustworthy and, instead, assume that past positions offer a more credible prediction of the party's policies. Our model summarizes these factors with a simple constraint on the speed of adjustment. Adjustment speeds may reflect a parties' resources; that is, parties with greater financial resources and larger pools of activists may find it easier to adjust their positions than parties that lack these resources. For much of the analysis, we assume that parties face identical adjustment constraints, but we discuss asymmetric adjustment speeds in the final section of paper.

In our dynamic model, parties are uncertain of the median's policy preferences at the time they choose their policies. Specifically, parties' believe that the median's ideal policy is distributed on the interval $[-1, 1]$ according to a symmetric distribution Γ that admits a continuous density. The main result of the paper characterizes a steady state equilibrium of the dynamic game and relates this equilibrium to parties' beliefs about voter preferences. We show that what matters for policy outcomes is $\Gamma'(0)$; that is, the density of Γ at 0. Thus, only a narrow aspect of (parties' beliefs about) voter preferences affects policy outcomes: the likelihood that the median voter prefers the most moderate policy (policy 0).

Duggan (2014) shows that the static Wittman games has a pure strategy Nash equilibrium if Γ is log-concave.[2] Our dynamic game has a pure strategy equilibrium even if this condition is violated. Moreover, this equilibrium converges to a steady state in which parties choose policies that are local (but not global) pure strategy equilibria of the static game. If parties could change their policies abruptly, the local equilibrium would provide

[2] See Duggan (2014), Theorem 14. Duggan does not assume the game is zero sum. However, his argument extends to our model.

incentives to deviate. In our dynamic model, parties cannot realize these gains from deviating because of the interaction of two of its features: parties must adjust their policies gradually and their opponents can react quickly to these adjustments. These two features, together, imply that the rival party's policy adjustments will neutralize any gains from a deviation. Thus, our main result establishes that the unique local equilibrium of the static game (which always exists) is also the unique steady state of the frequent elections version of our model.

For the infrequent elections version of our model, we show that the unique local equilibrium is *a* steady state and, when parties are sufficiently office motivated, it is the only steady-state outcome. Finally, we investigate what happens when one party can devote much greater resources to policy adjustments than its rival; that is, we assume that one party can adjust much more quickly than the other. We establish a lower bound for the advantaged party's equilibrium payoff and show that this bound is close to what the advantaged party would receive in the equilibrium of a two-stage version Wittman's model in which the advantaged party chooses its policy after observing the choice of its rival.

In a multi-dimensional policy setting, Schofield (2001) and Schofield and Sened 2002) analyze local equilibria and provide conditions for their existence. Our model complements Schofield's approach by establishing conditions under which a local equilibrium is the unique outcome of a dynamic policy game. Kramer (1977) analyzes policy trajectories of a model in which successive policies are chosen to maximize their vote share. He characterizes the limit points of this dynamic process in a multi-dimensional policy space. While Kramer's (1977) setting, purpose and dynamic model are quite different than ours, both his approach and ours provide a common insight: a dynamic model often leads to predictions that are sharper and less permissive than the corresponding static model. Finally, our dynamic model is related to Anderson (1984) and Bergin (2006) who introduce inertia and quick response to dynamic games. Bergin (2006) analyzes repeated games in continuous time and proves a folk-theorem for those games. In contrast to their models, we analyze a zero-sum game which yields a unique equilibrium payoff and, under certain conditions, a unique equilibrium outcome.

2. Wittman Games

Consider the following one-shot policy game: there are two parties, $1, 2$. Party 1 prefers policies to the left and chooses a policy $x \in [-1, 0]$ while party 2 prefers policies to the right and chooses a policy $y \in [0, 1]$. At the time parties choose policies, the median's ideal point is uncertain and distributed on the interval $[-.5, .5]$. The realized median elects the party that has chosen a policy closer to her ideal point. Let Γ be the restriction of the cumulative distribution function of the median's distribution to the interval $[-.5, .5]$.

For a given pair of policies, x, y, let $m = (x + y)/2$ be the policy midpoint. If o is the realized median, then party 1 wins if $o < m$ and party 2 wins if $o > m$. Otherwise, the election is a tie. If there is a winner, the winning party's policy is implemented; if the election is tied; that is, if $o = m$ or if $x = y = 0$, then each party's policy is implemented with probability .5. When policy z is implemented, party 2's payoff is $\hat{v}(z) + \beta$ if it wins and $\hat{v}(z) - \beta$ if it loses while party 1's payoff is $-\hat{v}(z) + \beta$ if it wins and $-\hat{v}(z) - \beta$ if it loses. The parameter $\beta \geq 0$ measures the payoff from holding office. Therefore, party 2's (expected) payoff from the policy profile (x, y) is

$$\pi(x, y) = \hat{v}(x) \cdot \Gamma(m) + \hat{v}(y) \cdot (1 - \Gamma(m)) + \beta \cdot (1 - 2\Gamma(m)) \tag{1}$$

We assume that the game is symmetric. Specifically, symmetry requires that

$$\begin{aligned} \hat{v}(z) &= -\hat{v}(-z) \\ \Gamma(z) &= 1 - \Gamma(-z) \end{aligned} \tag{2}$$

We take advantage of the symmetry of $\hat{v}$ and simplify the notation as follows. Relabel party 1's strategy by defining it as the distance from 0. That is, we will replace x with $-x$ and let both players choose strategies in $[0, 1]$. Then, the expected payoff of party 2 becomes

$$u(x, y) = -v(x) \cdot \Gamma(\Delta) + v(y) \cdot (1 - \Gamma(\Delta)) + \beta \cdot (1 - 2\Gamma(\Delta)) \tag{1'}$$

where $\Delta = (y - x)/2$ and v is the restriction of $\hat{v}$ to $[0, 1]$. We call v the policy valuation and Γ the median distribution. Note that after this relabeling, a strategy $x, y \in [0, 1]$ can be identified with the degree of partisanship; $x = 0$ corresponds to the most moderate policy while $x = 1$ corresponds to the most partisan policy. A symmetric strategy profile (x, y) with $x = y$ does not represent a situation in which the two parties choose identical policies unless $x = y = 0$. Rather, it represents a situation in which the two parties' policies are *equally partisan*; that is, equally appealing to the mean of the median voter distribution.

We say that a function on set $A \subset I\!R^n$ is *smooth* if it has a twice continuously differentiable extension to an open set that contains A. Throughout, we assume that Γ and v are smooth, Γ is symmetric, v is strictly concave (i.e., $v'' < 0$), $\Gamma' \geq 0, v' > 0$. If these properties are satisfied, we call the u defined in (1') above a *Wittman game.* We let $\mathcal{U}$ denote the set of Wittman games. Henceforth, we will simply say u or the game u when we mean the one-shot zero-sum game $([0, 1], [0, 1], -u, u)$ for $u \in \mathcal{U}$. We let $\omega = (\omega_1, \omega_2)$ denote a generic (action) profile; that is, an element of the unit square $\Omega := [0, 1] \times [0, 1]$.

Definition 1: *Profile ω is an equilibrium of u if $u(x, \omega_2) \geq u(\omega_1, \omega_2) \geq u(\omega_1, y)$ for all (x, y); it is a strict equilibrium if both inequalities are strict whenever $(x-\omega_1)\cdot(y-\omega_2) \neq 0$.*

Call $u \in \mathcal{U}$ regular if $u_{22}(x, y) < 0$ and $-u_{22}(x, y) \leq |u_{12}(x, y)|$ for all $x, y \in [0, 1]$ where u_i is the partial derivative of u with respect to its i'th argument and u_{ij} is the partial derivative of u_j with respect to its i'th argument. Hence, u is regular implies u is strictly concave (convex) in y (x) and that players' best response functions have slopes no greater than 1 in absolute value.

Lemma 1: *Every regular game has a unique equilibrium. This equilibrium is symmetric and strict.*

Proof: Since u is strictly concave in its second argument, for each x, there is a unique $\phi(x)$ that maximizes $u(x, \cdot)$. The function ϕ maps the unit interval into the unit interval and, by the theorem of the maximum, is continuous. Therefore, ϕ has a fixed-point x_*. Since u is symmetric, (x_*, x_*) is a symmetric equilibrium. Suppose $(x, y) \neq (x_*, x_*)$ is also an equilibrium. Since the best response ϕ is single-valued, we must have $x \neq x_* \neq y$. But since u is a zero-sum game, (x, x_*) must also be an equilibrium, contradicting the single valuedness of ϕ. Hence, (x_*, x_*) is the only equilibrium. That (x_*, x_*) is strict follows from the fact that u is strictly concave in its second argument and strictly convex in its first.$\square$

A non-regular Wittman game may have no equilibrium. In particular, if Γ is not log-concave, then u need not have a (pure strategy) equilibrium. However, as we demonstrate in Lemma 2 below, all Wittman games have a unique *local* equilibrium.

Definition 2: *Profile ω is a local equilibrium of u if there is $\epsilon > 0$ such that $u(x, \omega_2) \geq u(\omega_1, \omega_2) \geq u(\omega_1, y)$ for all (x, y) in an $\epsilon-$neighborhood of ω; it is a strict local equilibrium if both inequalities are strict whenever $(x - \omega_1) \cdot (y - \omega_2) \neq 0$.*

Lemma 2: *A Wittman game has a unique local equilibrium. This local equilibrium is symmetric and strict.*

Proof: First, we will establish the existence of a symmetric strict local equilibrium. Let u_2 denote the derivative of u with respect to its second argument and let

$$\psi(x) := u_2(x, x) = \frac{1}{2}v'(x) - (v(x) + \beta)\Gamma'(0)$$

Define $\psi^*(x) := \max\{0, \min\{1, x + \psi(x)\}\}$. The assumed properties of v imply that ψ is continuous and decreasing. Therefore, ψ^* must have a unique fixed-point x_*. At that fixed-point, one of the following must hold: (i) $\psi(0) < 0$ and $x_* = 0$; (ii) $\psi(1) > 0$ and $x_* = 1$ or (iii) $\psi(x_*) = 0$. If the unique fixed-point x^* satisfies (i) or (ii), then clearly (x_*, x_*) is a symmetric strict local equilibrium. If (iii) holds, let u_{22} denote the second derivative of u with respect to its second argument. Since Γ is symmetric, $\Gamma''(0) = 0$ and therefore,

$$u_{22}(x_*, x_*) = \frac{1}{2}v''(x_*) - v'(x_*)\Gamma'(0) < 0$$

ensuring that x_* is a local best response to x_* and therefore, (x_*, x_*) is a symmetric local strict equilibrium.

To conclude the proof, we will show that there are no asymmetric local equilibria. Consider any $y \neq x$. Without loss of generality assume $y > x$. Note that

$$u_2(x, y) = v'(y)(1 - \Gamma(\Delta)) - \frac{1}{2}(v(y) + v(x) + 2\beta)\Gamma'(\Delta)$$
$$u_1(x, y) = -v'(x)\Gamma(\Delta) + \frac{1}{2}(v(y) + v(x) + 2\beta)\Gamma'(\Delta)$$

Since $x < 1$ and $y > 0$, the first order conditions imply

$$0 \leq v'(y)(1 - \Gamma(\Delta)) - \frac{1}{2}(v(y) + v(x) + 2\beta)\Gamma'(\Delta)$$
$$0 \geq v'(x)\Gamma(\Delta) - \frac{1}{2}(v(y) + v(x) + 2\beta)\Gamma'(\Delta)$$

and therefore,

$$v'(y)(1 - \Gamma(\Delta)) \geq v'(x)\Gamma(\Delta)$$

But $y > x$ implies $\Gamma(\Delta) > 1/2$ and $v'(y) \leq v'(x)$, contradicting the above inequality. □

Note that symmetry, in our notation, means that parties are choosing equally partisan policies and hence, each party will win the election with the same probability $1/2$. If $x = y = 0$, then both parties are choosing the same, moderate, policy. However, if $0 < x = y$, then the party's policies differ; party 1 chooses the policy $z_1 = -x$ while party 2 chooses the policy $z_2 = x$.

To illustrate our results, we will use the following example throughout the paper. Let

$$\Gamma(\Delta) = \frac{1}{2} + 4\Delta^3 \tag{3}$$

The cumulative Γ defined in (3) above has the property that $\Gamma'(0) = 0$. For any policy valuation v, the Wittman game with Γ as defined in (3) is not regular. The local equilibrium is $\omega_* = (1, 1)$ and, thus, parties choose their most partisan policies. Note that in this equilibrium neither party has a local incentive to moderate because $\Gamma'(0) = 0$; that is, at symmetric profiles the probability of winning is locally unresponsive to changes in the policy. However, this local equilibrium is not an equilibrium: for $\beta > 0$, the best response to $x = 1$ is a moderate policy $y < 1$. Thus, the game has no (pure strategy) equilibrium. Note that if β is large, parties care almost exclusively about being in office. The limiting case (as $\beta \to \infty$) is not a Wittman game but a standard Downs-Hotelling game in which parties care only about winning. That game has a unique equilibrium in which parties choose $x = y = 0$.

This example illustrates that Wittman games may fail to have equilibria and, in addition, that a small policy preference can lead to a local equilibrium that radically departs from the standard Downs-Hotelling prediction of policy convergence.

3. Dynamic Policy Games

In this section, we define a dynamic Wittman game in continuous time. In such a game, parties choose policies at every moment in time and their instantaneous payoffs are as described in the previous section. We constrain the speed with which parties can adjust their policies by assuming that policy trajectories must be Lipschitz continuous with a fixed Lipschitz constant. What matters in our setting is not the absolute adjustment speeds but the relative adjustment speeds of the two parties. We begin with the assumption that the two parties can adjust policies at the same maximal speed; that is, their strategies must be Lipschitz continuous with respect to the same constant. We analyze the case of asymmetric speeds in the section 6.

The continuous time formulation implies that an election takes place every instant. We can interpret this as a setting in which national parties choose an ideological position and contest many different elections at the national, state, and local level. However, the assumption of frequent elections is inessential for our results. In section 5, we assume *infrequent elections* and establish results similar to those we obtained with frequent elections. Thus, the continuous-time setting is important only in so far as it allows us to capture party positioning in real time.

Parties' strategies are real-valued functions, chosen every period k, that specify a continuous evolution of their policies over the time interval $k\lambda$ to $(k+1)\lambda$. We assume that parties choose these functions sequentially so that the second mover observes the opponent's choice. To preserve symmetry, the identity of the first mover changes in every period.

Fix an initial policy profile $o \in \Omega$, a utility $u \in \mathcal{U}$, a period-length $\lambda > 0$ and a common discount factor $e^{-\lambda r}$. At the start of each period, parties have their current positions, $\omega \in \Omega$. In odd periods, $k = 1, 3, \ldots$, party 1 moves first and decides how it will adjust its policy during the next λ units of real-time. Hence, it chooses a function $f : [0, \lambda] \to [0, 1]$ such that $f(0) = \omega_1$ and

$$|f(\tau) - f(t)| \leq |\tau - t| \tag{4}$$

for all $\tau, t \in [0, \lambda]$. Inequality (4) places bounds the speed of parties' policy adjustments. Let H denote the set of all functions f that satisfy inequality (4) and let $H(x)$ be the subset of that set with $f(0) = x$. After observing party 1's adjustment function, party 2 chooses $\hat{f} \in H(\omega_2)$. Parties reverse the order of moves in even periods.

There are two types of histories, histories that mark the beginning of a period and histories that mark the middle of a period (when the first but not the second mover has chosen). We refer to the first type of histories as *tight* histories and write P^k for the tight $k-$period histories. We call histories of the second type, *flush* histories and let Q^k the set of flush $k-$period histories. We consider the initial policy o a part of every history. In particular, $P^0 = \Omega, Q^0 = H$ are the tight and flush period $0-$period histories.

For any two functions $g, \hat{g} : [0, \tau] \to [0, 1]$, let $[g\hat{g}]$ denote the function $q : [0, \tau] \to \Omega$ such that $q(t) = (g(t), \hat{g}(t))$. Then, a tight $1-$period history is a function $[f_1^1 f_2^1]$ and $P^1 = H \times H$ are the tight $1-$period histories. The policy profile at the beginning of period 0 is $[f_1^1 f_2^1](0)$; the policy profile at the beginning of period 1 is $[f_1^1 f_2^1](\lambda)$. For $k \geq 1$,

$$P^k = \left\{ (f_1^j, f_2^j)_{j=1}^k \in (H \times H)^k : f_i^j(0) = f_i^{j-1}(\lambda) \, \forall j = 2, \ldots k \, \forall i = 1, 2 \right\}$$

We identify $p \in P^k$ with the function $p : [0, k\lambda] \to \Omega$ such that $p(\tau) = [f_1^j f_2^j](t)$ for $\tau = j\lambda + t$ and define $\bar{\omega}(p) = p(k\lambda)$ as the policy at the end of that history. For $k \geq 1$, the flush $k-$period histories are

$$Q^k = \begin{cases} \{(p, f) : p \in P^k, f \in H(\bar{\omega}_1(p))\} & \text{if } k \text{ is even} \\ \{(p, f) : p \in P^k, f \in H(\bar{\omega}_2(p))\} & \text{if } k \text{ is odd} \end{cases}$$

For $(p, f) \in Q^k$ we let $\bar{\omega}(p, f) = \bar{\omega}(p)$. Then, let $P_1 = \bigcup_{k \geq 0} P^{2k}, Q_1 = \bigcup_{k \geq 0} Q^{2k+1}, P_2 = \bigcup_{k \geq 0} P^{2k+1}$, $Q_2 = \bigcup_{k \geq 0} Q^{2k}$. The set of all histories after which player i moves is $M_i = P_i \cup Q_i$. Finally, let $P := P_1 \cup P_2$, $Q := Q_1 \cup Q_2$, $M := M_1 \cup M_2$.

The strategy set of player i is:

$$\Sigma_i = \{\sigma_i : M_i \to H | \, \sigma_i(\nu) \in H(\bar{\omega}_i(\nu))\}$$

and $\Sigma = \Sigma_1 \times \Sigma_2$. A *trajectory* is a function $\theta = (\theta_1, \theta_2) : I\!R_+ \to \Omega$ such that

$$|\theta_i(\tau) - \theta_i(t)| \leq |\tau - t|$$

for $i = 1, 2$. We use Θ to denote the set of trajectories and, for $\theta \in \Theta$, we let θ^k be the implied $k-$period history. That is, θ^k is the element of P^k such that $\theta^k(t) = \theta(t)$ for all $t \leq k\lambda$. Let $\Theta(o) = \{\theta \in \Theta \,|\, \theta(0) = o\}$ be the set of feasible trajectories given initial policies o. We say that the strategy profile $\sigma \in \Sigma$ induces the trajectory $\theta \in \Theta(o)$ given the policy profile o if

$$\theta^{n+1} = (\theta^n, [\sigma_1(\theta^n)\sigma_2(\theta^n, \sigma_1(\theta^n))])$$
$$\theta^{n+2} = (\theta^{n+1}, [\sigma_1(\theta^{n+1}, \sigma_2(\theta^{n+1}))\sigma_2(\theta^{n+1})])$$

for all $n = 2k$, $k = 0, 1, \ldots$. It is easy to verify that given any history $\nu \in M$, a strategy profile induces a unique trajectory $\theta_{\nu\sigma}$.

At each time t an election is held. A party's policy choice at time t determines its chances of winning the election and also represents the policy the party implements should it win the election. Therefore, player 2's payoff for the trajectory θ is

$$U(\theta) = r \cdot \int_0^\infty u(\theta(t))e^{-rt}dt$$

and player 1's payoff is $-U(\theta)$. Henceforth, "payoff" means player 2's payoff. A strategy profile $\sigma = (\sigma_1, \sigma_2)$ is a (subgame perfect) equilibrium if

$$U(\theta_{\nu\hat{\sigma}}) \leq U(\theta_{\nu\sigma}) \leq U(\theta_{\nu\sigma'})$$

for all $\hat{\sigma} = (\sigma_1, \hat{\sigma}_2)$, $\sigma' = (\sigma_1', \sigma_2)$, $\sigma_1' \in \Sigma_1$, $\hat{\sigma}_2 \in \Sigma_2$ and all $\nu \in M$. We call (u, r, o, λ) a discrete game. The theorem below establishes the existence of an equilibrium for the discrete game. Its proof and all of the remaining proofs are in the appendix.

Theorem 1: *Every discrete-time game has an equilibrium and a unique equilibrium payoff.*

The uniqueness of the equilibrium payoff in Theorem 1 is an implication of the zero sum payoffs, the existence of a pure strategy equilibrium relies on the sequential structure of the game.

4. Policy Choice in Real-Time

In this section, we analyze a version of the policy game described above in which parties can react to each other's actions in real time; that is, when parties can adjust their policy trajectories very frequently. Formally, a sequence of discrete games, (u, r, o, λ_n) *converges to real time* if $\lim \lambda_n = 0$. The trajectory $\theta \in \Theta(o)$ is a *real-time trajectory* of (u, r, o) if, for every sequence games (u, r, o, λ_n) converging to real time, there exists a corresponding sequence of equilibria $\{\sigma^n\}$ such that $\{\theta_{o\sigma^n}\}$ converges to θ.[3] Similarly, we say that v is a real-time payoff of (u, r, o) if, for every sequence of games (u, r, o, λ_n) converging to real time, there exists a corresponding sequence of equilibrium payoffs converging to v.

Real-time trajectories reflect situations in which parties can quickly respond to opponent's policy changes but the rate of their policy changes is constrained; that is, when party 1 moves to a more partisan policy, party 2 can respond to this move almost instantaneously, but the rate of the move and the rate of the response are at most 1. Let $\|\omega - \omega'\| = \max\{|\omega_1 - \omega_1'|, |\omega_2 - \omega_2'|\}$.

Definition 4: *The policy profile ω_* is a steady-state if there is $\epsilon > 0$ such that $\|\omega_* - o\| \leq \epsilon$ implies there is a real-time trajectory $\theta \in \Theta(o)$ with $\theta(t) = \omega_*$ for all $t \geq 1$.*

If, in addition, the real time trajectory is unique, then ω_* is a local *strong* steady state. If the ϵ in the definition above can be set to 1, then ω_* is a *global steady-state* or a *strong global steady state* if real time trajectory is unique. Theorem 2, below, shows that if u is a regular Wittman game, then every real time trajectory converges to its unique equilibrium.

Theorem 2: *The unique equilibrium of a regular game is the strong global steady-state.*

The proof of Theorem 2 reveals that if u regular, then the corresponding the dynamic game has a unique equilibrium. Moreover, this equilibrium converges to the unique equilibrium of the static game in the minimal feasible time. That is, if $\omega_* = (x_*, x_*)$ is the unique equilibrium of the (regular) u, then at all times t such that $t \geq \|\omega - \omega_*\||$, both

[3] Convergence is in the sup norm.

parties choose the policy x_*. Even if u is not regular, the local equilibrium is a strong steady state:

Theorem 3: *The unique local equilibrium of a Wittman game is a strong steady-state.*

Theorem 3 shows that, given *any* Wittman game, if the initial state is close to the local equilibrium, then policies converge to the local equilibrium and remain there. Thus, despite the fact that the local equilibrium is not a static best response for either party, real-time competition renders any deviation unprofitable.

To illustrate Theorem 3, consider the example above with Γ as defined in equation (3) and $\beta > 1$. In the static game, at the local equilibrium $\omega_* = (1, 1)$, each party has an incentive to deviate to a moderate policy. For example, if a party deviates to policy $x = 0$, then its payoff will be β instead of 0. Thus, when parties are mostly concerned about holding office (i.e., β is large), and parties do not have the opportunity to respond to their opponent's policy changes (i.e., in a static game) the local equilibrium cannot be sustained. However, in the real-time game, the local equilibrium remains the unique outcome for the following reason. If a party were to deviate to a moderate policy it would have to do so gradually and its opponent could react to the deviation. Hence, in the dynamic game the question is whether or not the party can achieve a *policy trajectory* that increases its payoff. In Lemma 3, in the appendix, we establish the following property for all Wittman games: when parties choose policies with a similar degree of partisanship, the party closer to the local equilibrium obtains a higher payoff than the party farther from its local equilibrium action. As a consequence, the opposing party can counter a move towards the moderate policy by "following" the deviator so the gap between the degree of partisanship of the two policies remains small. Along this trajectory, the initial deviator loses and the follower gains. Thus, there is a "second mover advantage" to any departure from the local equilibrium which ensures that the local equilibrium remains the unique equilibrium outcome.

Notice that the above described second mover advantage may only hold if parties choose policies that exhibit a nearly identical degree of partisanship. For this reason, Theorem 3 above holds in the real time limit but may fail away from it. When λ is large, a deviating party may create sufficient distance between it and the opponent to gain the advantage.

5. Infrequent Elections

Next, we introduce a variation of our model in which elections take place at fixed, discrete intervals. We let elections take place at integer valued times $\ell = 1, 2, \ldots$. As in the previous two sections, parties adjust their policies continuously and their speed of adjustment is constrained. However, payoffs depend only on the policies chosen at integer-valued times.

The policy trajectories between elections; that is, in the interval $(\ell, \ell+1)$ represent the parties' adjustments during an election campaign. The definition of histories and party strategies remain unchanged from the previous section. The only difference is that we assume that the period length λ is $1/\bar{m}$ where $\bar{m}$ is a positive integer. Thus, between any two elections parties can adjust their policy trajectories $\bar{m}$-times. The $\ell-$th election is based on the parties' policy choices at time ℓ and the winning party implements its policy choice at that time. Therefore, player 2's payoff from the trajectory θ is

$$\hat{U}(\theta) = \left(1 - e^{-r}\right) \sum_{\ell=1}^{\infty} u(\theta(\ell)) e^{-(\ell-1)r} \tag{5}$$

and player 1's payoff is $-\hat{U}(\theta)$. Even though parties' policy choices are payoff relevant only at election time, it is essential that parties choose their policies in real time *between* elections. This ensures that parties can react to the opponent's policy adjustments and adjust their own policies by the time the election is held.

We refer to the game with elections at integer times as the *infrequent elections game.* It has the same parameters, (u, r, o, λ), as the discrete game defined in section 3. Theorem 2A in the appendix establishes the existence of equilibrium and the uniqueness of the equilibrium payoff.[4] The definitions of real-time trajectories and of steady states remain unchanged. Theorem 4, below, shows that the local equilibrium is a steady state of the infrequent elections game.

Theorem 4: *The unique local equilibrium of a Wittman game is a global steady state of the infrequent elections game.*

[4] Theorem 2A proves existence of stationary pure strategy equilibria. Stationarity means that the strategy depends only on the current state, the opponent's policy choice and the number of periods since the last election.

In one sense, Theorem 4 yields a weaker conclusion than Theorem 3: the local equilibrium of a Wittman game is a steady state rather than a strong steady state, reflecting the non-uniqueness of equilibria with infrequent elections. The reason for the difference is that the policy trajectories between elections have no effect on payoffs and, therefore, moving away from the local steady state has no short-run payoff consequences. On the other hand, there is a sense in which Theorem 4 strengthens Theorem 3: the local equilibrium is a global a steady state; that is, it is a steady state irrespective of the initial condition. Again, the reason is that the initial condition is not payoff relevant as long as parties have enough time between elections to converge to the steady state.

Recall that we have normalized the parties' maximal adjustment speed to 1. This normalization is without loss of generality for Theorems 2 and 3 but not for Theorem 4. Specifically, Theorem 4 would continue to hold if the adjustment speed were faster than 1 but may fail for slower adjustment speeds. With slow adjustment speeds, parties may simply be unable to reach the local equilibrium between elections. In that case, we would need to restrict the initial policy profile to an appropriate neighborhood of the local equilibrium, as in Theorem 3. Thus, our definition of the infrequent elections game implicitly assumes that parties have enough time between consecutive elections to reach any desired policy.

Theorem 5, below, gives conditions under which the local equilibrium of the static game is the unique policy that is implemented in every real-time trajectory of the infrequent elections game. As we noted above, the fact that policy choices between elections are not payoff relevant implies that real-time trajectories cannot be unique. For example, if parties start out close to the local equilibrium they may choose to linger at the initial state before converging to the local equilibrium; or, alternatively, they may converge to the local equilibrium immediately. Under appropriate conditions, both scenarios can be real-time trajectories. However, since both trajectories lead to the same election outcome, this non-uniqueness does not affect voter or party payoffs and, therefore, is inessential. If every real-time trajectory such that at each integer-valued time, both parties choose the local equilibrium, then we say that the local equilibrium is the unique *outcome*. Theorem 5 shows that this is the case when parties care mostly about being in office; that is, when β is sufficiently large.

We say that the outcome ω is the unique outcome for $(\Gamma, v, \beta) \in \mathcal{U}$ in the game with infrequent elections if, for every $r > 0$ and for every real time trajectory θ of (Γ, v, β, r),

$$\theta(\ell) = \omega \text{ for all } \ell = 1, 2, \ldots$$

Theorem 5: *For every Wittman game, there is $\bar{\beta} > 0$ such that $\beta > \bar{\beta}$ implies that the local equilibrium is the unique outcome of the game with infrequent elections.*

When parties care mostly about winning, there is little scope for trading off a reduced win probability for a more partisan policy. As a result, parties' win probabilities must be approximately equal in every election. This, in turn, implies that the two policies are of a similar level of partisanship, that is, $|x - y|$ must be small at the time of the election. As we show in the appendix, the party closer to the local equilibrium receives a higher payoff in the neighborhood of a symmetric profile. The proof of Theorem 5 uses this fact to show that the only symmetric policy profile that can be sustained as an election outcome is the local equilibrium. Thus, a large β implies that the policy profile must be symmetric which, in turn, is the basis for the uniqueness argument in the proof of Theorem 5.

Even when parties care mostly about winning, the local equilibrium may lead to partisan policies. Example (3) above, illustrates this point. In this example, the gains from a deviation to the moderate policy increase with β. Therefore, it might seem that moderate policies should be sustainable as an equilibrium outcome of the dynamic game. However, this intuition does not apply when parties compete in real time; a move to a more moderate policy leads to a corresponding move by the opposing party that ultimately results in a small utility loss without significantly altering the party's election probability.

6. Asymmetric Adjustment Speeds and Equilibrium Policies

As we observed in the introduction, the constraint on the speed of policy adjustment can be interpreted as a resource constraint. Put differently, a party's adjustment speed may be a proxy for its access to funding or the size of pool of activists. Well financed parties may find it easier to communicate policy than their poorly financed competitors and, thus, when parties differ in their access to resources we would expect them to face different constraints in their adjustment speeds.

In this section, we consider parties with different resource constraint; that is, different speeds of adjustment and investigate the effect of this asymmetry on election outcomes. As in section 2, each party can adjust its policy at discrete times (periods) and each interval between adjustments has length λ. Given a current policy positions $\omega \in \Omega$, party i chooses a function $f : [0, \lambda] \to [0, 1]$ such that $f(0) = \omega_i$ and

$$\alpha_i |f(\tau) - f(t)| \le |\tau - t| \tag{4'}$$

for all $\tau, t \in [0, \lambda]$. Inequality (4') bounds the speed of a party's policy adjustment at $1/\alpha_i$. Up to this point, we have assumed that $\alpha_1 = \alpha_2 = 1$; in this section, we assume $1 = \alpha_1 \ge \alpha_2 = \alpha$; that is, party 2 can adjust α times faster than party 1. As in the previous section, elections are infrequent; they take place at time $\ell = 1, 2, \ldots$. Thus, the payoff function is identical to the payoff function in the infrequent election game defined in Equation (5). We refer to game described above as an asymmetric resource election. Theorem 1 extends to asymmetric resource elections, as do the definitions of histories, strategies and real-time trajectories.[5]

A straightforward case is when party 2 has an overwhelming advantage in adjustment speed. In that case, party 2 can adapt optimally to party 1's choice. That is, the resulting policy choice is as in a two-stage game in which party 1 chooses first and then party 2 makes its choice after observing party 1's action. Thus, the equilibrium payoff is

$$u_0 := \min_{x \in [0,1]} \max_{y \in [0,1]} u(x, y)$$

[5] Theorem 2A, in the appendix, proves the existence of a stationary pure strategy equilibrium for the case of infrequent elections. The proof of Theorem 2A does not use the fact that party's adjustment speeds are identical and, therefore, would extend to the case of asymmetric adjustment speeds.

If the game is regular, then

$$\min_{x\in[0,1]} \max_{y\in[0,1]} u(x,y) = \max_{x\in[0,1]} \min_{y\in[0,1]} u(x,y)$$

and, therefore, the difference in adjustment speeds is irrelevant. It is only for non-regular Wittman games that the differential adjustment speeds play a role. As an illustration, consider following Wittman game

$$\begin{aligned} v(x) &= x \\ \beta &= 1/2 \\ \Gamma(\Delta) &= \frac{1}{2} + 4\Delta^3 \end{aligned} \tag{6}$$

This game is not regular and has no equilibrium. Define $\hat{x} \in [1/2, 1]$ such that $u(\hat{x}, 0) = u(\hat{x}, 1)$. For the parameters as specified in (6), $\hat{x} \approx .815$ solves the minmax problem and and $u_0 \approx .084$. Therefore, party 2 benefits from its faster adjustment speed. At $\hat{x} = .814$, party 2 is indifferent between its most moderate and its most extreme policy and both are best responses to $\hat{x}$.[6]

Our final theorem provides payoff bounds for the asymmetric resource election. Define $B_\alpha(z) = [0,1] \cap [z - \alpha/2, z + \alpha/2]$ and define:

$$u_\alpha = \min_{z\in[0,1]} \max_{y\in[0,1]} \min_{x\in B_\alpha(z)} u(x,y)$$

We can interpret the quantity u_α as the result of a three stage policy setting game. First, party 1 chooses a position z, then party 2 chooses a policy and, finally, party 1 chooses a policy in an $\alpha/2-$neighborhood of z.

Theorem 6: *Every real-time payoff of the asymmetric resource election is in* $[u_\alpha, u_0]$.

To see why party 2 can achieve the lower bound in Theorem 6, consider the following two-phase strategy for party 2: In phase 1, party 2 moves to the policy 1/2 and stays at 1/2 until phase 2. Phase 2 starts at t^*, where t^* leaves party 2 with just enough time to

[6] The utility function as specified in (6) is linear and, therefore, does not satisfy our assumption that $v'' < 0$. All the results below still hold for the utility function $v(x) = x + \epsilon(x - x^2)$ for ϵ sufficiently small. We have chosen the linear example to simplify the exposition.

get from 1/2 to any other policy in $[0,1]$ before the next election. During phase 2, party 2 moves as fast as it can to some y that maximizes $\min u(x,y)$ where the minimum is taken over all policies that party 1 can reach before the next election starting from z, its position at time t^*. Party 2 stays at policy y once it reaches it. Clearly, the best that party 1 can do against this strategy is to choose z so as to minimize $\min u(x,y)$ over x's that it can reach from z after the start of phase 2. Note that $t \geq \ell + 1 - \alpha/2 - \lambda$ where $\ell + 1$ is the time of the next election which means that phase 2 will take at most $\alpha/2 + \lambda$ units of time. That is, party 1 will have at most $\alpha/2 + \lambda$ units of time to get from its policy position z at t^* to its policy position at election time $\ell + 1$. Hence party 1's payoff is at best

$$u_\gamma = \min_{z \in [0,1]} \max_{y \in [0,1]} \min_{x \in B_\gamma(z)} u(x,y)$$

for $\gamma = \alpha + 2\lambda$. Continuity ensures that u_γ converges to u_α as λ converges to 0, yielding the bound in Theorem 6.

With the parameters as specified in (6), the lower bound in Theorem 6 is tight for α sufficiently small. That is, in any real-time trajectory player 2's payoff is exactly u_α. To see why the bound is tight in the above example, consider the following strategy for party 1. Between elections, party 1 chooses a trajectory that converges to a policy z that satisfies

$$u(z - \alpha/2, 0) = u(z + \alpha/2, 1)$$

Once party 1 reaches z, it stays there until time $\ell - \alpha/2$. Between $\ell - \alpha/2$ and ℓ, party 1 chooses a trajectory that increases its policy whenever party 2' policy is above 1/2 and reduces its policy when party 2's policy is below 1/2. Given this strategy of party 1, party 2 can do no better than u_α, as we show in the appendix. This example establishes that the lower bound of party 2's payoff in Theorem 6 is tight.

7. Conclusion

When parties compete in real time, the local equilibrium of the (static) Wittman game emerges as a steady state of the corresponding dynamic game. In our setting, local equilibria are symmetric; both parties choose equally partisan policies. In the interior case, the local equilibrium (x_*, x_*) solves the following simple equation:

$$v'(x_*) = 2\Gamma'(0)(v(x_*) + \beta)$$

Thus, the distribution of voter ideal points Γ affects the local equilibrium through the value $\Gamma'(0)$; that is, the change in party 1's win probability if it marginally moderates it's policy at a symmetric profile. The quantity $\Gamma'(0)$ is a proxy for the probability that the median turns out to be the voter who favors the most moderate policy (of either party). If this probability is small, then the local equilibrium will lead to more partisan policies (more polarization) than if this probability is large. To confirm this intuition, recall the example of section 2, with Γ (defined in equation (3)) such that $\Gamma'(0) = 0$. The local equilibrium for this example is $(1, 1)$; the most partisan policy profile. Hence, our model relates polarization to parties' estimates of the probability that the median voter will favor the most moderate policy.

Note that properties of Γ other than the slope at 0 do not affect the local equilibrium and hence do not affect the real-time steady state. For example, consider any bimodal distribution Γ such that the most likely realization of the median is in the range of moderate left policies $(-z, -z/2)$ or in the range of moderate right policies $(z/2, z)$ (where $0 < z \leq 1$) but the probability that the median is at 0 is small. In that case, the value of z is immaterial for the election outcome and, in this sense, electoral competition in real time is unresponsive to changes in voter preferences.

8. Appendix A: Proof of Theorem 1

First, we will state and prove the following version of the principle of optimality. A dynamic decision problem is a collection $\alpha = (\delta, w, b, X, T, S)$ where X and S are arbitrary sets, $\delta \in (0,1)$ is the discount factor, $w : X \times S \to I\!R$ is the *utility function*, the function b from S to the set of all nonempty subsets of X is the *feasibility function*, and $T : X \times S \to S$ is the *transition rule.*

Let $X^*(s)$ denote the set of all sequences $\{x_n\}$ in X such that $x_n \in b(s_{n-1})$ for all $n > 1$, where $s_1 = s$ and $s_n = T(x_{n-1}, s_{n-1})$. For any sequence $\xi = \{x_n\} \in X^*(s)$, define the sequence $\{s_n\}$ as follows: $s_1 = s$ and, for all $n > 1$, $s_n = T(x_{n-1}, s_{n-1})$. Finally, for $\xi = \{x_n\}$, let

$$W_\alpha(\xi, s) = \sum_n w(x_n, s_n)\delta^{n-1}$$

A *decision rule* is a function $\rho : S \to X$ such that $\rho(s) \in b(s)$ for all $s \in S$. For any decision rule ρ and $s \in S$, construct the sequence $\{(x_n, s_n)\}$ as follows: $s_1 = s$, $x_1 = \rho(s)$, $s_n = T(x_{n-1}, s_{n-1})$ and $x_n = \rho(s_n)$. Then, let $\xi_{s\rho} = \{x_n\}$.

We call the function W^* the value function if $W^*(s) = \max_{\xi \in X^*(s)} W_\alpha(\xi, s)$. Clearly, there can be at most one value function. We say that the policy ρ is optimal if $W_\alpha(\xi_{s\rho}, s) = W^*(s)$ for all $s \in S$.

Definition: *A bounded function $\hat{W} : S \to I\!R$ such that*

$$\hat{W}(s) = \max_{x \in b(s)} w(x, s) + \delta \hat{W}(T(x, s)) \text{ for all } s \in S$$

is a recursive value. A policy ρ such that $\rho(s) \in \arg\max_{x \in b(s)}[w(x, s) + \delta \hat{W}(s)]$ is unimprovable for the recursive value $\hat{W}$.

Optimality Lemma: *(i) If $\hat{W}$ is a recursive value, then it is the value. (ii) A policy is optimal if it is unimprovable for some recursive value.*

Proof: (i) Let $\hat{W}$ be a recursive value and for $s \in S$, pick $\rho(s) = \arg\max_{x \in b(s)} w(x, s) + \delta \hat{W}(T(x, s))$.

We claim that $\hat{W}(s) = W_\alpha(\xi_{s\rho}, s)$ for all s. If not, let $\epsilon = \sup_s |\hat{W}(s) - W_\alpha(\xi_{s\rho}, s)|$. Since both W_α and $\hat{W}$ are bounded, $\epsilon < \infty$. Choose s' such that $\sup_s |\hat{W}(s) - W_\alpha(\xi_{s\rho}, s)| - (1-\delta)\epsilon < |\hat{W}(s') - W_\alpha(\xi_{s'\rho}, s')|$. Then,

$$|\hat{W}(s') - W_\alpha(\xi_{s'\rho}, s')| = \delta|\hat{W}(T(x,s')) - W_\alpha(\xi_{\rho T(x,s')}), T(x,s')| \leq \delta\epsilon$$

and hence $\delta\epsilon + (1-\delta)\epsilon > \epsilon$, a contradiction.

Let $W^o(s) = \sup_{\xi \in X(s)} W_\alpha(\xi, s)$ for all s. Since w is bounded, $W^o(s)$ is well-defined and the function W^o is itself bounded.

To see that $W^o(s) \geq \hat{W}(s)$, assume the contrary and choose ϵ such that $0 < \epsilon < (1-\delta)\sup_{s\in S}[\hat{W}(s) - W^o(s)]$ and s' such that $\sup_{s\in S}[\hat{W}(s) - W^o(s)] - \epsilon < \hat{W}(s') - W^o(s')$. Let y solve $\max_{x \in b(s')}[w(x,s') + \delta\hat{W}(T(x,s'))]$. Then, $W^o(s') \geq w(y,s') + \delta W^o(T(y,s'))$ and hence $\hat{W}(s') - W^o(s') \leq \delta[\hat{W}(T(y,s')) - W^o(T(y,s')] \leq \delta \sup_{s\in S}[\hat{W}(s) - W^o(s)]$. Thus, we have

$$\sup_{s\in S}[\hat{W}(s) - W^o(s)] - \epsilon < \hat{W}(s') - W^o(s') \leq \delta \sup_{s\in S}[\hat{W}(s) - W^o(s)]$$

a contradiction.

To prove $W^o(s) \leq \hat{W}(s)$, again we assume the contrary and choose ϵ such that $0 < \epsilon < \frac{1-\delta}{2}\sup_{s\in S}[W^o(s) - \hat{W}(s)]$ and s' such that $\sup_{s\in S}[W^o(s) - W^*(s)] - \epsilon < W^o(s') - \hat{W}(s')$. Choose $\{x_n\} \in X(s')$ such that $W_\alpha(\{x_n\}, s') > W^o(s') - \epsilon$. Then, since $\hat{W}(s') \geq w(x_1, s') + \delta\hat{W}(T(x_1, s'))$ and $w(x_1, s') + \delta W^o(T(x_1, s')) \leq W_\alpha(\{x_n\}, s')$, we have $W^o(s') - \hat{W}(s') - \epsilon \leq \delta[W^o(T(x_1, s')) - \hat{W}(T(x_1, s')] \leq \sup_{s\in S}[W^o(s) - \hat{W}(s)]$. Therefore,

$$\sup_{s\in S}[W^o(s) - \hat{W}(s)] - \epsilon < W^o(s') - \hat{W}(s') \leq \delta \sup_{s\in S}[W^o(s) - \hat{W}(s)] + \epsilon$$

a contradiction. So, $\hat{W} = W^o$. By the claim above, $W_\alpha(\xi_{s\rho}, s) = \hat{W}(s) = W^o(s)$ and hence, $W^* = \hat{W} = W^o$ proving (i).

(ii) Suppose ρ is unimprovable for some recursive value. Then, by (i), it is unimprovable for W^*. Let $\{x_n\} = \xi_{s\rho}$, $s_1 = s$, $s_n = T(x_{n-1}, s_{n-1})$ and, since W^* is a recursive value,

$$W^*(s) = \sum_{n=1}^{m} w(x_n, s_n)\delta^{n-1} + \delta^m W^*(s_{m+1})$$

Since W^* is bounded, we conclude that $W^*(s) = \sum_{n=1}^{\infty} w(x_n, s_n)\delta^{n-1} = W_\alpha(\xi_{s\rho}, s)$. $\square$

Stationary Strategies

Let $j = 3 - i$ and for $i = 1, 2$, define $M_i = \Omega \cup \{(\omega, f) \in \Omega \times H \mid f(0) = \omega_j\}$. Then, let $\mathcal{D}_i$ be the set of all function $d_i : M_i \to H$ such that $d_i(\omega) \in H(\omega_i)$ and $d_i(\omega, f) \in H(\omega_i)$ for all ω, f. A strategy σ_i for player i is stationary if there is a $d_i \in \mathcal{D}_i$ such that

$$\sigma_i(p) = d_i(\bar{\omega}(p)) \text{ for } p \in P_i$$

$$\sigma_i(p, f) = d_i(\bar{\omega}(p), f) \text{ for } (p, f) \in Q_i$$

Clearly, there is a one-to-one correspondence between the set of stationary strategies of player i and $\mathcal{D}_i$.

8.1 Proof of Theorem 1:

We will prove the following stronger version of Theorem 1:

Theorem 1A: *A discrete-time game (with frequent elections) has a stationary equilibrium and a unique equilibrium payoff.*

Proof: Let $\mathcal{V}$ be the set of all bounded continuous functions on Ω. This set, endowed with the sup norm, is a complete metric space. Define

$$\begin{aligned}
J_V(f_1, f_2) &= \int_{t=0}^{\lambda} u(f_1(t), f_2(t))e^{-rt}dt + e^{-r\lambda}V(f_1(\lambda), f_2(\lambda)) \\
Q_V^1(f_1, y) &= \max_{f_2 \in H(y)} J_V(f_1, f_2) \\
Q_V^2(f_2, x) &= \min_{f_1 \in H(x)} J_V(f_1, f_2) \\
\Lambda_V^1(\omega) &= \min_{f_1 \in H(\omega_1)} Q_V^1(f_1, \omega_2) \\
\Lambda_V^2(\omega) &= \max_{f_2 \in H(\omega_2)} Q_V^2(f_2, \omega_1) \\
\Lambda_V &= \Lambda^1_{\Lambda_V^2}
\end{aligned}$$

First, we show that $\Lambda_V^2(\omega)$ is well-defined and that the function Λ_V^2 is in $\mathcal{V}$. Since every function in H is 1-Lipschitz, H is compact by the Ascoli-Arzela Theorem as are the sets $H(x)$ for $x \in [0, 1]$. The function $J_V(\cdot)$ is continuous as is the correspondences C_1, C_2

defined by $C_1(f, x) = C_2(x, f) = H(x)$. Hence, the functions $J_V(f, \cdot)$ and $-J_V(\cdot, f)$ attain their suprema on $H(x)$ and, by Berge's Maximum Theorem, the function Q^i is continuous for $i = 1, 2$. The same arguments ensure that Λ_V^1 is well-defined and continuous.

Next, we will show that $V \to \Lambda_V$ is a contraction mapping. Suppose f_2 solves $\max Q^1(f, \omega_1)$ subject to $f \in H(\omega_2)$ and f_1 solves $\min J(f, f_2)$ subject to $f \in H(\omega_1)$ for some $V \in \mathcal{V}$ and consider $V_o \in \mathcal{V}$. Then,

$$\Lambda_V^2(\omega) - \Lambda_{V_o}^2(\omega) \leq e^{-r\lambda}(V(f_1(\omega_1), f_2(\omega_2)) - V_o(f_1(\omega_1), f_2(\omega_2)) \leq e^{-r\lambda}\|V - V_o\|$$

A symmetric argument ensures $\Lambda_{V_o}^2(\omega) - \Lambda_V^2(\omega) \leq e^{-r\lambda}\|V - V_o\|$ for all ω and hence $\|\Lambda_V^2 - \Lambda_{V_o}^2\| \leq e^{-r\lambda}\|V - V_o\|$. Again by symmetry $\|\Lambda_V^1 - \Lambda_{V_o}^1\| \leq e^{-r}\|V - V_o\|$ and it follows that $\|\Lambda_V - \Lambda_{V_o}\| \leq e^{-2r\lambda}\|V - V_o\|$ proving that $V \to \Lambda_V$ is a contraction mapping. Then, by Banach's Fixed-Point Theorem, there is a unique $V \in \mathcal{V}$ such that $\Lambda_V = V$.

For the remainder of this proof, let $V_1 = V$ and $V_2 = \Lambda_{V_1}^2$. Choose D_1, D_2 such that

$$D_1(f, \omega, j) \in \begin{cases} \arg \min\limits_{f_1 \in H(\omega_1)} Q_{V_2}^1(f_1, \omega_2) & \text{if } j = 1 \\ \arg \min\limits_{f_1 \in H(\omega_1)} J_{V_1}(f_1, f) & \text{if } j = 2. \end{cases}$$

$$D_2(f, \omega, j) \in \begin{cases} \arg \max\limits_{f_2 \in H(\omega_2)} Q_{V_1}^2(f_2, \omega_1) & \text{if } j = 2 \\ \arg \max\limits_{f_2 \in H(\omega_2)} J_{V_2}(f, f_2) & \text{if } j = 1. \end{cases}$$

Let $\alpha = (\delta, b, w, X, T, S, \delta)$ be the following dynamic decision problem for party 1: $\delta = e^{-r\lambda}$, $S = \Omega \times \{1, 2\}$, $X = H$, $b(\omega, f, j) = H(\omega_2)$ and

$$w(f, \omega, j) = -r \int_{t=0}^{\lambda} u(f(t), D_2(f, \omega, j)(t))e^{-rt}dt$$
$$T(f, \omega, j) = (\bar{\omega}(f, D_2(f, \omega, j)), 3 - j)$$

It is easy to verify that $\hat{W} : S \to I\!R$ such that

$$\hat{W}(\omega, j) = \begin{cases} V_1 & \text{if } j = 1 \\ V_2 & \text{if } j = 2 \end{cases}$$

is a recursive value for α. Then, part (ii) of the Optimality Lemma implies that $\hat{W}$ is the value for α. Define ρ_1 such that

$$\rho_1(\omega, j) \in \begin{cases} \arg \min\limits_{f_1 \in H(\omega_1)} Q_{V_2}^1(f_1, \omega_2) & \text{if } j = 1 \\ \arg \min\limits_{f_1 \in H(\omega_1)} J_{V_1}(f_1, D_2(f_1, \omega, j)) & \text{if } j = 2. \end{cases}$$

By construction, ρ_1 is unimprovable for the recursive value $\hat{W}$; hence, part (ii) of the Optimality Lemma implies that ρ_1 is an optimal policy. Define the stationary strategy $d_i, i = 1, 2$ as follows: $d_i(\omega) = D_i(f, \omega, i)$ (and note that D_i is independent of the first argument in state (ω, i)) and $d_i(p, f) = D_i(f, \bar{\omega}(p), 3 - i)$. Since ρ_1 is an optimal policy, it follows that d_1 is a best response to d_2. A symmetric argument ensures that d_2 is a best response to d_1 and hence (d_1, d_2) is an equilibrium of (u, r, ω, λ) for every $\omega \in \Omega$.

To prove the uniqueness of the equilibrium payoff, assume that (σ_1, σ_2) and $(\hat{\sigma}_1, \hat{\sigma}_2)$ are two equilibrium strategies. Then, $U(\sigma_1, \sigma_2) \leq U(\hat{\sigma}_1, \sigma_2) \leq U(\sigma, \hat{\sigma}_2) \leq U(\hat{\sigma}_1, \hat{\sigma}_2)$; that is, $U(\sigma_1, \sigma_2) \leq U(\hat{\sigma}_1, \hat{\sigma}_2)$. A symmetric argument yields the reverse inequality and hence $U(\sigma_1, \sigma_2) = U(\hat{\sigma}_1, \hat{\sigma}_2)$. □

9. Proofs of Theorems 2 and 3

Let $\omega_* = (x_*, x_*)$ be the unique local equilibrium, the existence of which is guaranteed by Lemma 2. Let $N_\epsilon := \{(x, y) \in \Omega : \max\{|x - x_*|, |y - x_*|\} \leq \epsilon\}$ and, for $u \in \mathcal{U}$, let u_ϵ denote the restriction of u to N_ϵ.

Lemma 3: *Let $u \in \mathcal{U}$. Then, there exists $\epsilon > 0$ such that (i) u_ϵ is regular and (ii) $(x - x_*)(y - x_*) \geq 0, |x - x_*| > |y - x_*|$ and $\epsilon \geq |y - x|$ imply $u(x, y) > 0$.*

Proof: (i) Note that $u_2(x, y) = v'(y)(1 - \Gamma(\Delta)) - (v(x) + v(y) + 2\beta)\Gamma'(\Delta)/2$ and, therefore, $\Gamma' > 0, \Gamma''(0) = 0, \Gamma(0) = 1/2$ imply

$$u_{12}(x, x) = v'(x)\Gamma'(0)/2 - v'(x)\Gamma'(0)/2 + \Gamma''(0)(v(x) + v(y) + 2\beta) = 0$$
$$u_{22}(x, x) = v''(y)/2 - v'(y)\Gamma'(0) < 0$$

Since $v'', v', \Gamma', \Gamma''$ are continuous functions, part (i) follows.

(ii) Choose $\epsilon_o > 0$ such that u_{ϵ_o} is regular. Suppose $|x - x_*| < \epsilon_o$ and the hypotheses (of part (ii) of the lemma) are satisfied for $\epsilon = \epsilon_o$. Then, $u(x, x_*) > 0$ and $u(x, x) = 0$ and, therefore, concavity of u implies that $u(x, y) > 0$.

To conclude the proof, we will (1) find $\epsilon_1, \epsilon_2 \in (0, \epsilon_o]$ such that if the hypotheses are satisfied for $\epsilon = \epsilon_1$ ($\epsilon = \epsilon_2$) and $x \leq x_* - \epsilon_o$ ($x \geq x_* + \epsilon_o$), then $u(x, y) > 0$ and (2) set $\epsilon = \min\{\epsilon_1, \epsilon_2\}$.

Since the two cases are symmetric, we will only consider $x \leq x_* - \epsilon_o$. Recall that $u_2(x,x) = v'(x)/2 - (v(x) + \beta)\Gamma'(0)$ and since $u_2(x_*, x_*) = 0$, the concavity and monotonicity of v implies $u_2(x,x) > 0$ for $x < x_*$. Let $l_o = [0, x_* - \epsilon_o]$ and

$$I = \{(x,x) : x \in l_o\}$$
$$C = \{(x,y) : x \in l_o \text{ and } u_2(x,y) = 0\} \cup \{(x_* - \epsilon_o, x_*)\}$$

Since u_2 is continuous, C is compact. Hence, I, C are nonempty, disjoint compact sets. It follows that $\epsilon_1 := \min\{\|\omega - \omega'\| : \omega \in I, \omega' \in C\} > 0$. Clearly, $\epsilon_1 \leq \epsilon_o$ and, since $u_2(x,x) > 0$, we have $u(x,y) > 0$ whenever $x \leq x_* - \epsilon_o$, $y \in [x, x + \epsilon_1]$. Hence, $(x - x_*)(y - x_*) \geq 0, |x - x_*| > |y - x_*|$ and $\epsilon_1 \geq |y - x|$ imply $u(x,y) > 0$ for $x \leq x_* - \epsilon_o$. $\square$

Assume u_ϵ is regular, let $l_\epsilon = [0,1] \cap [x_* - \epsilon, x_* + \epsilon]$ and, for $x \in l_\epsilon$, define

$$\phi_\epsilon(x) = \arg\max_{y \in N_\epsilon} u(x,y)$$

Note that if u is regular, then ϕ_1 is party 2's best response function and, since u_ϵ is regular, ϕ_ϵ is well-defined. Moreover, $(x, \phi_\epsilon(x)), (\phi_\epsilon(x), x) \in N_\epsilon$ whenever $x \in l_\epsilon$. Hence, with some abuse of terminology, we will call ϕ_ϵ the best-response function whenever u_ϵ is regular.

We will construct a special class of trajectories to use in the proof of Theorem 2-5. We call these, rush-to-the-best-response (RBR) trajectories. In an RBR trajectory both parties try to reach their respective best response functions ϕ_ϵ as quickly as possible, assuming their rival will do the same. Once the first of the two parties reaches its best response, it continues at a pace that keeps it on the best response conditional on the rival continuing to move, as fast as possible, towards its own best response. The rival reaches ϕ_ϵ at the local equilibrium after which both parties stay put forever.

Formally: let $\Omega_0 = \{\omega_*\}$ where $\omega_* = (x_*, x_*)$ is the local equilibrium. Let Ω_i^ϵ denote the graph of the party i's best response function, excluding ω_*. That is, $\Omega_1^\epsilon := \{\omega \in \Omega : \omega_1 = \phi_\epsilon(\omega_2)\} \backslash \Omega_0$ and $\Omega_2^\epsilon = \{\omega \in \Omega : \omega_2 = \phi_\epsilon(\omega_1)\} \backslash \Omega_0$. Finally, $\Omega_3^\epsilon := \Omega \backslash (\Omega_0 \cup \Omega_1^\epsilon \cup \Omega_2^\epsilon)$. The four sets, $\Omega_0^\epsilon, \Omega_1^\epsilon, \Omega_2^\epsilon, \Omega_3^\epsilon$, form a partition of Ω. Then, the RBR trajectories are defined as follows:

For $i, j = 1, 2$, $j \neq i$ and $t \geq 0$, let

$$\upsilon_i^\epsilon(o) = \begin{cases} 1 & \text{if } o_i < \phi_\epsilon(o_j) \\ 0 & \text{if } o_i = \phi_\epsilon(o_j) \\ -1 & \text{if } o_i > \phi_\epsilon(o_j) \end{cases}$$

We let $\min \emptyset = \infty$ and define,

$$\begin{aligned} Y_i^\epsilon(o, t) &= o_i + \upsilon_i^\epsilon(o) \cdot t \\ Y^\epsilon(o, t) &= (Y_1^\epsilon(o, t), Y_2^\epsilon(o, t)) \\ \tau_i^\epsilon(o) &= \min\{t : Y_i^\epsilon(o, t) = x_*\} \\ \tau^\epsilon(o) &= \min\{t : Y^\epsilon(o, t) \notin \Omega_3^\epsilon\} \end{aligned}$$

Then, define $Z^\epsilon : \Omega \times I\!R_+ \to \Omega$ (where $Z := Z^1$) as follows:

For $o \in \Omega_0^\epsilon$, $Z^\epsilon(o, t) = o$ for all t.

For $o \in \Omega_i^\epsilon$ and $j = 3 - i$, $Z_j^\epsilon(o, t) = Y_j^\epsilon(o, t)$ and $Z_i^\epsilon(o, t) = \phi_\epsilon(Z_j^\epsilon(o, t))$ for $t \leq \tau_i^\epsilon(o)$; $Z^\epsilon(o, t) = (x_*, x_*)$ for $t > \tau_i^\epsilon(o)$.

For $o \in \Omega_3^\epsilon$, $Z^\epsilon(o, t) = Y^\epsilon(o, t)$ for $t \leq \tau^\epsilon(o)$. Then, if $Z^\epsilon(o, \tau^\epsilon(o)) \in \Omega_k^\epsilon$ for $k = 0, 1, 2$, let $Z^\epsilon(o, t) = Z^\epsilon(Z^\epsilon(o, \tau^\epsilon(o)), t - \tau^\epsilon(o))$ for all $t > \tau^\epsilon(o)$; otherwise, $Z^\epsilon(o, t) = Z^\epsilon(o, \tau^\epsilon(o))$ for all $t > \tau^\epsilon(o)$.

Lemma 4: *Assume u_ϵ is regular, $o \in N_\epsilon$, and let θ^1, θ^2 be trajectories such that $\theta^1(0) = \theta^2(0) = o$ and $\theta_i^i(t) = Z_i^\epsilon(o, t)$ and $\theta(t) \in N_\epsilon$ for all $t \in [0, T]$. Then:*

(i) $Z^\epsilon(o, t) \in N_\epsilon$ for all t and $Z^\epsilon(o, t) = \omega_*$ for all $t \geq \|o - \omega_*\|$.

(ii)

$$\int_0^T e^{-rt} u(\theta^1(t))dt \leq \int_0^T e^{-rt} u(Z^\epsilon(o, t))dt \leq \int_0^T e^{-rt} u(\theta^2(t))dt \qquad (A1)$$

Moreover, the first inequality is strict if $Z^\epsilon(o, t) \neq \theta^1(t)$ for some $t \in [0, T]$ and the second inequality is strict if $Z^\epsilon(o, t) \neq \theta^2(t)$ for some $t \in [0, T]$.

Proof: Part (i): First, we show that $Z^\epsilon(o, t) = \omega_* = (x_*, x_*)$ for all $t \geq \|o - \omega_*\|$. If $o \in \Omega_0^\epsilon$, the Lemma is immediate. For $o \in \Omega_i^\epsilon$, at every t, party $3 - i$ is moving as fast as it can towards its best response function while party i is moving with just the right speed to ensure that the two parties stay on the graph of party i's best response. This continues until party i reaches x_* at which point the parties are at the local equilibrium ω_*. Therefore, they reach ω_* no later than $|o_i - x_*| \leq \|o - \omega_*\|$. If $o \in \Omega_3$, then both parties are moving towards their best response functions. When the first party reaches its best response, we are in the case considered above. Hence, the total time until the parties reach ω_* is no greater than $\|o - \omega_*\|$.

Call a party that reaches its best response at least as early as its opponent, an advantaged party and its opponent a disadvantaged party. Without loss of generality, assume that party 2 is an advantaged party and that party 1 is a disadvantaged party. Note that party 1 always moves as fast as possible towards its best response, reaches it at x_*, no sooner than party 2, and starts off within ϵ of x_*. Since both parties stay at x_* once they reach it, it follows $Z(o,t) = (x_*, x_*)$ for all $t \geq \epsilon$.

Next, we show that $Z^\epsilon(o,t) \in N_\epsilon$ for all $t \leq \epsilon$. We noted above, that party 1, a disadvantaged party, is always is moving towards x_*. Thus, $Z_1^\epsilon(o,\cdot)$ either monotonically increases or monotonically decreases to x_* and, therefore, $|Z_1^\epsilon(o,\cdot) - x_*| \leq \epsilon$ for all t. To complete the argument, we must show that party 2, the advantaged party, always stays within ϵ of x_* as well. But this is immediate: since party 2 reaches x_* by $t = \epsilon$, it must be within ϵ of x_* at $t \leq \epsilon$.

Part (ii): Since u_ϵ is regular, the best responses are best responses to all $x \in [x_* - \epsilon, x_* + \epsilon]$. Then, regularity ensures that for all $i \neq j$, $\omega, \hat{\omega} \in N_\epsilon$ such that $\omega_j = \hat{\omega}_j$, $(\omega_i - x_*)(\hat{\omega}_i - x_*) \geq 0$,

$$|\omega_i - \phi_\epsilon(\omega_j)| \leq |\hat{\omega}_i - \phi_\epsilon(\omega_j)| \text{ implies } u(\omega) \geq u(\hat{\omega})$$

and the second inequality is strict whenever the first one is strict. It follows that the second inequality in (A1) holds and is strict whenever $\theta^2(t) \neq Z(o,t)$ for some $t \leq T$. For the first inequality, note that, if $\theta_2(t) \neq x_*$ for all $t \leq T$, the argument of the previous paragraph ensures that $u(\theta^1(t)) < u(Z(o,t))$ whenever $\theta^1(t) \neq Z(o,t)$. If $\theta_2(t) = x_*$ for some $t' \leq T$, let τ be the smallest such t' and let t^* be the first time t at which party 2 reaches its best response; that is, the first t such that $Z_2(o,t) = \phi_\epsilon(Z_1(o,t))$. The argument of the previous paragraph ensures that $u(\theta(t)) \leq u(Z(o,t)$ for all $t \leq \tau$ and that this inequality is strict at t^* whenever $\theta_2^1(t) \neq Z_2(o,t)$ for some $t \leq t^*$. For $t > t^*$, recall that $Z_2(o,t)$

is the best response to $Z_1(o,t)$, and therefore, $u(\theta(t)) \leq u(Z(o,t)$ for all $t > t^*$ and this inequality is strict whenever $\theta_2^2(t) \neq Z_2(o,t)$. This proves that the first inequality holds and holds strictly if $\theta_2^2(t) \neq Z_2(o,t)$ at some $t \leq T$. □

Proof of Theorem 2: Define $d_i(\omega)$ and $d_i(\omega, f)$ such that $d_i(\omega, f)(t) = d_i(\omega)(t) = Z_i(\omega, t)$ for all $t \in [0, \lambda]$ and all $\omega \in \Omega, f \in H$. By Lemma 3, $d_i(\omega) \in H$ and hence $d = (d_1, d_2)$ is a stationary strategy. We claim that d is the unique equilibrium of (u, r, o, λ). Let $d^* = (d_1^*, d_2^*)$ be a stationary equilibrium, the existence of which is ensured by Theorem 1A.

Since u is regular, so is u_ϵ for $\epsilon = 1$. For any σ_1, σ_2, let $\sigma^1 = (\sigma_1, d_2)$ and $\sigma^2 = (d_1, \sigma_2)$. Then, part (ii) of Lemma 4 implies that

$$U(\theta_{o\sigma^2}) \leq U(\theta_{od}) \leq U(\theta_{o\sigma^1}) \tag{A2}$$

and

$$U(\theta_{o\sigma^2}) < U(\theta_{od}) < U(\theta_{o\sigma^1}) \tag{A3}$$

whenever $\theta_{o\sigma^1} \neq \theta_{od} \neq \theta_{o\sigma^2}$. Applying (A2) to $\sigma^1 = (d_1^*, d_2)$ and $\sigma^2 = (d_1, d_2^*)$ reveals that $U(\theta_{od^*}) = U(Z(o, \cdot))$ for all o and hence, d^* is an equilibrium of (u, r, o, λ) for all $o \in \Omega$. Then, (A3) implies d must be the unique equilibrium of (u, r, o, λ). Part (i) of Lemma 4 implies that $Z(o,t) = \omega_*$ for all $t \geq 1$ completing the proof of the claim.

Note that $Z(o, \cdot)$ is independent of λ and hence it is the real-time trajectory of (u, r, o) by the claim above. Then, part (i) of Lemma 4 ensures that (x_*, x_*) is the strong steady-state. □

Next, we define the stationary strategy $d_i^n, i = 1, 2$, referred to as the *i's pursuit strategy* below. With the pursuit strategy, i moves towards the policy j selected at the end of the last period; that is, at the end of the last tight history, whenever the distance between the end-of-period policy positions is greater than $(n-1)\lambda$ and moves toward x_* when this distance is no greater than $(n-1)\lambda$. Formally, define $h_i^\omega : [0, \lambda] \to [0, 1]$ as follows:

$$h_i^\omega(t) = \begin{cases} \min\{\omega_i + t, x_*\} & \text{if } \omega_i \leq x_* \\ \max\{\omega_i - t, x_*\} & \text{if } \omega_i > x_* \end{cases}$$

For $j = 3 - i$, and $n \geq 3$, define $d_i^n(\omega) \in H(\omega_i)$ as follows

$$d_i^n(\omega)(t) = \begin{cases} \omega_i + t & \text{if } \omega_i < \omega_j - (n-1)\lambda \\ h_i^\omega(t) & \text{if } |\omega_j - \omega_i| \leq (n-1)\lambda \\ \omega_i - t & \text{if } \omega_i > \omega_j + (n-1)\lambda \end{cases}$$

Let $d_i^n(\omega, f) = d_i^n(\omega)$ for all (ω, f) such that $\omega_j = f(\lambda)$. Let σ_j^n the dynamic discrete game strategy associated with d_i^n; that is, $\sigma_i^n(p) = d_i^n(\bar{\omega}(p))$ and $\sigma_i^n(p, f) = d_i^n(\omega(p), f)$ for all histories $p \in P_i$ and $(p, f) \in Q_i$.

Lemma 5: *Assume $n > 2$ and $\lambda < 1/3$. Let $\sigma_j = \sigma_j^n$, $j \neq i$, $\theta = \theta_{o\sigma}$ and $t \geq 1$, then one of the following two conditions must hold:*

$$x_* \leq \theta_j(t) \leq \theta_i(t) \leq \theta_j(t) + (n+1)\lambda \qquad (a)$$

$$\theta_j(t) - (n+1)\lambda \leq \theta_i(t) \leq \theta_j(t) \leq x_* \qquad (b)$$

Moreover, if $\theta(t)$ satisfies (a) or (b), then $\theta(t')$ satisfies (a) or (b) for all $t' \geq t$.

Proof: Suppose (a) holds at $t = k\lambda$ for some $k = 0, 1, \ldots$. Then, verifying that either (a) holds or (b) holds for all $t = (k+1)\lambda$ is straightforward. Hence, by induction we have: if (a) or (b) hold for some $t = k\lambda$, then (a) or (b) hold at $t' = k'\lambda$ for $t' \geq t$. Next, we will show that if (a) or (b) holds at $t = k\lambda$, then (a) or (b) hold must also hold at $k\lambda + \epsilon$ for all $\epsilon \in (0, \lambda)$. By symmetry, we can assume, without loss of generality, that (a) holds at $t = k\lambda$. If $\theta_j(k\lambda) < \theta_i(k\lambda) + (n-1)\lambda$, the desired conclusion is obvious. Otherwise, either $x_* \geq \theta_j(k\lambda) - \lambda$, in which case $\theta_j(k\lambda + \epsilon) = x_*$ and the desired conclusion follows, or $x_* < \theta_j(k\lambda) - \lambda$, in which case $\theta_j(k\lambda + \epsilon) = \theta_j(k\lambda) - \lambda\epsilon$ and, again the desired conclusion follows.

So, to conclude the assertion, it is enough to show that if (a) or (b) holds at $t = k\lambda + \epsilon$ for $\epsilon \in (0, \lambda)$, then (a) or (b) must also hold at either $k\lambda$ or $(k+1)\lambda$. Again, we assume without loss of generality, that (a) holds at t. If $\theta_j(k\lambda) = \theta_j(k\lambda + \epsilon)$, then $\theta_j(k\lambda) = x_* = \theta_j(k\lambda + \epsilon)$ and hence (a) or (b) must hold at $k\lambda$. If $\theta_j(k\lambda) < \theta_j(k\lambda + \epsilon)$, then either $\theta_j((k+1)\lambda) = x_*$ and again, we are done or $\theta_j(k\lambda) < \theta_i(k\lambda) + (n-1)k$ and hence $t = (k+1)\lambda$ satisfies (a). Finally, if $\theta_j(k\lambda) > \theta_j(k\lambda + \epsilon)$, then either $\theta_j((k+1)\lambda) = x_*$ and again, we are done or (a) still holds at $t = (k+1)\lambda$.

To conclude the proof of the lemma, we will show that there is some $t \leq 1$ at which either (a) or (b) holds. We will consider only the case in which $o_j \leq o_i$. Let m be the first k such that $\theta_j(k\lambda) \geq \theta_i(k\lambda) - (n-1)\lambda$. Clearly, $m\lambda \leq 1$. If $x_* \leq \theta_j(m\lambda)$, then (a) holds at $m\lambda$ and we are done. If $\theta_j(k\lambda) < x_* \leq \theta_i(k\lambda)$, let m^* be the smallest integer such that $x_* \leq m^*\lambda$. Clearly, $m^* > m$ and we must have $\theta_j(m^*\lambda) = x_*$ and hence at $m^*\lambda$, (a) or (b) holds; or at some k such that $m < k < m^*$, (b) holds. □

Proof of Theorem 3: Choose $\epsilon > 0$ so that u_ϵ is regular and (ii) of Lemma 3 holds. Without loss of generality, let $\lambda < {}^{\epsilon}\!/_{6}$. Then, choose n, an integer multiple of 3, such that $n \leq {}^{\epsilon}\!/_{\lambda} - 1$. Define,

$$\epsilon_1 = (n-1)\lambda$$

$$m = n/3$$

$$\epsilon_2 = m\lambda$$

Note that $0 < 2\epsilon_2 \leq \epsilon_1 \leq \epsilon$ and $(n+1)\lambda \leq \epsilon$.

Let $o \in N_{\epsilon_2}$ be the initial state and note that since $\epsilon_2 \leq \epsilon$, (ii) of Lemma 3 holds. Define the following (non-stationary) strategy σ_i. This strategy has two phases.

Phase 1: The first $m-1$ periods comprise phase 1 and in that phase σ_i follows $Z_i(o, \cdot)$ irrespective of j's action. That is: for $p \in P_i^k$ and for $(p, f) \in Q_i^k$ such that $0 \leq k < m$ and $t \in [k\lambda, (k+1)\lambda]$

$$\sigma_i(p)(t) = \sigma_i(p, f)(t) = Z_i^\epsilon(o, t)$$

Phase 2: Phase 2 consists of all periods m and later. In phase 2, player i plays her pursuit strategy; that is, for $p \in P_i^k$ and for $(p, f) \in Q_i^k$ such that $k \geq m$

$$\sigma_i(p) = \sigma_i(p, f) = d_i^n(\bar{\omega}(p)).$$

Since $Z^\epsilon(o, t) = Z^{\epsilon_2}(o, t)$ for all $o \in N_{\epsilon_2} \subset N_\epsilon$, part (i) of Lemma 4 implies that $Z^\epsilon(o, t) = (x_*, x_*)$ for all $t \geq \epsilon_2$. Let $\sigma = (\sigma_1^*, \sigma_2^*)$ be an equilibrium of (u, r, o, λ). To complete the proof of Theorem 3, we will show that $\theta_{o\sigma^*}(t) = Z^\epsilon(o, t)$ for all t.

Let $\hat{\sigma} = (\sigma_1^*, \sigma_2)$, let $\sigma = (\sigma_1, \sigma_2)$ and let $\tilde{\sigma} = (\sigma_1, \sigma_2^*)$. Since $o \in N_{\epsilon_2}$ and $m\lambda = \epsilon_2 \leq \epsilon_1/2$, $\theta_{o\sigma'}(t) \in N_{\epsilon_1}$ for all $t \leq m\lambda$ and all $\sigma' \in \Sigma$. Therefore, part (ii) of Lemma 4 implies

$$\int_0^{m\lambda} e^{-rt} u(\theta_{o\tilde{\sigma}}(t)) \leq \int_0^{m\lambda} e^{-rt} u(\theta_{o\sigma^*}(t)) \leq \int_0^{m\lambda} e^{-rt} u(\theta_{o\hat{\sigma}}(t))$$

Let $(x, y) = \theta_{o\hat{\sigma}}(m\lambda))$. First, note that $y = x_*$ and since $(x, y) \in N_{\epsilon_1} = N_{(n-1)\lambda}$, (x, y) satisfies (a) or (b) in Lemma 5. Then, Lemmas 3 and 5 imply $u(\theta_{o\hat{\sigma}}(t)) \geq 0$ for $t \geq m\lambda$. An analogous argument for player 1 implies that $u(\theta_{o\tilde{\sigma}}(t)) \leq 0$ for $t \geq m\lambda$. Since σ^* is an equilibrium, it follows that

$$\int_0^\infty e^{-rt} u(\theta_{o\sigma^*}(t)) = \int_0^\infty e^{-rt} u(\theta_{o\sigma}(t))$$

Since $\theta_{o\sigma} = Z^{\epsilon_2}(o, \cdot)$, Lemma 4 and the above equality imply that $\theta_{o\sigma^*} = Z^{\epsilon_2}(o, \cdot)$, as desired. □

10. Proofs of Theorems 4-6

For the case with infrequent elections we must modify the definition of a stationary strategy. Let $\bar{m} = 1/\lambda$. A stationary strategy depends on the state $\omega \in \Omega$, the number of periods since the last election $m \in N = \{1, \ldots, \bar{m}\}$, and (if $p \in Q_i^k$) on the $f \in H$ chosen by the opponent. Let $j = 3 - i$ and, for $i = 1, 2$, let $M_i = \Omega \cup \{(\omega, f) \in \Omega \times H \mid f(0) = \omega_j\}$ be as above. Let

$$\mathcal{D}_i := \{d_i : M_i \times N \to H : d_i(\omega, f, m) \in H(\omega_i), d_i(\omega, m) \in H(\omega_i)\}$$

A strategy σ_i for player i is stationary if there is a $d_i \in \mathcal{D}_i$ such that for all $k = 1, \ldots,$

$$\sigma_i(p) = d_i(\bar{\omega}(p), m) \text{ for } p \in P_i^k, m \equiv k \bmod m$$

and

$$\sigma_i(p, f) = d_i(\bar{\omega}(p), m, f) \text{ for } p \in Q_i^k, m \equiv k \bmod m$$

Theorem 2A: *A discrete-time game with infrequent elections has a stationary equilibrium and a unique equilibrium payoff.*

Proof: Let $\mathcal{V}$ be the set of all bounded continuous functions on $\Omega \times N$ where $N = \{1, \ldots, \bar{m}\}$. This set, endowed with the sup norm, is a complete metric space. Let

$$J_V(f_1, f_2, m) = \hat{u}(f_1(0), f_2(0), m) + e^{-r\lambda} V(f_1(\lambda), f_2(\lambda), \tau(m))$$

such that $\hat{u}(\cdot,1) := u(\cdot)$ and $\hat{u}(\cdot,m) := 0$ for $m > 1$; $\tau(m) = m+1$ for $1 \leq i \leq m-1$ and $\tau(\bar{m}) = 1$. Next, define

$$\begin{aligned} Q_V^1(f_1, y, m) &= \max_{f_2 \in H(y)} J_V(f_1, f_2, m) \\ Q_V^2(f_2, x, m) &= \min_{f_1 \in H(x)} J_V(f_1, f_2, m) \\ \Lambda_V^1(\omega, m) &= \min_{f_1 \in H(\omega_1)} Q_V^1(f_1, \omega_2, m) \\ \Lambda_V^2(\omega, m) &= \max_{f_2 \in H(\omega_2)} Q_V^2(f_2, \omega_1, m) \\ \Lambda_V &= \Lambda^1_{\Lambda_V^2} \end{aligned}$$

By the same argument as the one given in the proof of Theorem 1A, $\Lambda_V^2(\omega, m)$ is well-defined and that the function Λ_V^2 is in $\mathcal{V}$. Moreover, there is a unique $V \in \mathcal{V}$ such that $\Lambda_V = V$. In the following, let $V_1 = V$ and $V_2 = \Lambda_{V_1}^2$. The remainder of the proof mirrors the proof of Theorem 1A above. The only difference is the inclusion of the state variable $m \in N$.

Choose D_2 such that

$$D_2(f, \omega, m, j) \in \begin{cases} \arg \max_{f_2 \in H(\omega_2)} Q_{V_1}^2(f_2, \omega_1, m) & \text{if } j = 2 \\ \arg \max_{f_2 \in H(\omega_2)} J_{V_2}(f, f_2, m) & \text{if } j = 1. \end{cases}$$

Similarly, choose D_1 such that

$$D_1(f, \omega, m, j) \in \begin{cases} \arg \min_{f_1 \in H(\omega_1)} Q_{V_2}^1(f_1, \omega_2, m) & \text{if } j = 1 \\ \arg \min_{f_1 \in H(\omega_1)} J_{V_1}(f_1, f, m) & \text{if } j = 2. \end{cases}$$

Next, we define the dynamic decision problem $\alpha = (\delta, b, w, X, T, S, \delta)$ (for party 1) as follows: $\delta = e^{-r\lambda}$, $S = \Omega \times \{1, \ldots, n\} \times \{1, 2\}$, $X = H$, $b(\omega, f, m, j) = H(\omega_2)$ and

$$\begin{aligned} w(f, \omega, m, j) &= \hat{u}(\omega, m) \\ T(f, \omega, m, j) &= (\bar{\omega}(f, D_2(f, \omega, m, j)), \tau(m), 3 - j) \end{aligned}$$

It is easy to verify that $\hat{W} : S \to I\!R$ such that

$$\hat{W}(\omega, m, j) = \begin{cases} V_1 & \text{if } j = 1 \\ V_2 & \text{if } j = 2 \end{cases}$$

is a recursive value for α. Then, part (ii) of the Optimality Lemma implies that $\hat{W}$ is the value for α. Define ρ_1 such that

$$\rho_1(\omega, m, j) \in \begin{cases} \arg\min_{f_1 \in H(\omega_1)} Q^1_{V_2}(f_1, \omega_2, m) & \text{if } j = 1 \\ \arg\min_{f_1 \in H(\omega_1)} J_{V_1}(f_1, D_2(f_1, \omega, m, j), m) & \text{if } j = 2. \end{cases}$$

By construction ρ_1 is unimprovable for the recursive value $\hat{W}$ and hence, part (ii) of the Optimality Lemma implies that ρ_1 is an optimal policy. Define the stationary strategy $d_i, i = 1, 2$ as follows: $d_i(\omega, m) = D_i(f, \omega, m, i)$ (and note that D_i is independent of the first argument in state (ω, m, i)) and $d_i(\omega, m, f) = D_i(f, \bar{\omega}(p), m, 3 - i)$. Since ρ_1 is an optimal policy in the decision problem α, it is immediate that d_1 is a best response to d_2. A symmetric argument ensures that d_2 is a best response to d_1 and hence (d_1, d_2) is an equilibrium. Uniqueness of the equilibrium payoff follows from an argument identical to the one given in the proof of Theorem 1A above. □

Proof of Theorem 4: Let ϵ, λ and n be as in the proof of Lemma 5. Assume player 2 chooses the pursuit strategy d_2^n as defined prior to the proof of Lemma 5. (The pursuit strategy is independent of m, the number of periods that have elapsed since the last election.) Then, for any strategy of the opponent, (a) or (b) of Lemma 5 must hold for all $t \geq 1$. Lemma 3 then implies that along any trajectory θ consistent with player 2 choosing d_2^n we have $u(\theta(\ell)) \geq 0$ for all $\ell = 1, 2, \ldots$. Therefore, player 2's equilibrium payoff is non-negative and an analogous argument for player 1 implies that equilibrium payoffs of both players must be zero.

We claim that the strategy profile $\hat{\sigma} = (\hat{\sigma}_1, \hat{\sigma}_2)$, defined below, is an equilibrium. Let (d_1, d_2) be a stationary equilibrium strategy, the existence of which is guaranteed by Theorem 2A. In $\hat{\sigma}$, both players follow their pursuit strategies as long as no one deviates; after any deviations both players revert to the stationary strategies (d_1, d_2).

To see that $\hat{\sigma}$ is an equilibrium, we need only verify that no party benefits by being the first one to deviate. Suppose party 1 were to deviate first in period k. Then, the argument of the first paragraph establishes that by continuing with its pursuit strategy, party 2 would guarantee a payoff of at least 0. Hence, party 2's payoff following the stationary equilibrium strategy must be greater or equal to 0. This implies that party 1 cannot benefit from the deviation. Reversing the roles of party 1 and party 2 establishes that $\hat{\sigma}$ is an equilibrium. To conclude the proof of the theorem, we note that when both parties employ their pursuit strategies one of them is always moving toward x_* as fast as possible; any party that is not always moving toward x_* as fast as possible, reaches x_* before its opponent. Therefore, ω_* is reached at time 1 and both parties remain at that state thereafter. □

Proof of Theorem 5: Choose $\epsilon > 0$ small enough so that condition (ii) of Lemma 3 holds and let $\lambda < \epsilon/2$. Recall that

$$u(x, y) = -v(x)\Gamma(\Delta) + v(y)(1 - \Gamma(\Delta)) + \beta(1 - 2\Gamma(\Delta))$$

where $\Delta = (y - x)/2$. Furthermore, observe that Γ is strictly decreasing and $\Gamma(0) = 1/2$. Therefore, we may choose $\bar{\beta}$ such that for $\beta > \bar{\beta}$, $u(x, y) = 0$ implies that $|x - y| < \epsilon$. Then, condition (ii) of Lemma 3 implies that $u(x, y) = 0$ if and only if $x = y$. Hence, we conclude that if $(x, y) = \theta(\ell)$ for some $\ell = 1, 2, \ldots$, then $x = y$.

In proof Theorem 4, we established that the equilibrium payoff after any history is 0. Then, (1) the best payoff that any party can achieve in any election must be 0 and, by the argument of the previous paragraph, (2) $\theta(\ell) = (x, x)$ for some x. Assume that $x \neq x_*$ for some ℓ. Without loss of generality, assume $\ell = 1$ and $x < x_*$. For $1 \leq k \leq \bar{m}$, define $x_k = x - k\lambda$. We claim that $\theta_i((n - k)\lambda) = x_k$.

First, we prove that $\theta_i((\bar{m}-1)\lambda) = x_1$; if not, then, at the end of period $\bar{m}-1$, the last period before the election, $\theta_i((\bar{m} - 1)\lambda) \in (x - \lambda, x + \lambda]$ for some $i = 1, 2$; without loss of generality assume $i = 2$. If party 2 is the second mover in period $\bar{m} - 1$ then it can deviate so that the new trajectory $\hat{\theta}_2$ satisfies $\hat{\theta}_2(1) = y \in (x, x_*), |y - x| \leq \epsilon/2$ yielding a payoff $u(x, y) > 0$. This contradicts (1) above. If party 2 is the first mover in period $\bar{m} - 1$, then the preceding argument establishes that $\theta_1((\bar{m} - 1)\lambda) = x_1$. Again, party 2 can deviate so that the new trajectory is $\hat{\theta}_2(1) \in (x, x_*)$ within $\epsilon/2$ of x. Since $\theta_1((\bar{m} - 1)\lambda) = x_1$, any feasible trajectory for player 1 satisfies $x - \epsilon/2 \leq \hat{\theta}_1(1) \leq x$ and, therefore, $u(\hat{\theta}(1)) > 0$, contradicting (1) above. We have shown that (i) $\theta_i((\bar{m} - 1)\lambda) = x_1$ for $i = 1, 2$ and (ii) party $j = 3 - i$ can guarantee itself a strictly positive payoff at any $\hat{\theta}((\bar{m} - 1)\lambda)$ such that $\hat{\theta}_i((\bar{m} - 1)\lambda) = x_1, \hat{\theta}_j((\bar{m} - 1)\lambda) \in (x_1, x_1 + \lambda)$.

Replacing x with x_1 above, the same argument establishes that (i) $\theta_i((\bar{m}-2)\lambda) = x_2$ and (ii) party $j = 3 - i$ can guarantee itself a strictly positive payoff at $\hat{\theta}((\bar{m}-2)\lambda)$ such that $\hat{\theta}_i((\bar{m}-2)\lambda) = x_2, \hat{\theta}_{3-i}((\bar{m}-1)\lambda) \in (x_2, x_2+\lambda)$. Successively replacing x_k with x_{k+1} and repeating this argument then implies that $\theta_i(0) = x_{\bar{m}} = x - 1$; that is, $x = 1$, contradicting the fact that $x < x_* \leq 1$. □

Proof of Theorem 6: Below, we specify a strategy, σ_2, for player 2 that guarantees the desired bound. Let $K(\ell, \lambda, \gamma)$ be the largest integer k such that $k < [\ell + 1 - \gamma]/\lambda$. Thus, period $K(\ell, \lambda, 0) = \bar{m}(\ell+1) - 1$ is the last chance that the parties have for adjusting their policies before the $\ell + 1$'th election and period $K(\ell, \lambda, 1) + 1 = \bar{m}\ell$ is the first chance that the parties have for adjusting their policies after the ℓ'th election. To describe party 2's behavior, consider k such that $K(\ell, \lambda, 1) + 1 \leq k \leq K(\ell, \lambda, 0)$ and $\eta \in P_2^k \cup Q_2^k$. Then, define σ_2 as follows:

For $k \leq K(\ell, \lambda, \lambda + \alpha/2)$, let

$$\sigma_2(\eta)(t) = \begin{cases} \max\{\bar{\omega}_2(\eta) - t, 1/2\} & \text{if } \bar{\omega}_2(\eta) \geq 1/2 \\ \min\{\bar{\omega}_2(\eta) + t, 1/2\} & \text{if } \bar{\omega}_2(\eta) < 1/2 \end{cases}$$

For $k > K(\ell, \lambda, \lambda + \alpha/2)$ and $(x, y) = \bar{\omega}(\eta)$, let $\hat{y}$ be a static best response to x. Then, let

$$\sigma_2(\eta)(t) = \begin{cases} \max\{y - t, \hat{y}\} & \text{if } y \geq \hat{y} \\ \min\{y + t, \hat{y}\} & \text{if } y < \hat{y} \end{cases}$$

It is straightforward to verify that this yields the following lower bound to player 2's payoff:

$$\min_{z \in [0,1]} \max_{y \in [0,1]} \min_{x \in B_\gamma(z)} u(x, y)$$

where $\gamma = 2\lambda + \alpha$. Since λ converges to 0 as the game converges to real time, the payoff bound follows. □

Proof that the lower bound in the Example is tight: In this section, we show that the lower bound in Theorem 6 is tight; that is, we find an example in which party 1 can guarantee a real-time payoff u_α. Let

$$\begin{aligned} v(x) &= x \\ \beta &= 1/2 \\ \Gamma(\Delta) &= \frac{1}{2} + 4\Delta^3 \\ \alpha &= .05 \end{aligned} \tag{6}$$

Let $z \in [.5, 1]$ be the unique value that satisfies

$$u(z - .05, 0) = u(z + .05, 1)$$

and note that for the parameters above $z \approx .8436$. Let $\lambda_n = .05/n$ so that there are $\bar{m} = 20n$ periods between elections and $k \geq 1$ periods correspond to $0.05k$ units of time. The strategy we will construct for party 1 is as follows: as soon as an election takes place, party 1 moves as fast as possible to z; once it reaches z, party 1 stay put until the n period before the election. In each of the last n periods, party 1 moves as fast as possible towards $x = z - .05$ if the policy position of party 2 at the start of the current period is less than or equal 1/2; otherwise, it moves as fast as possible towards $x = z + .05$.

We describe the strategy for the first $20n$ periods. The strategy repeats after each election. Let $\eta \in P_1^k \cup Q_1^k$ and let $(x, y) = \bar{\omega}(\eta)$. Then, for $k \leq 19n$,

$$\sigma_1(\eta)(t) = \begin{cases} \max\{\bar{\omega}_1(\eta) - t, z\} & \text{if } \bar{\omega}_1(\eta) \geq z \\ \min\{\bar{\omega}_1(\eta) + t, z\} & \text{if } \bar{\omega}_1(\eta) < z \end{cases}$$

and for $19n \leq k \leq 20n - 1$ and $\eta \in P_1^k \cup Q_1^k$, let

$$\sigma_1(\eta)(t) = \begin{cases} \bar{\omega}_1(\eta) - t & \text{if } \bar{\omega}_2(\eta) \leq 1/2 \\ \bar{\omega}_1(\eta) + t & \text{if } \bar{\omega}_2(\eta) > 1/2 \end{cases}$$

Hence, with σ_1, party 1 "mirrors" party 2's movement around 1/2 with its own movement around z. It follows that if (x, y) is the state at election time, then x is within $1/n$ of $z + \alpha(y - 1/2)$. Therefore, party 1 can ensure a real-time payoff of $\max_{y \in [0,1]} u(z + \alpha(y - 1/2), y)$. Routine calculations reveal that since $\alpha = .1$, this maximum is attained at $y = 0$ or $y = 1$. Hence, the real-time payoff is at most $u(z + \alpha/2, 1)$. Similar calculations establish that $u_\alpha = \max_{y \in [0,1]} \min_{x \in B_\alpha(z)} u(x, 1) = \min_{x \in B_\alpha(z)} u(x, 1) = u(z + \alpha/2, 1)$ and hence, by Theorem 6, $u_* = u(z + \alpha/2, 1)$, proving that the bound of Theorem 6 is tight. □

References

Adams, J and Z. Somer-Topcu (2009): "Policy adjustment by parties in response to rival parties' policy shifts: spatial theory and the dynamics of party competition in twenty-five post-war democracies," *British Journal of Political Science,* 39, pp. 825–846.

Anderson, R. M. (1984): "Quick-response equilibrium." Working Papers in Economic Theory and Econometrics # IP-323, Center for Research in Management, University of California, Berkeley

Bergin, J, (2006): "The folk theorem revisited," *Economic Theory*, 27, pp. 321332.

Dalton, R. and I. McAllister (2014): "Random Walk or Planned Excursion? Continuity and Change in the Left-Right Positions of Political Parties," Comparative Political Studies, (December), pp. 1-29.

Duggan, J (2014): "A Survey of Equilibrium Analysis in Spatial Models of Elections," mimeo. University of Rochester.

Hinich, M., J. Ledyard, and P. Ordeshook (1972) "Nonvoting and the Existence of Equilibrium Under Majority Rule," *Journal of Economic Theory*, 4, 144153.

Hinich, M., J. Ledyard, and P. Ordeshook (1973) "A Theory of Electoral Equilibrium: A Spatial Analysis Based on the Theory of Games," *Journal of Politics*, 35, 154193.

Kramer, G., "A Dynamic Model of Political Equilibrium," *Journal of Economics Theory*, 16(2), pp. 310-334.

McDonald, M.D., Mendes, S.M. (2001): "The policy space of party manifestos." In: Laver, M. (Ed.) Estimating the policy position of political actors. London: Routledge.

Schofield, N., (2001): "Generic Existence of Local Political Equilibrium," in Approximation, Optimization and Mathematical Economics, M. Lassonde, ed. Springer: Heidelberg (March 2001), pp. 297-308.

Schofield, N. and I. Sened, (2002): "Local Nash Equilibrium in Multiparty Politics," Annals of Operations Research: 109 (August), pp. 193-210

Osborne, M. J. (1995): "Spatial Models of Political Competition under Plurality Rule: A Survey of Some Explanations of the Number of Candidates and the Positions They Take," *Canadian Journal of Economics*, 27, 261-301.