The Political Economy of Dynamic Elections: A Survey and Some New Results*

John Duggan† César Martinelli‡

May 23, 2017

Abstract

We survey and synthesize the political economy literature on dynamic elections in the two traditional settings, spatial preferences and rent-seeking, under perfect and imperfect monitoring of politicians' actions. We define the notion of stationary electoral equilibrium, which encompasses previous approaches to equilibrium in dynamic elections since the pioneering work of Barro (1973), Ferejohn (1986), and Banks and Sundaram (1998). We show that repeated elections mitigate the commitment problems of both politicians and voters, so that a *responsive democracy result* holds in a variety of circumstances; thus, elections can serve as mechanisms of accountability that successfully align the incentives of politicians with those of voters. In the presence of term limits, however, the possibilities for responsiveness are attenuated. We also touch on related applied work, and we point to areas for fruitful future research, including the connection between dynamic models of politics and dynamic models of the economy.

Keywords: dynamic elections, electoral accountability, median voter, political agency, responsiveness

*This paper is an extended version of "The Political Economy of Dynamic Elections: Accountability, Commitment, and Responsiveness," which is forthcoming at *Journal of Economic Literature.* The reader will find proofs and discussion omitted from the published article.

†Duggan: University of Rochester, dugg@ur.rochester.edu

‡Martinelli: George Mason University, cmarti33@gmu.edu

Contents

1 Introduction

By its very nature, representative democracy entails the delegation of power by society to elected officials who may use this power in ways that are not necessarily in agreement with the interests of the electorate. A main concern for representative democracy is then to devise means to discipline politicians in office to achieve desirable policy outcomes for citizens. Political thinkers since Madison, if not earlier, have considered the possibility of re-election to be an essential device in this regard.[1] An active and growing literature on electoral accountability has taken up this subject in the context of explicitly dynamic models. The ultimate goal of this literature is to improve our understanding of the operation of real-world political systems and the conditions under which democracies succeed or fail. This, in turn, may facilitate the design of political institutions that produce desirable sequences of policies. The literature is developing, but it has the potential to inform us about the interplay between politics and dynamic processes such as economic growth and cycles, the evolution of income inequality, and transitions to democracy (or in the opposite direction, to autocracy).

In this article, we survey and synthesize the literature on electoral accountability, focusing on the interplay between disciplining incentives, provided by the possibility of future re-election, and incentives for opportunistic behavior in the present. Drawing from this literature, we show that repeated elections can be effective in mitigating the commitment problem faced by politicians whose ideal policies are different from those desired by the majority. Moreover, we show that when office incentives are important enough and politicians and other citizens place sufficient weight on the future, *responsive democracy* is possible, in the sense that elected politicians choose policies that converge to the majority winning policy.

Although superficially similar to median voter results in the traditional Hotelling-Downs competition framework, the mechanism underlying responsive democracy is different: candidates cannot make binding campaign promises, and they do not compete for votes in the Hotelling-Downs sense; rather, they are citizen candidates whose policy choices must maximize their payoffs in equilibrium, and the responsiveness result is driven by competition with the prospect of outside challengers, who themselves are converging to the median. Both incentives and selection are important for this result: some politicians' short run incentives may be tempered by the desire to be re-elected, inducing them to compromise by choosing policies that are more desirable for voters; and politicians who are not willing to compromise will be removed, until a compromising candidate is elected. Though

[1] *The Federalist* 57, in particular, offers a discussion of the role of re-election in the selection of politicians and the control of politicians while in office.

we frame our discussion in terms of representative democracies, and consequently focus on elections as the means to discipline politicians, note that political accountability is to some extent at work in nondemocratic polities through protest, coups, and revolutions.

Convergence to the majority winner in repeated elections arises from a politician's concern for reputation and relies on the absence of term limits and the assumption of incomplete information. The desire to be re-elected may induce politicians to mimic types whose preferences are closer to those of the median voter, and if the reward for political office is large enough, then the desire for re-election induces politicians to approximate the median voter's ideal policy. Thus, repeated elections engender the possibility of responsive democracy, despite the paucity of instruments that the voters can yield relative to the principal-agent model in complete contract settings. We generally assume that politicians' preferences are private information, i.e., *adverse selection,* but we consider alternative assumptions about the observability of politicians' actions. In the *perfect monitoring* model, policy choices of politicians are observable, while in the model of *imperfect monitoring,* or *moral hazard,* policy choices are observed only with some noise. We do not attempt to explore each informational assumption under general specifications of preferences, but we will survey the most relevant specifications from the point of view of existing work on the topic.

Accordingly, throughout this review, we alternate the focus between two different environments that have received much attention in the literature. The first is the classical *spatial preferences* environment derived from Harold Hotelling (1929) and studied in the social choice tradition since the seminal work of Duncan Black (1948) and Anthony Downs (1957). In this environment, voters have conflicting policy preferences over a unidimensional policy space, and politicians have a short-run incentive to adopt their preferred policies rather than those favored by the median voter. As explained, above, this short-run incentive can be overcome in a repeated elections setting. The second environment is the *rent-seeking* environment studied in the public choice tradition exemplified by Robert Barro (1973) and John Ferejohn (1986). In this environment, politicians have a short-run incentive to shirk from effort, or equivalently to engage in rent-seeking activities that hurt other citizens, while in office. In a repeated elections setting, the incentive of re-election may induce politicians to exert high levels of effort as the office incentive becomes more important, overcoming short-run incentives to shirk even in the presence of adverse selection and moral hazard problems.

The spatial preferences and rent-seeking environments emphasize different conflicts of interest giving rise to short-term temptation—conflicts of interests between citizens or between the citizens at large and politicians in office—which capture important and related challenges to the well functioning of democracy. For instance, in the context of economic development, Acemoglu and coauthors (e.g., Acemoglu et al. 2005, Acemoglu and Robinson 2012) argue that nondemocratic institutions tend to serve an entrenched elite at the expenser of the citizens at large. As a consequence, these institutions suffer from a hold-up problem: they cannot commit to not expropriate wealth, so economic actors fail to make productive investments, with lower growth as a consequence. The authors claim that democratic political institutions can lead to more secure property rights and higher growth. This argument implicitly assumes, however, that political representatives in democratic systems can commit to the protection of property rights, but a premise of the electoral accountability approach is precisely that this is impossible. From the viewpoint of this literature, electoral democracy in itself does not prevent elected politicians from serving the interests of an elite because of the possibility of capture, and a central problem that arises is to understand the extent to which democratic institutions can indeed solve the *commitment problem of politicians.*

The electoral accountability literature shows that a key disciplining device for preventing politicians in office from serving themselves, an elite, or even the citizens' myopic interests is the existence of a viable opposition in the form of credible outside challengers. Electoral democracy in itself is not enough to solve the hold-up problem, but it can lead office holders to moderate their policy choices when politicians in office face the possibility of replacement. Although incumbents cannot commit now to moderate future policies, the anticipation of future challengers and the incentive to win re-election can serve to discipline politicians. In the absence of term limits, these incentives are maintained throughout an incumbent's tenure, and voters may rationally expect incumbents to choose moderate policies in the future.

The absence of a term limit is important for the possibility of responsive democracy. Elections can provide a commitment mechanism for politicians because an office holder must provide a majority of voters with an expected payoff at least equal to what they would obtain from an untried challenger; this is true at the time a politician decides whether to compromise her policy choice, and because voters know the politician will have the same incentives in the next period, they can rationally expect her to compromise in the future. When a term limit is in place, however, politicians always choose their ideal policies (or zero effort) in the last term of office, so prior to the last term, voters cannot expect an incumbent to compromise if re-elected, and the policy responsiveness result begins to unravel. But this logic is incomplete. Assuming for simplicity that a two-period term limit is in effect, it could still be that voters re-elect an incumbent after her first term of office if her policy choice (or effort level) passes some threshold, inducing the politician to compromise in her first term, even though she chooses her ideal policy in the second term. Now it is the *commitment problem of voters* at work: if first-term politicians were expected to compromise, then a majority of voters would strictly prefer to elect a challenger rather than re-elect an incumbent, so such a threshold cannot be supported in equilibrium.

Interestingly, this logic does not apply in a two-period model, because elected challengers are also expected to shirk, so the two-period model and the infinite-horizon model with a two-period term limit possess fundamentally different incentive properties. We show that a version of the responsive democracy result does in fact obtain in the two-period model, with policy choices in the first period reflecting the preferences of the median voter as politicians become more office motivated. Thus, somewhat paradoxically, the two-period model better approximates the infinite-horizon model with no term limit than it does the infinite-horizon model with term limits. Of course, the infinite-horizon model with term limits is not necessarily a realistic model for representative democracies, since politicians' careers usually extend beyond their term in office, so the idea that an incumbent will simply act in a completely self-serving fashion in the final term of office is perhaps extreme.

The remainder of this article is organized as follows. Section 2 overviews the classical static framework of electoral competition and provides notation and background results used throughout. Section 3 presents a basic two-period model of electoral accountability in the spatial preferences and in the rent-seeking environments, and it serves to introduce issues related to imperfect observability of preferences and policy choices in the sequel. Section 4 presents the infinite-horizon framework, encompassing much of the recent literature and introducing the concept of stationary electoral equilibrium in the dynamic model. Section 5 summarizes the literature dealing with adverse selection in infinite-horizon models. Section 6 summarizes the literature dealing with political moral hazard in infinite-horizon models. Section 7 reviews some of the applied literature connected to electoral accountability, and it shows how the dynamic electoral framework presented here can be used to model pandering and political business cycles, generalizing some existing work in these areas. Section 8 concludes by identifying areas for future research that are critical to the development of dynamic political economy as a field.

2 Classical electoral competition

In this section, we present a static electoral framework, review classical results in the theory of elections, and set notation and background results for the analysis of dynamic elections to follow.

2.1 Hotelling-Downs model

We begin with a basic model of electoral competition, tracing back to Hotelling (1929) and Downs (1957), that assumes the political actors are two parties and are *office-motivated,* in the sense that both parties seek to win election without regard to policy outcomes per se. The two parties simultaneously announce policy platforms; each voter casts a ballot for the party offering her preferred platform; and parties seek to maximize their chances of winning the election. We denote the policy space by X, and for simplicity we assume throughout that $X \subseteq \mathbb{R}$. A continuum N of voters is partitioned in a finite set $T = \{1,\dots,n\}$ of types, with $n \geqslant 2$, and each voter type $j \in T$ has policy preferences given by the utility function $u_j{:}X \to \mathbb{R}$. Assume:

(A1) For each $j \in T$, u_j has unique maximizer $\hat{x}_j \in X$, which is the *ideal policy* of the type j citizen, and furthermore types are indexed in order of their ideal policies, i.e.,

$$\hat{x}_1 < \hat{x}_2 < \cdots < \hat{x}_n. \tag{1}$$

(A2) For all $j \in T$ and all $x,y \in X$ with $x > y$, the utility difference $u_j(x) - u_j(y)$ is strictly increasing in j, i.e., preferences are *supermodular.*

These assumptions admit two simple formulations of utility that we rely on for special cases. A common specification is *quadratic utility,* in which case $u_j(x) = -(x-\hat{x}_j)^2 + K$, where K is a constant; this functional form determines ideal policy $\hat{x}_j$, and utility differences are $y^2 - x^2 + 2\hat{x}_j(x-y)$, which is strictly increasing in the ideal policy when $x > y$, fulfilling (A1) and (A2). Another is *exponential utility,* whereby $u_j(x) = -e^{x-\hat{x}_j} + x + K$, which determines ideal policy $\hat{x}_j$ and also satisfies (A1) and (A2).

The distribution of types in the electorate is given by $(q_1,\dots,q_n)$, where $q_j > 0$ is the fraction of type j voters. We assume the generic property that types cannot be divided into exactly equal parts, i.e., there is no $S \subseteq T$ such that $\sum_{j\in S} q_j = \frac{1}{2}$. This implies that there is a unique *median* type, which we denote $m \in T$, defined by the inequalities

$$\sum_{j:j<m} q_j < \frac{1}{2} \quad \text{and} \quad \sum_{j:m<j} q_j < \frac{1}{2}.$$

By (A2), voter preferences are order restricted, and a result of Rothstein (1991) implies that the median type m is pivotal in pairwise voting,[2] in the sense that a

[2] See also Gans and Smart (1996) for analysis of a single-crossing condition that is equivalent to Rothstein's order restriction.

majority of voters strictly prefer policy x to policy y if and only if $u_m(x) > u_m(y)$. In particular, the ideal policy $\hat{x}_m$ of the median voter type defeats all other policies in pairwise majority voting, i.e., it is the *Condorcet winner.*

The two parties, A and B, simultaneously announce platforms x_A and x_B; importantly, we assume that the winning party is bound to its election platform. Each voter casts her ballot for the party offering the preferred platform, and the probability that party A wins, which is denoted $P(x_A,x_B)$, therefore satisfies:[3]

$$P(x_A,x_B) \quad = \quad \begin{cases} 1 & \text{if } u_m(x_A) > u_m(x_B), \\ 0 & \text{if } u_m(x_A) < u_m(x_B), \\ \frac{1}{2} & \text{if } x_A = x_B. \end{cases}$$

We do not impose any restriction when the parties offer distinct platforms and the median type is indifferent. Consistent with the assumption of office motivation, we assume party A's payoffs are given by $P(x_A,x_B)$, and party B's payoffs are $1 - P(x_A,x_B)$. A *Nash equilibrium* (in pure strategies) is a pair (x_A^*,x_B^*) of policies such that neither party can increase its probability of winning by deviating unilaterally.

Next, we state the well-known median voter theorem establishing that under the above weak conditions, strategic incentives of office-motivated candidates lead to the adoption of the Condorcet winner, a phenomenon we refer to as *responsive democracy.*

Proposition 2.1 *Assume (A1) and (A2). In the unique Nash equilibrium of the Hotelling-Downs model, we have* $x_A^* = x_B^* = \hat{x}_m$.

An especially important application of the model with win-motivated parties is to the determination of tax rates and public good provision. Romer (1975) applies the median voter theorem to a model of lump sum transfers and linear taxes with Cobb-Douglas utilities. Roberts (1977) extends the analysis to more general voter preferences and establishes that the voter with median income is pivotal; this is true even when preferences over tax rates fail to be single-peaked, because it can be shown that voter preferences are nonetheless order restricted. Meltzer and Richard (1981) provide a model in which the assumptions of the latter paper are satisfied, and they examine the effect of varying the decisive voter (e.g., through a change in the franchise) and the relative productivity of the median voter.

2.2 Calvert-Wittman model

The basic model of elections is extended by Calvert (1985) and Wittman (1977, 1983) to model political actors as candidates with policy preferences. We add the

[3]Technically, we assume that voters of the median type split their votes to create a tie, which is decided by the toss of a fair coin.

following standard continuity and convexity assumption:

(A3) The policy space X is convex, and for all $j \in T$, u_j is continuous and strictly quasi-concave.

Viewing candidates as citizens, we let one candidate be type $a \in T$ and the other type $b \in T$, and we assume that the political candidates have opposed preferences, i.e., $\hat{x}_a < \hat{x}_m < \hat{x}_b$. Given platforms x_a and x_b, the payoffs of candidate a are now given by

$$P(x_a, x_b)(u_a(x_a) + \beta) + (1 - P(x_a, x_b))u_a(x_b),$$

where $\beta \geqslant 0$ is an office benefit term that captures all non-policy rewards to holding office,[4] and candidate b's payoffs are analogous. Because we allow politicians to care about both policy and holding office, politicians have *mixed motivations.*

The median voter theorem extends to the Calvert-Wittman model.

Proposition 2.2 *Assume (A1)–(A3). In the unique Nash equilibrium of the Calvert-Wittman model, we have* $x_a^* = x_b^* = \hat{x}_m$.

We see that the Downsian responsive democracy result generalizes even to the case in which candidates have policy agendas that differ from the median voter's; thus, static elections, in which candidates can make binding campaign promises, lead to centrally located policy outcomes.

2.3 Probabilistic voting

We have thus far assumed that political actors have full information about the preferences of voters. A variation on the classical model, referred to as models of "probabilistic voting," assumes that a parameter of the voters' preferences is unobserved by the candidates at the time platforms are chosen. These models differ with respect to the particular parameterization used (candidates may have unobserved valences, or voters may have unobserved ideal policies) and the nature of the distribution of the parameters; early work is due to Hinich (1977), Coughlin and Nitzan (1981), Lindbeck and Weibull (1993), and Roemer (1997).

A simple way of introducing uncertainty is to assume an aggregate preference shock $\omega \in \mathbb{R}$ to voter preferences that is unobserved by politicians. Let ω be distributed according to a continuous distribution F with full support. We strengthen (A3) to

(A4) For all $j \in T$, u_j is strictly concave,

[4]The convention in the literature is to mention the alternative terminology of "ego rents," which we have now done as well.

and we assume the shock is linear: the utility of the type j voter from policy x is $u_j(x) + \omega x$. If utilities are quadratic, then ω can be viewed as simply a parameter that shifts each type j voter's ideal policy by the amount $\omega/2$. Given distinct platforms x_a and x_b, voters are indifferent between the platforms with probability zero; thus, for almost all shocks ω, candidate a wins if and only if the set

$$\{j \in T : u_j(x_a) + \omega x_a > u_j(x_b) + \omega x_b\}$$

contains a majority of voter types. By our supermodularity assumption (A2), this occurs if and only if the median type prefers candidate a's platform, i.e., $u_m(x_a) + \omega x_a > u_m(x_b) + \omega x_b$.

Therefore, assuming $x_a < x_b$, candidate a wins if and only if

$$\omega \quad < \quad \frac{u_m(x_a) - u_m(x_b)}{x_b - x_a},$$

and the function

$$H(x_a, x_b) \quad \equiv \quad F\left(\frac{u_m(x_a) - u_m(x_b)}{x_b - x_a}\right)$$

gives the probability that candidate a wins. Then candidate a's payoff is

$$H(x_a, x_b)(u_a(x_a) + \beta) + (1 - H(x_a, x_b))\, u_a(x_b),$$

with candidate b's payoffs defined analogously.

Due to non-convexities of payoffs, discussed below, equilibrium may require mixed strategies on the part of candidates. Nevertheless, in the model with *pure policy motivation,* i.e., $\beta = 0$, Roemer (1997) establishes existence in pure strategies when the probability of winning is log concave. It is straightforward to show that, in contrast to the median voter theorem, candidates adopt distinct equilibrium platforms.

Proposition 2.3 *Assume (A1)–(A4). In the probabilistic voting model with pure policy motivation, assume that for all x_a and x_b with $x_a \leqslant x_b$, the functions $H(x_a, x_b)$ and $1 - H(x_a, x_b)$ are, respectively, log-concave in x_a and in x_b. Then there is a Nash equilibrium, and in every Nash equilibrium (x_a^*, x_b^*), we have $x_a^* < x_b^*$.*

The case of mixed motives, with $\beta > 0$, becomes complicated by the possibility that one candidate's best response may be to "jump over" the other in order to capture the office benefit β with higher probability. To extend the existence result to mixed motives and to provide an exact equilibrium characterization, we consider the *symmetric probabilistic voting model* as the special case such that $X \subseteq \mathbb{R}$ is an interval centered at zero; for all $j \in T$, u_j is quadratic with $\hat{x}_a = -\hat{x}_b$; the ex ante ideal policy of the median voter is zero, i.e., $\hat{x}_m = 0$; and for all x, $F(x) = 1 - F(-x)$. In this case, the probability that the perturbed ideal policy of the median voter is less than x is just the probability that $\omega \leqslant 2x$, which is just $F(2x)$. Note that (up to labeling of types), assumptions (A1)–(A4) are satisfied in this special case. The following is established by Bernhardt, Duggan, and Squintani (2009).

Proposition 2.4 *In the symmetric probabilistic voting model with mixed motivation, where (A1)–(A4) are satisfied, assume that u_a and u_b are differentiable and that F is log-concave. Then there is a unique symmetric Nash equilibrium, $(x^*, -x^*)$, and x^* is defined as follows: if $|u_a'(0)| \leqslant 2\beta f(0)$, then $x^* = 0$; and otherwise, x^* is the unique negative solution to*

$$-\frac{u_a'(x)}{u_a(x) + \beta - u_a(-x)} \quad = \quad 2f(0).$$

An implication is that increased office benefit leads candidates to adopt more moderate platforms. In fact, if candidates are sufficiently office motivated or the location of the median voter is known with high enough precision, i.e., $2\beta f(0) \geqslant |u_a'(0)|$, then we obtain exact coincidence of policy platforms, and analogous to the median voter theorem, the candidates both locate at the median of the distribution of medians in the unique equilibrium. Thus, an ex ante form of the responsive democracy result extends to the model with probabilistic voting and sufficiently office-motivated candidates.

The best response problem of a candidate with mixed motives is analogous to that of a first-term office holder in the moral hazard model covered in Subsection 3.4, so it is instructive to consider the non-convexity problem mentioned above and the role of log concavity in solving this problem. It is clear that because the candidate's objective function involves the term $H(x_a, x_b)u_a(x_a)$, it need not be quasi-concave. We can gain insight by transforming the problem into a constrained optimization problem in which the candidate chooses policy x and a winning probability p as follows:

$$\max_{(x,p)} p(u_a(x) - u_a(x_b) + \beta)$$
$$\text{s.t. } p \leqslant H(x, x_b),$$

where we omit the constant term $u_a(x_b)$ and (for expositional purposes) restrict the problem to $x \leqslant x_b$. The solutions to this problem correspond to the best policies of candidate a given x_b, subject to the restriction $x \leqslant x_b$. Although the objective function above is nicely behaved, the constraint set is not in general convex, and it

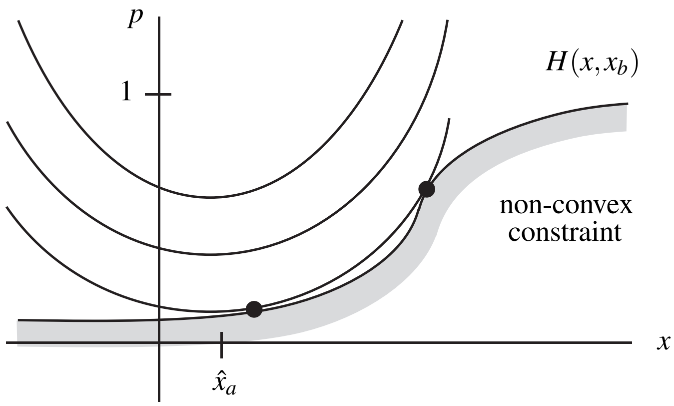

Figure 1: Multiple best responses

is possible in principle that the best response problem has multiple solutions; see Figure 1.

We can, however, translate the constrained optimization problem to log form as follows:

$$\max_{(x,p)} \ln(p) + \ln(u_a(x) - u_a(x_b) + \beta)$$
$$\text{s.t. } \ln(p) \leqslant \ln(H(x,x_b)).$$

The objective function of the transformed problem continues to be concave, and we assume $\ln(H(x,x_b))$ is concave in x, which implies that the constraint set is convex; see Figure 2. Thus, candidate a has a unique optimal policy subject to $x \leqslant x_b$, and when the politician is policy motivated, this policy will be globally optimal, obviating the need for mixed strategies.

2.4 Dynamic Hotelling-Downs model

The classical framework of electoral competition, in its diverse forms, has an important implication: in a representative democracy, competition leads politicians to adopt moderate policy platforms when office benefit is sufficiently great. This regularity is predicated on the assumptions that candidates have the ability to commit their policy choices and that elections are temporally isolated. In reality, however, elections are repeated, and we cannot dismiss the effect of linkages across time and the importance of time preferences in determining plausible sequences of policies. For instance, Bertola (1993) and Alesina and Rodrik (1994) appeal to the median voter theorem within each period in the context of growth models; more in line

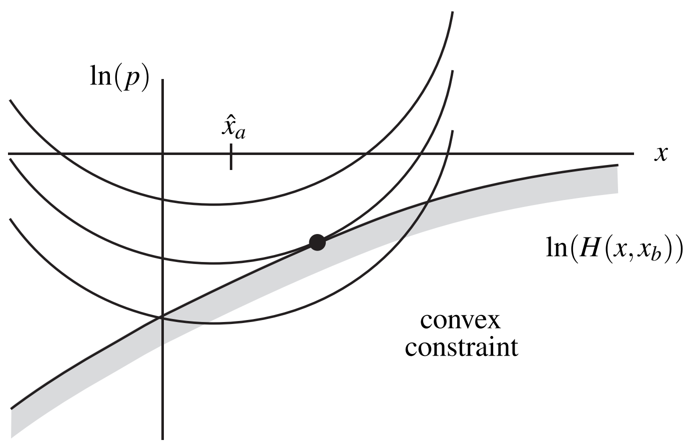

Figure 2: Log concavity

with the treatment here, Basseto and Benhabib (2006) provide conditions for the order restriction to be satisfied over sequences of policies in a dynamic economy.

Under reasonable assumptions, it turns out that if candidates can commit to *sequences* of policies, then the median voter results persists in a strong form. To formalize this, we return to the Hotelling-Downs model and strengthen (A2) to:

(A5) There exist constants θ_j and κ_j for each type $j \in T$ and functions $v: X \to \mathbb{R}$ and $c: X \to \mathbb{R}$ such that for all $x \in X$,

$$u_j(x) = \theta_j v(x) - c(x) + \kappa_j,$$

where $\theta_1 < \theta_2 < \cdots < \theta_n$.

Extending voter preferences to lotteries via expected utility, a straightforward argument (see Duggan 2014b) shows that given any two lotteries on the policy space, say L and L', the difference in expected utility, $\mathbb{E}_L[u_j(x)] - \mathbb{E}_{L'}[u_j(x)]$, is monotonic in the type j. Therefore, voter preferences over lotteries are order restricted, and again the median type m is pivotal in pairwise voting. For example, we obtain quadratic utility $u_j(x) = -(x - \hat{x}_j)^2 + K_j$ by setting $v(x) = 2x$, $c(x) = x^2$, $\theta_j = \hat{x}_j$, and $\kappa_j = -\hat{x}_j^2 + K_j$. For another example, we obtain exponential utility $u_j(x) = -e^{x-\hat{x}_j} + x + K_j$, via a scalar transformation by $e^{-\hat{x}_j}$, upon setting $v(x) = x$, $c(x) = e^x$, $\theta_j = e^{\hat{x}_j}$, and $\kappa_j = e^{\hat{x}_j} K_j$. In Basseto and Benhabib's (2006) economy, all households trade off a measure of distortions against the redistribution implied by the distortions, with households of different wealth disagreeing about the optimal trade-off; we can think about $\theta_j v(x)$ as the redistributed gains (or losses) associated to policy, and about $c(x)$ as the associated distortion losses.

To apply these observations to the dynamic policy model, assume that in an initial election, two office-motivated parties simultaneously announce sequences, $\mathbf{x}_A$ and $\mathbf{x}_B$, of policy platforms. Thus, party A's platform is $\mathbf{x}_A = (x_A^1, x_A^2, \ldots) \in X^\infty$, and likewise for party B's platform. Assume the discount factor $\delta \in [0,1)$ is common to all voters and that voters evaluate sequences of policies according to their discounted utility; for example, party A's platform is preferable to B's for type j voters if and only if

$$(1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u_j(x_A^t) \quad > \quad (1-\delta)\sum_{t=1}^{\infty}\delta^{t-1}u_j(x_B^t),$$

where $1-\delta$ is a normalizing constant. The left-hand side of the latter inequality is equivalent to the type j voter's expected utility from the lottery L that puts probability $(1-\delta)\delta^{t-1}$ on policy x_A^t, and the right-hand side is equivalent to the lottery L' that puts probability $(1-\delta)\delta^{t-1}$ on x_B^t. That is, the discounted utility from a sequence of policies is mathematically equivalent to the expected utility from a particular lottery, and by (A5) it follows that the median type m is pivotal in pairwise votes over policy streams.

A dynamic median voter theorem for the model with unlimited commitment is immediate: when all policy streams are feasible, the unique Nash equilibrium is for both parties to commit to the ideal policy stream $(\hat{x}_m, \hat{x}_m, \ldots)$ for the median voter. But a more general result is possible. Assume that the set of feasible policy streams is $\mathbf{Y} \subseteq X^\infty$, perhaps reflecting the productivity of a durable capital good in a growth economy, and assume that the median voter type has unique ideal feasible policy stream $\hat{\mathbf{x}}_m$.

Proposition 2.5 *Assume (A1) and (A5). In the unique Nash equilibrium of the dynamic Hotelling-Downs model with commitment to streams of policies, we have* $\mathbf{x}_A = \mathbf{x}_B = \hat{\mathbf{x}}_m$.

A premise of representative democracy is, however, that politicians have discretionary power once in office, and the assumption that parties or candidates can commit to policy for an infinite sequence of periods (or even a single period) can reasonably be questioned. Duggan and Fey (2006) maintain the Downsian assumption that parties can commit to policy choices in the current period. They show that the median voter theorem is sensitive to the time preferences of voters and parties: when voters and parties are not too patient, there is a unique subgame perfect path of play (even if complex punishments are possible), and in equilibrium both parties locate at the median; but when players place more weight on future periods than the current one, arbitrary paths of policies can be supported in equilibrium. Alesina (1988) studies a repeated two-party model with probabilistic voting and shows that when candidates cannot commit to policies, Nash-reversion equilibria can be used to support non-trivial equilibria in which candidates' choices diverge from their ideal policies on the equilibrium path of play.

2.5 Citizen-candidate model

The commitment assumption is dropped entirely in the citizen-candidate models of Osborne and Slivinski (1996) and Besley and Coate (1997), where campaigns are viewed as non-binding. In this setting, voters elect a candidate to office, that politician selects a policy, and the game ends. In equilibrium, the winning candidate simply chooses her ideal policy, and in two-candidate equilibria, each citizen simply votes for the candidate whose ideal policy is preferred. Thus, policy choices degenerate, and there is no scope for responsive democracy in the model.

Once we introduce dynamics into the electoral framework, however, informational considerations rise to the fore and can play an important role in escaping the shirking equilibrium. It may be that politicians' preferences are difficult to ascertain before they are elected, and that the policy choices made by politicians while in office may be observed only with noise. The literature on electoral accountability, which is the subject of the remainder of this review, addresses these issues: elections are modeled as a repeated game in which politicians are citizen-candidates (who cannot make binding campaign promises) and have private information about political variables (either their preferences or policy choices or both) relevant to voters. These aspects of elections interact in complex and interesting ways, permitting the analysis of a simple class of equilibria and informing our understanding of the possibility of responsive democracy.

3 Two-period accountability model

3.1 Timing and preferences

This subsection introduces the basic ideas and themes of the accountability literature in a simple model. As in the previous section, we consider a continuum of citizens, N, partitioned into a finite set of types $T = \{1, \ldots, n\}$, with $n \geqslant 2$ and $q_j > 0$ denoting the fraction of type $j \in T$ in the population. Now, there are two periods, $t = 1, 2$. In period 1, a politician is randomly drawn from the population of citizens, with each type j having probability $p_j > 0$, and chooses a policy $x_1 \in X$, where X is a convex (possibly unbounded) subset of $\mathbb{R}$. In period 2, the politician in office, the *incumbent*, faces a randomly drawn *challenger*, with each type j having probability p_j. The winner of the election chooses a policy $x_2 \in X$, and the game ends.

Each period, the policy choice x_t generates a policy outcome y_t in a nonempty, convex (possibly unbounded) outcome space $Y \subseteq \mathbb{R}$. Technically, neither politicians' types nor actions are directly observable by voters, but policy outcomes are. We consider two possibilities: under *perfect monitoring*, the policy outcome is deterministic and equal to the policy choice; under *imperfect monitoring*, the policy outcome depends stochastically on the policy choice. We capture both environments by assuming that outcomes are realized from a distribution function $F(\cdot|x)$ given policy choice x. Under perfect monitoring, we set $Y = X$ and let the distribution of outcomes be degenerate on x, and under imperfect monitoring, we set $Y = \mathbb{R}$ and assume that $F(\cdot|x)$ is continuous with jointly differentiable, strictly positive density $f(y|x)$.

As in the citizen-candidate model, we assume that neither the incumbent nor the challenger can make binding promises before an election. A related point, which does not arise in the static model of elections, is that we also assume voters cannot commit their vote, so that voting as well as policy making must be time consistent. Figure 3 illustrates the timeline of events in the two-period model. First, nature chooses the incumbent's type. Once in office, the incumbent chooses the first-period policy action x_1. Next, a publicly observed outcome y_1 is realized. Then voters vote to re-elect the incumbent or not. Finally, the winner of the election chooses the second-period policy x_2, and the policy outcome y_2 is realized.

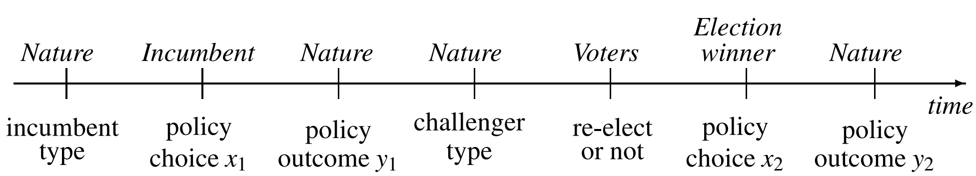

Figure 3: Timeline in two-period model

Given policy choice x and outcome y in any period, type j citizens obtain a payoff of $u_j(y)$ if not in office and a payoff of $w_j(x) + \beta$ if they hold office during the period, where $u_j : Y \to \mathbb{R}$ and $w_j : X \to \mathbb{R}$ are type-dependent functions, and $\beta \geqslant 0$ represents the benefits of holding office. Total payoffs for voters and politicians are the sum of per-period payoffs.

We consider two possible specifications of payoffs in the model.

Spatial preferences We assume that citizens of each type possess policy preferences, and that holding office does not change a citizen's policy preferences, although it may convey a positive benefit. We assume $X = [\underline{x}, \overline{x}]$ is a closed and bounded interval, and we assume perfect monitoring, so that $Y = [\underline{x}, \overline{x}]$. Utility for policies has the simple form

$$u_j(x) \;=\; w_j(x) \;=\; \theta_j v(x) - c(x) + \kappa_j,$$

where $v{:}X \to \mathbb{R}$ is a continuously differentiable, concave, and strictly increasing function, $c{:}X \to \mathbb{R}$ is a continuously differentiable, strictly convex, and strictly increasing function, $\theta_1 < \theta_2 < \cdots < \theta_n$ are type-dependent parameters, and κ_j is a constant that can depend on type. To ensure interior ideal policies, we assume $\theta_1 v'(\underline{x}) > c'(\underline{x})$ and $\theta_n v'(\overline{x}) > c'(\overline{x})$. Without loss of generality, we assume $u_j = w_j$, v, and c take non-negative values. Under our assumptions, each voter type j has an ideal policy $\hat{x}_j$, and these ideal policies are ordered by type, as in (1). We again assume a generic distribution of types among voters, so there is a unique median type m. Assumptions (A1)–(A5) are satisfied by voters' preferences in this environment, and for example, we admit the quadratic and exponential functional forms,

$$u_j(x) = -(x - \hat{x}_j)^2 + K_j \quad \text{and} \quad u_j(x) = -e^{x - \hat{x}_j} + x + K_j,$$

with constant K_j appropriately chosen.[5] In particular, the median type is pivotal in pairwise voting over lotteries over policy. In this version of the model, citizen types can be interpreted as ideological groups with different policy preferences; an alternative is that citizens have common preferences but that the costs and benefits of policy choices are distributed unevenly among citizens, e.g., when all citizens prefer more public good but are taxed differentially due to variation in income.

Rent-seeking In this environment, all voters have increasing preferences over policy outcomes, while a politician who holds office incurs a cost for higher policy choices. We assume $X = \mathbb{R}_+$ and imperfect monitoring, so that $Y = \mathbb{R}$. Utility has the simple form

$$u_j(y) \;=\; u(y) \quad \text{and} \quad w_j(x) \;=\; v(x) - \frac{1}{\theta_j} c(x) + \kappa_j,$$

where $u{:}Y \to \mathbb{R}$ is continuous and strictly increasing, $v{:}X \to \mathbb{R}$ is continuously differentiable, concave, and strictly increasing, $c{:}X \to \mathbb{R}_+$ is continuously differentiable, strictly convex, and strictly increasing, and $0 < \theta_1 < \theta_2 < \cdots < \theta_n$ and

[5] We argue this following (A5), above. The fact that the exponential form is obtained via a scalar transformation by $e^{-\hat{x}_j}$ does not affect the analysis of the spatial preferences model.

$\kappa_1, \ldots, \kappa_n$ are type-dependent parameters. We assume that if in office, each politician type has an optimal policy $\hat{x}_j$, which must be unique, and furthermore we assume that $\kappa_1, \ldots, \kappa_n$ are specified so that the politicians' maximized utility is weakly increasing in type, i.e., for all $j < n$, we have $w_j(\hat{x}_j) \leqslant w_{j+1}(\hat{x}_{j+1})$. To ensure interior ideal policy choices, we also assume $\theta_1 v'(0) > c'(0)$. As in the spatial preferences environment, the ideal policies of office holders are ordered according to type, as in (1). Again, assumptions (A1)–(A5) are satisfied by voters' preferences, and we admit the quadratic and exponential functional forms,[6]

$$w_j(x) = -(x - \hat{x}_j)^2 + K_j \quad \text{and} \quad w_j(x) = -e^{x - \hat{x}_j} + x + K_j.$$

Note that we can assume politicians share the voters' preferences over policy by setting the term $v(x) = \mathbb{E}[u(y)|x]$ equal to the expected utility from policy outcomes generated by the choice x, in which case an office holder differs from other citizens only by the cost term $(1/\theta_j)c(x)$. In this version of the model, policy can be viewed as a level of public good or (the inverse of) corruption, and politician types then reflect different abilities to provide the public good or a distaste for corruption while in office.

3.2 Electoral equilibrium

A *strategy for the incumbent of type* j is a pair $\pi_j = (\pi_j^1, \pi_j^2)$, where

$$\pi_j^1 \in \triangle(X) \quad \text{and} \quad \pi_j^2 : X \times Y \to \triangle(X),$$

specifying policy choices in period 1 and policy choices in period 2 for each possible previous policy choice and observed outcome.[7] Here, π_j^1 has the form of a mixed strategy (a distribution over policy choices), but in the infinite-horizon framework we will interpret π_j^1 as the distribution of pure strategies used by type j politicians; that is, given a subset $Z \subseteq X$ of policies, $\pi_j^1(Z)$ is the fraction of type j politicians who choose policies in Z. For tractability, we impose the restriction that the distribution π_j^1 has finite support for each type. A *strategy for the challenger of type* j is a mapping

$$\gamma_j : Y \to \triangle(X),$$

specifying policy choices in period 2 for each policy type and observed outcome. A *strategy for a voter of type* j is a mapping

$$\rho_j : Y \to [0,1],$$

[6] In the rent-seeking environment, we obtain quadratic preferences via a scalar transformation by $1/\hat{x}_j$. This does not affect supermodularity of politician payoffs, established in Proposition 3.5.

[7] Measurability of strategies or subsets of policies will be assumed implicitly, as needed, without further mention.

where $\rho_j(y)$ is the probability of a vote for the incumbent given outcome y. A *belief system for voters* is a probability distribution $\mu(\cdot|y_1)$ on $T \times X$ as a function of the observed outcome.

A strategy profile $\sigma = (\pi_j, \gamma_j, \rho_j)_{j \in T}$ is *sequentially rational* given beliefs μ if neither the incumbent nor the challenger can gain by deviating from the proposed strategies at any decision node, and if voters of each type vote for a candidate that makes them best off in expectation, given their belief system for any realization of y_1. The latter requirement is needed because in a model with a continuum of voters, no single voter's ballot can affect the outcome of the election; the requirement is consistent with optimization, and it would emerge in a type-symmetric equilibrium of the model if we were to specify that with small probability, the ballot of a type j voter would be randomly drawn to decide the election.[8] Beliefs μ are *consistent* with the strategy profile σ if for every y_1 on the path of play given $(\pi_j^1)_{j\in T}$, the distribution $\mu(j,x|y_1)$ is derived from $(\pi_j^1)_{j\in T}$ via Bayes' rule.[9,10] A *perfect Bayesian equilibrium* of the two-period model is a pair (σ,μ) such that the strategy profile σ is sequentially rational given the beliefs μ, and μ is consistent with σ.

Sequential rationality implies that challengers will choose their ideal policies with probability one, since they cannot hope to be re-elected, so that $\gamma_j(\hat{x}_j|y_1) = 1$ for all y_1. This implies that the expected payoff of electing the challenger for a voter of type j is

$$V_j^C \;=\; \sum_k p_k \mathbb{E}[u_j(y)|\hat{x}_k].$$

Similarly, sequential rationality implies $\pi_j^2(\hat{x}_j|x_1,y_1) = 1$ for all x_1 and all y_1, so the expected payoff from re-electing the incumbent is

$$V_j^I(y_1) \;=\; \sum_k \mu_T(k|y_1)\mathbb{E}[u_j(y)|\hat{x}_k],$$

where $\mu_T(j|y_1)$ is the marginal distribution of the incumbent's type given policy outcome y_1. Since the median voter is pivotal, the incumbent is thus re-elected if $V_m^I(y_1) > V_m^C$ and only if $V_m^I(y_1) \geqslant V_m^C$. Sequential rationality does not pin down the votes of voters when they are indifferent between the incumbent and challenger; we say the equilibrium is *deferential* if voters favor the incumbent when indifferent, so that the incumbent is re-elected if and only if $V_m^I(y_1) \geqslant V_m^C$.

[8] In the terminology of Fearon (1999), voters focus on the problem of "selection," rather than "sanctioning." See his essay for arguments in support of this behavioral postulate.

[9] Bayesian updating is well-defined, as we only consider equilibria in which the mixtures π_j^1 have finite support.

[10] In the model with perfect monitoring, we add the assumption that the marginal on policy choices, $\mu_X(j,x|y_1)$, places probability one on $x = y_1$. This emulates the model in which policy outcomes are observable and chosen directly by the office holder.

This general formulation of deferential equilibrium implies that there is an *acceptance set* of policy outcomes such that the incumbent is re-elected with probability one after realizations in this set and loses for sure after realizations outside the set:

$$A = \{y_1 \in Y : V_m^I(y_1) \geqslant V_m^C\}.$$

We say an equilibrium is *monotonic* if the acceptance set is closed, and if for every policy outcome belonging to the acceptance set, increasing the median voter's utility maintains inclusion in the acceptance set. Formally, for all $y \in A$ and all $y' \in Y$ such that $u_m(\alpha y' + (1-\alpha)y)$ is weakly increasing in $\alpha \in [0,1]$, we have $\alpha y' + (1-\alpha)y \in A$. In the environments we consider, this implies that A is convex, and in the spatial preferences model, that if A is nonempty, then $\hat{x}_m \in A$. The monotonicity condition imposes a link between the voters' utilities over policy outcomes and the informational content of those outcomes in the first period. There could of course be perfect Bayesian equilibria in which this link does not exist—in the spatial environment with perfect monitoring, for example, it could be that the median voter's ideal policy is not chosen in equilibrium, and that voters update negatively following a choice of the median policy off the path of play—but the posited linkage seems natural in the electoral context and simplifies the equilibrium analysis of the model.

An *electoral equilibrium* is a perfect Bayesian equilibrium that is deferential and monotonic. We consider the implications of this equilibrium concept in the context of the models with and without observable policy choices; as we will see, several interesting properties that emerge in the simple two-period model persist in the infinite-horizon model without term limits.

Before proceeding to the equilibrium analysis of the two-period model, note that the assumption that voters observe the incumbent's type directly could be obtained if we fixed the prior p on the challenger's type and allowed the prior beliefs about the incumbent's type to be degenerate. For example, if the voters' prior places probability one on the incumbent being type j, then Bayesian updating does not occur, and we have $\mu_T(j|y_1) \equiv 1$. This implies that the median voter's expected payoff $V_m^I(y_1)$ is constant, and thus either $A = \emptyset$ or $A = Y$, and the median voter's choice $\rho_m \in \{0,1\}$ is constant. Then the first-period office holder solves

$$\max_{x \in X} w_j(x) + \beta + \rho_m[w_j(\hat{x}_j) + \beta] + (1-\rho_m)V_j^C,$$

which has the unique solution $x = \hat{x}^j$. That is, the absence of uncertainty about the incumbent's type removes all reputational concerns of the politician, and the equilibria of the model devolve to the trivial myopic strategies such that each type of politician chooses her ideal policy. This observation holds regardless of whether monitoring is perfect or imperfect and regardless of the preference environment.

3.3 Adverse selection

In this subsection, we focus on the spatial preferences environment, and we assume that the first-period policy choice, x_1, is observable; in other words, the realized policy outcome is $y_1 = x_1$ with probability one. The two-period model with perfect information is analyzed by Reed (1994), who in contrast assumes rent-seeking preferences and examines the optimal re-election rule for voters; we return to this work at the end of the subsection. In the current model, note that a type j office holder's maximum payoff from choosing a policy in the acceptance set is

$$\max_{x \in A} u_j(x) + u_j(\hat{x}_j) + 2\beta,$$

and, assuming $\hat{x}_j \notin A$, the maximum payoff from shirking is

$$u_j(\hat{x}_j) + V_j^C + \beta.$$

Thus, the office holder's choice is dictated by the comparison between $\max_{x \in A} u_j(x)$ and V_j^C.

Of course, in equilibrium, if the first-period office holder's ideal policy belongs to the acceptance set A, then the politician will simply choose that ideal policy and be re-elected. Other office holder types may optimally choose a policy in A, in which case they choose the acceptable policy closest to their ideal policy; and the remaining types simply shirk, choosing their ideal policy and being replaced by an unknown challenger. Let

$$\begin{aligned} W &= \{j \in T : \hat{x}_j \in A\}, \\ C &= \{j \in T \backslash W : \max_{x \in A} u_j(x) + \beta \geqslant V_j^C\}, \\ L &= T \backslash (W \cup C). \end{aligned}$$

We refer to politician types in the set W as "winners," in the set C as "compromisers," and in the set L as "losers."

In this section, we first consider simple conditions such that there exists an electoral equilibrium in which all politician types choose the median ideal policy in the first period. It turns out that the incentives of the extreme types, $j = 1, n$, are critical in determining the possibility of this responsiveness result. These types are willing to compromise to the median if and only if

(B1) $$u_1(\hat{x}_m) + \beta \geqslant V_1^C \quad \text{and} \quad u_n(\hat{x}_m) + \beta \geqslant V_n^C.$$

or equivalently

$$\theta_j (v(\hat{x}_m) - \textstyle\sum_k p_k v(\hat{x}_k)) + \beta \quad \geqslant \quad c(\hat{x}_m) - \textstyle\sum_k p_k c(\hat{x}_k) \qquad (2)$$

for $j = 1, n$, and by linearity of the left-hand side, if the inequality holds for the extreme types, then all politician types are willing to compromise to the median. Condition (B1) holds when office benefit is sufficiently large, and even when $\beta = 0$, it holds when the distribution of challenger types is close to symmetric around the median and the median voter's utility function is close to symmetric around his or her ideal policy. In general, at an intuitive level, the condition holds as long as office benefit is large relative to asymmetries in the model.

Then we specify strategies so that the acceptance set is $A = \{\hat{x}_m\}$ and all politician types choose the median $\hat{x}_m$. These policy strategies are optimal by construction, and given these strategies, on the equilibrium path, the voters' beliefs are equal to their prior, and thus the median voter is indifferent between the incumbent and the challenger. Off the equilibrium path, we specify that voters believe the incumbent is the worst possible type for the median voter, so deviations from the median lead to an electoral loss. These strategies and beliefs form an electoral equilibrium and show how a responsive democracy result can arise in the two-period model. The result holds despite the fact that politicians cannot commit to policy platforms, but it is driven by the voters' incomplete information and the politicians' concern for reputation in the model.

Proposition 3.1 *In the two-period model of adverse selection with spatial preferences and perfect monitoring, assume (B1) holds. Then there is an electoral equilibrium with acceptance set $A = \{\hat{x}_m\}$ and such that every politician type chooses the median policy in the first period, i.e., for all j, we have $\pi_j^1(\hat{x}_m) = 1$.*

The equilibrium constructed above illustrates total compromise, in which every politician type chooses the median policy. To provide insight into electoral equilibria with partial compromise, which arise in the infinite-horizon model, we relax our restriction on parameters. Note that as long as $p_m < 1$, the median voter type will strictly prefer a type m politician to an unknown challenger, i.e., $u_m(\hat{x}_m) > V_m^C$. Let $G_L = \{j \leqslant m : u_m(\hat{x}_j) > V_m^C\}$ denote the set of "above average" types to the left of the median; let $G_R = \{j \geqslant m : u_m(\hat{x}_j) > V_m^C\}$ denote the set of above average types to the right; and let $G = G_L \cup G_R$ be the set of all above average types. Set $\ell = \min G_L$ and $r = \max G_R$. It is straightforward to see that in equilibrium, the above average types must be winning or compromising, i.e., $G \subseteq W \cup C$, for otherwise an above average type loses in equilibrium, but losing politicians choose their ideal policies in the first term, and the median voter would prefer to elect an incumbent after a policy choice that reveals that she is above average.

We next construct an equilibrium with acceptance set equal to a non-degenerate interval, $A = [\underline{x}(\beta), \overline{x}(\beta)]$. The type n politician weakly prefers to compromise at the median rather than shirk if and only if (2) holds for $j = n$. Note that (2) always holds for $j = m$, and thus if the inequality fails for $j = n$, then it must be that $v(\hat{x}_m) - \sum_k p_k v(\hat{x}_k) < 0$. This, in turn, implies that the inequality holds for all $j = 1,\dots,m$, and in particular, it holds for $j = 1$. We conclude that either the type 1 or type n politician is willing to compromise at the median. If both are willing to make this compromise, then (B1) holds, so we henceforth assume that the type n politician is willing to compromise and the type 1 politician is not:

$$\text{(B2)} \qquad u_1(\hat{x}_m) + \beta < V_1^C \quad \text{and} \quad u_n(\hat{x}_m) + \beta \geqslant V_n^C.$$

Note the implication that higher types are more willing to compromise at the median: since (2) holds for $j = n$ and fails for $j = 1$, it follows that $v(\hat{x}_m) - \sum_k p_k v(\hat{x}_k) > 0$, and thus higher θ_j increases the left-hand side of inequality (2). We then specify the upper endpoint of the acceptance set as the median policy, i.e., $\overline{x}(\beta) = \hat{x}_m$.

The lower endpoint of the acceptance set is determined in two cases. In case the type ℓ politician strictly prefers to compromise at the median rather than shirk, i.e.,

$$u_\ell(\hat{x}_m) + \beta \quad > \quad V_\ell^C, \tag{3}$$

we set $\underline{x}(\beta) = \hat{x}_m$. In the remaining case, we define the lower endpoint $\underline{x}(\beta)$ so that the type ℓ politician is indifferent between compromising and shirking:[11] it is the greatest solution to

$$u_\ell(x) + \beta \quad = \quad V_\ell^C. \tag{4}$$

Note that $\underline{x}(\beta)$ increases with β until $\underline{x}(\beta) = \hat{x}_m$, at which point the first case obtains. Letting $\tilde{x} = \sum_k p_k \hat{x}_k$ be the expected policy choice of a challenger in the second period, concavity of v and convexity of c imply

$$\theta_j(v(\tilde{x}) - \textstyle\sum_k p_k v(\hat{x}_k)) + \beta \ > \ 0 \ > \ c(\tilde{x}) - \textstyle\sum_k p_k c(\hat{x}_k),$$

and thus we have $\tilde{x} < \underline{x}(\beta)$. Moreover, since v is increasing, we conclude that $v(\underline{x}(\beta)) - \sum_k p_k v(\hat{x}_k) > 0$. This has important implications for the incentives of more extreme types to compromise: because the type j politician is indifferent between shirking and compromising at $\underline{x}(\beta)$, it follows that types $1,\dots,\ell-1$ will not be willing to compromise.

In addition, we assume that conditional on having drawn the incumbent from the right side of the median type (and including the median type m), the median

[11] Because u_ℓ is strictly concave, this indifference condition can have at most two solutions, one below and one above the type ℓ politician's ideal policy.

voter's expected payoff of re-electing the incumbent is at least as great as that from drawing the challenger at large:

$$\text{(B3)} \qquad \frac{\sum_{k:k\geqslant m} p_k u_m(\hat{x}_k)}{\sum_{k:k\geqslant m} p_k} \;\geqslant\; V_m^C.$$

This condition is satisfied if the probability of a median type challenger is sufficiently large relative to right-skewness of the distribution of challenger types: as long as the probability of a type m challenger is positive, it is satisfied when the distribution of challenger types is close to symmetric around (or skewed to the left of) the median and the median voter's utility function is close to symmetric around his or her ideal policy.

Under these conditions, we obtain a partial median voter theorem in which a set of compromising politician types choose the median policy in the first period, some choose their ideal policies and are re-elected, some may pool at an alternative closer to (but distinct from) the median than their ideal policies, and some other types shirk instead. The equilibrium has a *partitional* structure, in which the winning types are centrally located and form a "connected" set of types around the median, with compromising types surrounding the winning types, and losing types at one extreme of the policy space. This structure is illustrated in Figure 4. The top panel shows the equilibrium in case the type ℓ politician is willing to compromise to the median, and the bottom panel depicts the equilibrium in the second case, where where the lower endpoint is given by the type ℓ politician's indifference condition and the acceptance set is the darkened interval.[12]

Proposition 3.2 *In the two-period model of adverse selection with spatial preferences, assume (B2) and (B3). Then there is an electoral equilibrium with acceptance set* $A = [\underline{x}(\beta), \overline{x}(\beta)]$ *such that* $\overline{x}(\beta) = \hat{x}_m$ *and:*

(i) if β *satisfies (3), then* $\underline{x}(\beta) = \hat{x}_m$*, and there exists* $k \in \{2, \ldots, \ell\}$ *such that*

$$L = \{1, \ldots, k-1\} \;\; \textit{and} \quad C = \{k, \ldots, n\} \quad \textit{and} \quad W = \{m\}$$

(ii) otherwise, $\underline{x}(\beta)$ *is the greatest solution to (4), and there exists* $k \in \{\ell + 1, \ldots, m\}$ *such that*

$$L = \{1, \ldots, \ell-1\} \;\; \textit{and} \quad C = \begin{array}{c}\{\ell, \ldots, k-1\} \cup \\ \{m+1, \ldots, n\}\end{array} \quad \textit{and} \quad W = \{k, \ldots, m\}.$$

[12]The statement of Proposition 3.2 implicitly relies on the fact that if (4) holds, then $\ell < m$.

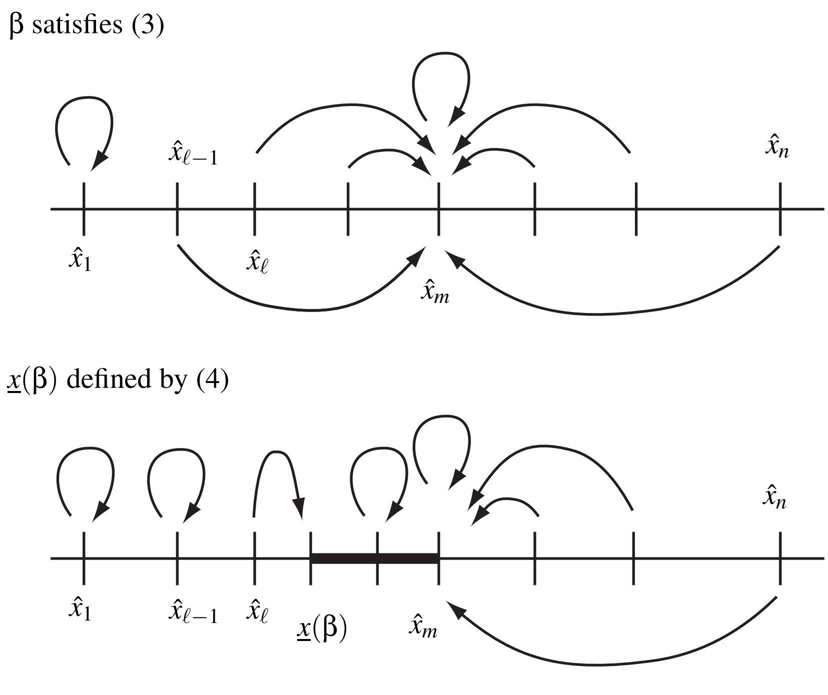

Figure 4: Partial compromise equilibria

For the equilibrium construction when β satisfies (3), we specify that the acceptance set is $A = \{\hat{x}_m\}$, and we specify that politician types $j \geqslant m$ choose the median policy $\hat{x}_m$, and types $j \leqslant m$ choose $\hat{x}_m$, unless they strictly prefer to shirk. Since higher types are more willing to compromise at the median it follows that there is a type k as in (i) such that types $j \geqslant k$ are willing to compromise at the median, while types $j < k$ are not. And since the set of types who compromise at the median truncates only below average types, the median voter is indeed willing to re-elect the incumbent, conditional on the choice of the median policy. To justify this claim more formally, note that

$$\begin{aligned} V_m^C &= \sum_j p_j u_m(\hat{x}_j) \\ &= \sum_{j=1}^{k-1} p_j u_m(\hat{x}_j) + \sum_{j=k}^{n} p_j u_m(\hat{x}_j) \\ &\leqslant (\sum_{j=1}^{k-1} p_j) V_m^C + \sum_{j=k}^{n} p_j u_m(\hat{x}_j), \end{aligned}$$

where the inequality uses the fact that all types $j = 1, \ldots, k-1$ are below average. This implies

$$V_m^C \leqslant \sum_{j=k}^{n} \left(\frac{p_j}{\sum_{h=k}^{n} p_h} \right) u_m(\hat{x}_j) = \sum_j \mu_T(j|\hat{x}_m) u_m(\hat{x}_j) = V_m^I(\hat{x}_m).$$

Finally, given policy choice $x \neq \hat{x}_m$, if this is the ideal policy of a politician type that does not compromise, i.e., $x = \hat{x}_j$ for some j outside $C \cup W$, then the politician is below average, so the median voter prefers to elect the challenger. And if x is off the equilibrium path, then we assign voter beliefs so that voters believe the incumbent is the worst possible type for the median voter, so once again the median voter prefers to elect the challenger.

When $\underline{x}(\beta)$ is defined by (4), we specify that the acceptance set is $A = [\underline{x}(\beta), \hat{x}_m]$ and that the type ℓ politician compromises at $\underline{x}(\beta)$, while above average types $\ell+1, \ldots, m$ choose the closest acceptable policy to their ideal policy, and below average types $1, \ldots, \ell-1$ shirk. We have noted that $v(\underline{x}(\beta)) - \sum_k p_k v(\hat{x}_k) > 0$, so higher types are more willing to compromise at $\underline{x}(\beta)$, and therefore these politician strategies are optimal given the acceptance set. Since only above average types choose $\underline{x}(\beta)$ and the ideal policies $\hat{x}_k, \ldots, \hat{x}_{m-1}$, the median voter weakly prefers to re-elect the incumbent after these choices are made. And by (B3), the median voter also weakly prefers the incumbent to the challenger conditional on $\hat{x}_m$. Since policy choices $\hat{x}_1, \ldots, \hat{x}_{k-1}$ are made by below average types, the median voter prefers to replace the incumbent with a challenger following these choices. Finally, for x off the equilibrium path, if x is outside the acceptance set, then we assign voter beliefs so that voters believe the incumbent is the worst possible type for the median voter; and if x belongs to the acceptance, then we specify that voter beliefs are equal to their prior, completing the equilibrium construction.

Proposition 3.1 implies that if the office benefit β is sufficiently large, then there is an electoral equilibrium in which all types compromise at the median, and Proposition 3.2 extends this result to allow for low office benefit, while imposing some limitation on skewness of the distribution of challenger ideal policies. Both of these results permit the existence of other electoral equilibria in which the responsive democracy result fails. Next, we address this issue by assuming sufficiently high β and strengthening (B3) to impose strict inequalities in both directions. Conditional on having drawn the incumbent from one side of the median type (and including the median type m), the median voter's expected payoff of re-electing the incumbent is strictly greater than drawing the challenger at large:

(B4) $$\frac{\sum_{k:k \leqslant m} p_k u_m(\hat{x}_k)}{\sum_{k:k \leqslant m} p_k} > V_m^C \quad \text{and} \quad \frac{\sum_{k:k \geqslant m} p_k u_m(\hat{x}_k)}{\sum_{k:k \geqslant m} p_k} > V_m^C.$$

This condition is satisfied if the probability of a median type challenger is sufficiently large relative to asymmetries in the model. The next proposition provides a strong responsive democracy result by establishing that when (B1) and (B4) hold, all politician types choose the median policy in the first period in *every electoral equilibrium.*

Proposition 3.3 *In the two-period model of adverse selection with spatial preferences, assume (B1) and (B4). Then in every electoral equilibrium* (σ,μ)*, each type* j *chooses the median policy in the first period, i.e.,* $\pi^1_j(\hat{x}_m) = 1$.

To see the result, note that in an electoral equilibrium σ, a finite number of policies, say $x^1 < x^2 < \cdots < x^\ell$, are chosen with positive probability in the first period, and we can write $\mu_T(j|i)$ for the voters' posterior belief that the incumbent is type j conditional on observing choice x^i. Let $\mu(i)$ denote the probability of observing x^i in the first period. Then we have the following helpful accounting identity, which holds for all strategy profiles:

$$\sum_{i=1}^{\ell}\mu(i)V^I_m(x^i|\sigma) \;=\; \sum_{i=1}^{\ell}\mu(i)\sum_{j=1}^{n}u_m(\hat{x}_j)\mu_T(j|i) \;=\; \sum_j p_j u_m(\hat{x}_j) \;=\; V^C_m,$$

where we use the fact that for all politician types j, we have

$$p_j \;=\; \sum_{i=1}^{\ell}\mu(j|i)\mu(i).$$

It follows that $V^I_m(x^i|\sigma) \geqslant V^C_m$ for some type j. Now using the assumption that σ is an electoral equilibrium, it is deferential, and so the acceptance set A is nonempty. By monotonicity, it contains the median policy $\hat{x}_m$, and (B1) implies that every politician type chooses a policy in the acceptance set and is re-elected. In this case, the accounting equation implies $V^I_m(x^i|\sigma) = V^C_m$ for all i, so that after every equilibrium policy choice in the first period, the median voter is indifferent between re-electing the incumbent and selecting a challenger.

Suppose that there are three or more distinct policies chosen with positive probability, so $\ell \geqslant 3$. Since the acceptance set is an interval and equilibrium policy choices are optimal, it follows that: for all j with $\hat{x}_j \leqslant x^1$, type j politicians choose x^1; for all j with $x^1 \leqslant \hat{x}_j \leqslant x^\ell$, type j politicians choose their ideal policy; and for all j with $x^\ell \leqslant \hat{x}_j$, type j politicians choose x^ℓ. Assume without loss of generality that $V^I_m(x^1|\sigma) \leqslant V^I_m(x^\ell|\sigma)$, and note that the median voter's expected payoff from re-election, $V^I_m(x^i|\sigma)$, is minimized at $x^i = x^1$. Moreover, by concavity, we actually have $V^I_m(x^2|\sigma) > V^I_m(x^1|\sigma)$. Indeed, if $x^2 = \hat{x}_m$, then the choice of x^2 reveals that the politician is the median type, so $V^I_m(x^2|\sigma) = u_m(\hat{x}_m)$, and the inequality follows. Otherwise, we have $x^2 < \hat{x}_m$, and there is some type j politician such that $x^2 = \hat{x}_j$, which implies $V^I_m(x^2|\sigma) = u_m(\hat{x}_j)$, whereas politician types pooling at x^1 are further from the median, so once again the inequality follows, contradicting our observation that all equilibrium policy choices in the first period determine the same expected payoff from re-election for the median voter.

Thus, we have $\ell \leqslant 2$. Since the type m politician chooses $\hat{x}_m \in A$ in an electoral equilibrium, we can assume $x^2 = \hat{x}_m$ without loss of generality. We know that each type j with $\hat{x}_j \leqslant x^1$ chooses x^1, and so the expected payoff to the median voter from re-election of the incumbent conditional on x^1 is maximized when all types $j < m$ choose x^1. Let σ' be this strategy profile. Using the accounting identity and (B4), we have

$$V^I_m(x^2|\sigma') \;=\; \frac{\sum_{k:k\geqslant m} p_k u_m(\hat{x}_k)}{\sum_{k:k\geqslant m} p_k} \;>\; V^C_m \;=\; \mu'(1)V^I_m(x^1|\sigma') + \mu'(2)V^I_m(x^2|\sigma').$$

This implies that $V^C_m > V^I_m(x^1|\sigma') \geqslant V^I_m(x^1|\sigma)$, a contradiction. Thus, we must have $\ell = 1$, and therefore every politician type chooses the median ideal policy, completing the proof.

Note that the centripetal effect of electoral incentives highlighted in Propositions 3.1–3.3 derives from the informational asymmetry in the model. An extremist office holder has incentives to pool at the median ideal policy in order to avoid appearing extremist, but if voters observed politicians' types, then this incentive would be removed. Thus, asymmetric information can facilitate responsive democracy, whereas full transparency leads to shirking. This observation anticipates an "anti-folk theorem" for the version of the infinite-horizon model with in which the incumbent's type is observed by voters, stated in Section 4.

The equilibria highlighted in Proposition 3.1–3.3 have a structure similar to equilibria in the infinite-horizon model. As in the current subsection, we will see that in the infinite-horizon model, when the office benefit is high enough, we obtain the responsive democracy result that all politician types pool at (or close to) the median voter's ideal policy in every equilibrium. Moreover, the partitional form of equilibrium in Proposition 3.2 extends to an interesting class of equilibria in the infinite-horizon model. A technical difference between the two-period and infinite-horizon model is that in the current setting, compromise incentives in (2) are linear in θ_j, which implies that at least one extreme type is willing to compromise at the median policy. In particular, the equilibrium in Proposition 3.2 cannot support losers on both sides of the median policy; in the infinite-horizon model, this is a possibility. Interestingly, equilibria of the two-period model do *not* approximate equilibria of the infinite-horizon model with a two-period term limit. In the latter model, we do not obtain the strong responsiveness results from Proposition 3.1–3.3

when the discount factor and office benefit are high, because voters face a commitment problem: if all politician types were to choose the median policy in their first term of office, then the median voter would always prefer to elect the challenger, but this removes the incentive of first-term office holders to compromise.

The two-period model with perfect monitoring is investigated by Reed (1994), but he considers the rent-seeking environment with continuously distributed types, and rather than analyzing electoral equilibria, he focusses on the distinction between performance and selection effects, and he considers the retrospective voting rule that maximizes expected effort. In response to a cutoff for re-election, politician types partition themselves in the first period into winning, compromising, and losing sets, exemplifying the partitional structure highlighted in Proposition 3.2, and they choose their ideal effort levels in the second period. A drawback of the optimal re-election rule, however, is that information is revealed by the policy choice of the incumbent in the first period, so that the cutoff may be time-inconsistent, in the sense that it can require voters to replace an incumbent who is superior to an untried challenger.

3.4 Moral hazard and adverse selection

We now suppose that in addition to a politician's type being private information, the first-period office holder's action x is not observed directly by voters; rather, we assume voters observe a noisy outcome y realized from a differentiable, positive density $f(\cdot|x_1)$. That is, we combine adverse selection and moral hazard in the two-period framework. Furthermore, we focus on the rent-seeking environment, where voters have common preferences that are monotonically increasing in y, while politicians internalize the cost of the policy x and have ideal policy choices $\hat{x}_1 < \cdots < \hat{x}_n$. Fearon (1999) studies a related model, the difference being that he assumes a random shock added directly to the voter's utility, and not to the underlying policy outcome.[13] Chapter 3 of Besley (2006) presents a two-period, two-type model in which the first-period office holder observes the values of a binary state of the world and preference shock, followed by a binary policy choice. Closer to the model of this section, Chapter 4 (coauthored with Michael Smart) of the book investigates a two-type model in which an office holder essentially chooses a level x of shirking, and voters observe this with noise, $x+\varepsilon$, but it is assumed that the first-period politician observes the policy shock ε before her choice; in addition, the policy choice of the good type of politician is fixed exogenously. Ashworth and Bueno de Mesquita (2014) consider the effect of varying voter information in two simplified models of adverse selection and moral hazard.

[13] The approaches are interchangeable when voters are risk neutral, but not otherwise.

Chapter 4 of Persson and Tabellini (2000) contains a simplified, two-period model of symmetric learning, in which politicians are parameterized by a skill level that is unobserved by voters and politicians themselves. In this setting, voters and politicians update their beliefs symmetrically along the equilibrium path, and signaling cannot occur. Moreover, voters are assumed to be risk neutral. Ashworth (2005) considers a three-period model of symmetric learning that further differs from ours in that the skill level of a politician evolves over time according to a random walk.[14] Ashworth and Bueno de Mesquita (2008) use a variant of the model, one in which the voter has quadratic policy utility and a stochastic partisan preference, to establish existence and comparative statics of incumbency advantage. We consider the symmetric learning environment separately in the infinite-horizon model in Subsection 6.4.

Other work, including Barganza (2000) and Canes-Wrone, Herron, and Shotts (2001), studies a two-type model in which politicians differ in ability. In the latter paper, the voter's desired policy depends on the realization of a state of the world, about which politicians are better informed. Politicians may have an incentive to pander to voters by knowingly choosing policies that are not in the voters' best interest. Maskin and Tirole (2004) study pandering in a two-type model in which politicians differ in preferences. Austen-Smith and Banks (1989) investigate the voters' ability to discipline politicians when all politicians have the same preferences, so that the model is one of pure moral hazard.

The development in this subsection is based on Duggan and Martinelli (2017), who establish existence and provide a characterization of electoral equilibria in the current framework. For simplicity we take the policy choice x to be a shift parameter on the density of outcomes, so, abusing notation slightly, the density can be written $f(y|x) = f(y-x)$ for strictly positive density $f(\cdot)$, and the probability that the realized outcome is less than y given policy x is simply $F(y-x)$. We assume that f satisfies the monotone likelihood ratio property (MLRP), i.e.,

$$\text{(C1)} \qquad \frac{f(y-x)}{f(y-x')} > \frac{f(y'-x)}{f(y'-x')} \quad \text{for all } x > x' \text{ and all } y > y'.$$

This implies that greater policy outcomes induce voters to update favorably their beliefs about the policy adopted by the incumbent in the first period. As is well-known, the MLRP implies that the density function is unimodal, and that both the density and the distribution functions are strictly log-concave. Moreover, we assume $Y = \mathbb{R}$ and

$$\text{(C2)} \qquad \lim_{y\to-\infty} \frac{f(y-x)}{f(y-x')} = \lim_{y\to+\infty} \frac{f(y-x')}{f(y-x)} = 0 \text{ for all } x > x'.$$

[14]Although the model assumes three periods, the first-term office holder has private information about her ability only in the second and third terms, as her action in office are hidden from voters.

As an example, $f(\cdot)$ may be a normal density with arbitrary mean and variance.

In this setting, electoral equilibrium implies that voters follow a simple retrospective rule: there exists $\overline{y} \in \mathbb{R} \cup \{-\infty, \infty\}$ such that they re-elect the incumbent if and only if $y \geqslant \overline{y}$, i.e., $A = [\overline{y}, \infty)$.[15] Electoral equilibria are then characterized by three conditions. First, the threshold $\overline{y}$ must be such that, anticipating that politicians choose their ideal policies in the second period, the expected utility of re-electing the incumbent conditional on observing y is greater than or equal to $\sum_k p_k \mathbb{E}[u(y)|\hat{x}_k]$ if and only if $y \geqslant \overline{y}$. Second, each politician type j, knowing that she is re-elected if and only if $y \geqslant \overline{y}$, mixes over optimal actions in the first period, i.e., the type j politician's policy strategy π_j places probability one on maximizers of

$$w_j(x) + (1 - F(\overline{y} - x))[w_j(\hat{x}_j) + \beta] + F(\overline{y} - x)V^C. \tag{5}$$

Third, updating of voter beliefs follows Bayes rule, i.e., after observing outcome y, the voters' posterior beliefs assign probability

$$\mu_T(j|y) \quad = \quad \frac{p_j \sum_x f(y-x)\pi_j(x)}{\sum_k p_k \sum_x f(y-x)\pi_k(x)}$$

to the incumbent being type j. Since the outcome density is positive, every outcome is on the path of play, so Baye's rule pins down the voters' beliefs; thus, in the remainder of this section, we summarize an electoral equilibrium by the strategy profile σ, leaving beliefs implicit.

We assume that all politicians are in principle interested in re-election, i.e.,

$$\text{(C3)} \qquad w_1(\hat{x}_1) + \beta \; > \; V^C,$$

so that if re-election is assured by choosing their ideal policies in the first period, then the benefits of re-election outweigh the costs. Note that an office holder can always choose her ideal policy, so it is never optimal for the politician to choose large policies for which $\mathbb{E}[u(y)|\hat{x}_j] > w_j(x) + \beta$. By (C3), it is never optimal to choose a policy below the politician's ideal policy, so there is at least one solution to the office holder's problem in the first period. Denoting by x_j^* such a solution, the necessary first order condition for a solution of the office holder's maximization problem (5) is

$$w_j'(x_j^*) \quad = \quad -f(\overline{y} - x_j^*)[w_j(\hat{x}_j) + \beta - V^C]. \tag{6}$$

That is, the marginal disutility in the current period from increasing the policy choice is just offset by the marginal utility in the second period, owing to the

[15] Or $A = \mathbb{R}$ if $\overline{y} = -\infty$.

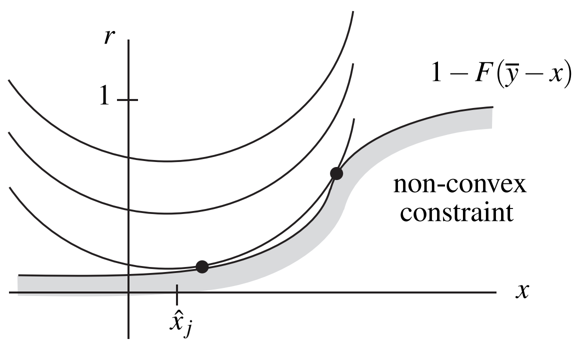

Figure 5: Politician's optimization problem

politician's increased chance of re-election. By (C3), the right-hand side of (6) is negative, and we see that for an arbitrary cutoff, the politician optimally exerts a positive amount of effort, i.e., chooses $x_j^* > \hat{x}_j$, in the first term of office.

We can gain some insight into the incumbent's problem by reformulating it in terms of optimization subject to an inequality constraint. Define a new objective function

$$U_j(x,r) \quad = \quad w_j(x) + r[w_j(\hat{x}_j) + \beta - V^C],$$

which is the expected utility if the politician chooses policy x and is re-elected with probability r, minus a constant term. Note that U_j is concave in (x,r) and quasi-linear in r. Of course, given x, there is only one possible re-election probability, namely $1 - F(\overline{y} - x)$. Defining the constraint function

$$g(x,r) \quad = \quad r - 1 + F(\overline{y} - x),$$

we can then formulate the politician's optimization problem as

$$\begin{gathered} \max_{(x,r)} U_j(x,r) \\ \text{s.t. } g(x,r) \leqslant 0, \end{gathered}$$

which has the general form depicted in Figure 5. Here, level sets of the objective function are well-behaved, but the constraint inherits the natural non-convexity of the distribution function F, leading to the possibility of multiple solutions. This, in turn, can lead to multiple optimal policies and the necessity of mixing in equilibrium, as encountered in the probabilistic voting model and depicted in Figure 1.

We exploit log concavity and impose further restrictions on the risk aversion of politicians to limit the need for mixing to at most two policy choices for each type.

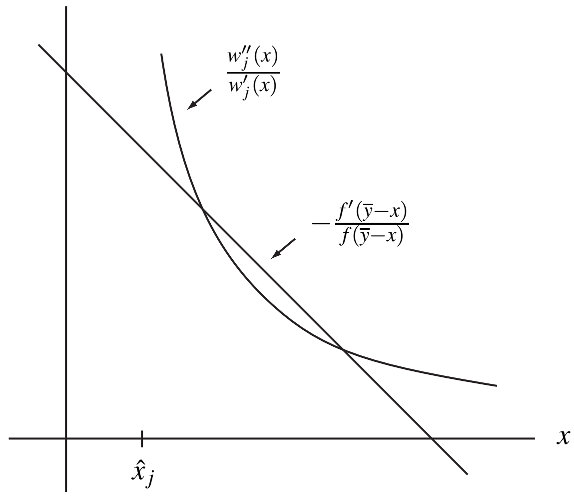

Figure 6: Quadratic-normal special case

Assume that for all j, all finite $\overline{y}$, and all $x, \tilde{x}, z$ with $\hat{x}_j < x < \tilde{x} < z$, we have

(C4)
$$\text{if } \frac{w_j''(x)}{w_j'(x)} \leqslant -\frac{f'(\overline{y}-x)}{f(\overline{y}-x)} \text{ and } \frac{w_j''(z)}{w_j'(z)} \leqslant -\frac{f'(\overline{y}-z)}{f(\overline{y}-z)},$$
$$\text{then } \frac{w_j''(\tilde{x})}{w_j'(\tilde{x})} < -\frac{f'(\overline{y}-\tilde{x})}{f(\overline{y}-\tilde{x})}.$$

That is, the set of $x > \hat{x}_j$ such that $\frac{w_j''(x)}{w_j'(x)} \leqslant -\frac{f'(\overline{y}-x)}{f(\overline{y}-x)}$ is convex, and if x and z satisfy the inequality, then every policy between them satisfies it strictly. To see the permissiveness of this condition, note that by log concavity of $f(\cdot)$, the term $-\frac{f'(\overline{y}-x)}{f(\overline{y}-x)}$ is strictly decreasing in x, and thus (C4) is satisfied if the coefficient of absolute risk aversion, $\frac{w_j''(x)}{w_j'(x)}$, does not decrease too rapidly to the right of the type j politicians' ideal policy. To illustrate, when the utility function w_j is quadratic, the coefficient of absolute risk aversion is $\frac{1}{x-\hat{x}_j}$, and when the density f is normal, the likelihood ratio $-\frac{f'(\overline{y}-x)}{f(\overline{y}-x)}$ simplifies to $\overline{y}-x$. Thus, (C4) is satisfied in the quadratic-normal special case, depicted in Figure 6. Likewise, in the case of exponential utility, the coefficient of risk aversion is $\frac{1}{1-\exp(\hat{x}_j-x)}$, and again (C4) is satisfied. Technically, (C4) works by precluding a local maximizer of the objective function of the politician between two local minimizers, and the proof simply applies the second order conditions associated to the politician's problem.

The usefulness of (C4) is delineated by Duggan and Martinelli (2017) in the next result, which implies that for arbitrary cutoffs, each type of office holder has at most two optimal policies. We let $x^*_j(\overline{y})$ denote the greatest solution to the incumbent's optimization problem and $x_{*,j}(\overline{y})$ the least, as a function of the cutoff. Intuitively, assuming the type j politician has two optimal policy choices, so $x_{*,j}(\overline{y}) < x^*_j(\overline{y})$, the politician is indifferent between the policies: the higher one is more costly but leads to a higher probability of winning, and these considerations exactly offset each other. Thus, we can think of the higher policy as "going for broke," while the lower policy as "taking it easy." Of course, standard continuity arguments imply that the correspondence of optimal policies has closed graph; in the present context, this means that the functions $x^*_j(\cdot)$ and $x_{*,j}(\cdot)$ are, respectively, upper and lower semi-continuous.

Proposition 3.4 *In the two-period model of moral hazard with rent-seeking, assume (C1)–(C4). Then for every cutoff $\overline{y} \in \mathbb{R}$ and every type j, there are at most two local maximizers of the objective function (5), and the greatest and least optimal policies, $x^*_j(\overline{y})$ and $x_{*,j}(\overline{y})$, are upper semi-continuous and lower semi-continuous, respectively, as a function of the cutoff.*

Suppose there are three distinct local maximizers of the type j politicians' objective function, say x', x'', and x''' with $x' < x'' < x'''$. Thus, there are local minimizers z' and z'' such that $x' < z' < x'' < z'' < x'''$. With (C3), inspection of the first order condition (6) at x reveals that $w'_j(z') < 0$ and $w'_j(z'') < 0$, and we can rewrite the first order condition at z' and z'' as

$$w_j(\hat{x}_j) + \beta - V^C \;=\; -\frac{w'_j(z')}{f(\overline{y}-z')} \;=\; -\frac{w'_j(z'')}{f(\overline{y}-z'')}.$$

By the necessary second order condition for a local minimizer, the second derivative at z' satisfies

$$0 \;\leqslant\; w''_j(z') - f'(\overline{y}-z')[w_j(\hat{x}_j) + \beta - V^C] \;=\; w''_j(z') - f'(\overline{y}-z')\left[-\frac{w'_j(z')}{f(\overline{y}-z')}\right],$$

or equivalently,

$$\frac{w''_j(z')}{w'_j(z')} \;\leqslant\; -\frac{f'(\overline{y}-z')}{f(\overline{y}-z')}.$$

Similarly, we have

$$\frac{w''_j(z'')}{w'_j(z'')} \;\leqslant\; -\frac{f'(\overline{y}-z'')}{f(\overline{y}-z'')}.$$

Since x'' is a local maximizer, the first order condition holds at x'', and the second derivative at x'' is non-positive, but then we have

$$\frac{w_j''(x'')}{w_j'(x'')} \quad \geqslant \quad -\frac{f'(y-x'')}{f(y-x'')},$$

contradicting (C4). We conclude that the objective function has at most two local maximizers, as desired.

We can illustrate Proposition 3.4 assuming the normal density and exponential utility with ideal policy equal to zero, $w_j(x) = -e^x + 1$. Then the first order condition is

$$e^x \quad = \quad \frac{\Delta}{\sigma\sqrt{2\pi}} e^{-\frac{(\overline{y}-x)^2}{2\sigma^2}},$$

where $\Delta = \beta - w_j(\hat{x}_j) - V^C = \beta - V^C > 0$. Multiplying by negative one and taking logs of both sides, this is a quadratic equation in x, with solutions

$$x \quad = \quad \overline{y} - \sigma^2 \pm \sigma\sqrt{\sigma^2 - 2\overline{y} + 2\ln\left(\frac{\Delta}{\sigma\sqrt{2\pi}}\right)}.$$

The solutions are real as long as office benefit is sufficiently high relative to the cutoff, and otherwise there is no solution to the first order condition, so that the politician optimizes at the corner by choosing zero effort. Alternatively, the solutions are real if the variance of the observed outcome is sufficiently small. Note that the optimal effort increases without bound as the office benefit becomes large; we return to the latter observation in our analysis of responsive democracy, below.

The next proposition establishes that the politicians' objective functions satisfy the important property that differences in payoffs are monotone in type. We say that $U_j(x, 1-F(\overline{y}-x))$ is *supermodular* in (j,x) if for all (j,x) and all (k,z) with $j > k$ and $x > z$, we have

$$\begin{aligned} &U_j(x, 1-F(\overline{y}-x)) - U_j(z, 1-F(\overline{y}-z)) \\ &\quad > \quad U_k(x, 1-F(\overline{y}-x)) - U_k(z, 1-F(\overline{y}-z)). \end{aligned}$$

An implication is that given an arbitrary value $\overline{y}$ of the cutoff, the optimal policy choices of the types are strictly ordered by type, i.e.,

$$\text{for all } j < n, \quad x_j^*(\overline{y}) \; < \; x_{*,j+1}(\overline{y}).$$

This ordering property will, in turn, be critical for establishing existence of equilibrium.

Proposition 3.5 *In the two-period model of moral hazard with rent-seeking, the type j politician's objective function, $U_j(x, 1 - F(\overline{y} - x))$, is supermodular in (j, x).*

To see the result, consider $j > k$ and $x > z$. We must show that

$$\begin{aligned} & w_j(x) - w_j(z) + (F(\overline{y} - z) - F(\overline{y} - x))[w_j(\hat{x}_j) + \beta - V^C] \\ & \quad > \quad w_k(x) - w_k(z) + (F(\overline{y} - z) - F(\overline{y} - x))[w_k(\hat{x}_k) + \beta - V^C]. \end{aligned}$$

Moreover, (C3) implies that $w_k(\hat{x}_k) + \beta - V^C > 0$. Since $x > z$, we have $F(\overline{y} - z) - F(\overline{y} - x) > 0$, and the assumption that $w_j(\hat{x}_j) \geqslant w_k(\hat{x}_k)$ yields

$$\begin{aligned} & (F(\overline{y} - z) - F(\overline{y} - x))[w_j(\hat{x}_j) + \beta - V^C] \\ & \quad > \quad (F(\overline{y} - z) - F(\overline{y} - x))[w_k(\hat{x}_k) + \beta - V^C]. \end{aligned}$$

It then remains to be shown that $w_j(x) - w_j(z) > w_k(x) - w_k(z)$. The latter inequality simplifies to

$$\left(\frac{1}{\theta_k} - \frac{1}{\theta_j}\right)(c(x) - c(z)) \quad > \quad 0.$$

Since $\theta_j > \theta_k$ and c is strictly increasing, the desired inequality holds.

The above ordering property is very useful in combination with the fact that given arbitrary policy choices $x_1 < x_2 < \cdots < x_n$ of the politician types in the first period, there is a unique outcome, which we denote $y^*(x_1, \ldots, x_n)$, such that conditional on realizing this value, the voters are indifferent between re-electing the incumbent and electing a challenger. Moreover, this extends to the case of mixed policy strategies $\pi_1, \ldots, \pi_n$ with supports that are strictly ordered by type, i.e., for all $j < n$,

$$\max\{x : \pi_j(x) > 0\} \quad < \quad \min\{x : \pi_{j+1}(x) > 0\}.$$

That is, there is a unique solution in $\overline{y}$ to the equation $V^I(\overline{y}) = V^C$, or more explicitly,

$$\sum_k \mu_T(k|\overline{y})\mathbb{E}[u(y)|\hat{x}_k] \quad = \quad \sum_k p_k \mathbb{E}[u(y)|\hat{x}_k]. \tag{7}$$

We let $y^*(\pi_1, \ldots, \pi_n)$ denote the solution to the voter's indifference condition as a function of policy choices.

In addition to uniqueness, the next proposition establishes that the cutoff lies between the choices of the type 1 and type n politicians, shifted by the mode of the density of $f(\cdot)$, which we denote by $\hat{z}$.

Proposition 3.6 *In the two-period model of moral hazard with rent-seeking, assume (C1)–(C4). Then for all mixed policy strategies $\pi_1, \ldots, \pi_n$ with supports that are strictly ordered by type and for all belief systems μ derived via Bayes rule, there is a unique solution to the voters' indifference condition (7), and the solution $y^*(\pi_1, \ldots, \pi_n)$ is continuous as a function of mixed policies. Moreover, this solution lies between the extreme policy choices shifted by the mode of the outcome density, i.e.,*

$$\min\{x : \pi_1(x) > 0\} + \hat{z} \ \leqslant\ y^*(\pi_1, \ldots, \pi_n) \ \leqslant\ \max\{x : \pi_n(x) > 0\} + \hat{z}.$$

For existence of a solution to the indifference condition, fix $\pi_1, \ldots, \pi_n$ with supports that are strictly ordered by type, and note that the left-hand side of (7) is continuous in $\overline{y}$. For any $j < n$, let $x_j = \max\{x : \pi_j(x) > 0\}$ be the greatest policy chosen with positive probability by the type j politicians, and let $x_n = \min\{x : \pi_n(x) > 0\}$ be the lowest policy chosen with positive probability by the type n politicians. For all j and all $x < x_j$ with $\pi_j(x) > 0$, (C2) implies that for sufficiently large $\overline{y}$, we have $f(\overline{y} - x) < f(\overline{y} - x_j)$. Then (C2) implies

$$\mu_T(j|\overline{y}) \ =\ \frac{p_j \sum_x f(\overline{y} - x)\pi_j(x)}{\sum_k p_k \sum_x f(\overline{y} - x)\pi_k(x)} \ \leqslant\ \frac{p_j}{\sum_k p_k \sum_x \frac{f(\overline{y}-x)}{f(\overline{y}-x_j)}\pi_k(x)} \ \leqslant\ \frac{p_j}{p_n \frac{f(\overline{y}-x_n)}{f(\overline{y}-x_j)}\pi_n(x_n)} \to 0$$

as $\overline{y} \to \infty$, which implies that $\mu_T(n|\overline{y})$ goes to one as the cutoff increases. In words, when the policies of the politicians are ordered by type, high realizations of the outcome become arbitrarily strong evidence that the incumbent is the best possible type. Similarly, $\mu_T(1|\overline{y})$ goes to one as $\overline{y}$ decreases without bound. Thus, the left-hand side of (7) approaches $\mathbb{E}[u(y)|\hat{x}_n]$ when the cutoff is large, and it approaches $\mathbb{E}[u(y)|\hat{x}_1]$ when the cutoff is small, and existence of a solution follows from the intermediate value theorem. Uniqueness follows from the fact that the left-hand side is strictly increasing in $\overline{y}$, from Lemma A.6 of Banks and Sundaram (1998). Standard continuity arguments imply that $y^*(\pi_1, \ldots, \pi_n)$ is continuous in its arguments.

To obtain the bound on the cutoff, consider any $\overline{y} > \max\{x : \pi_n(x) > 0\} + \hat{z}$. Recall that the posterior probability that the politician is type j, conditional on observing $\overline{y}$, is

$$\mu_T(j|\overline{y}) \quad = \quad \frac{p_j \sum_x f(\overline{y} - x)\pi_j(x)}{\sum_k p_k \sum_x f(\overline{y} - x)\pi_k(x)}.$$

Note that for all $k > j$ and all policies x_j with $\pi_j(x_j) > 0$ and x_k with $\pi_k(x_k) > 0$, we have $\hat{z} < \overline{y} - x_k < \overline{y} - x_j$. Since $f(\cdot)$ is single-peaked by (C1), we see that for all $x_1, \ldots, x_n$ such that each x_k is in the support of π_k, we have

$$f(\overline{y} - x_1) \ <\ f(\overline{y} - x_2) \ <\ \cdots \ <\ f(\overline{y} - x_n).$$

Therefore, in calculating posterior beliefs, the coefficients on prior beliefs are ordered by type, i.e.,

$$\frac{\sum_x f(\bar{y}-x)\pi_1(x)}{\sum_k p_k \sum_x f(\bar{y}-x)\pi_k(x)} \; < \; \cdots < \; \frac{\sum_x f(\bar{y}-x)\pi_n(x)}{\sum_k p_k \sum_x f(\bar{y}-x_k)\pi_k(x)},$$

and we conclude that the posterior distribution $\mu_T(\cdot|\bar{y})$ first order stochastically dominates the prior, contradicting the indifference condition. An analogous argument derives a contradiction for the case $\bar{y} < \min\{x : \pi_1(x) > 0\} + \hat{z}$, as desired.

To see the structure of $y^*(\pi_1,\ldots,\pi_n)$ for the special case of two types using pure policy strategies, the voters' cutoff is simply the solution to $\mu_T(2|y) = p_2$, i.e., the cutoff is such that conditional on the cutoff, the probability the incumbent is the high type is just equal to the prior probability. Letting x_1 and x_2 be the policies chosen by the two types, this means that $y^*(x_1,x_2)$ solves the equation

$$p_2 \quad = \quad \frac{p_2 f(y-x_2)}{p_1 f(y-x_1) + p_2 f(y-x_2)},$$

or after manipulating, it means that the likelihood of y is the same given the policy choices of the politicians, i.e., $f(y-x_1) = f(y-x_2)$. Adding the assumption that the density $f(\cdot)$ is standard normal, the cutoff is simply the midpoint of the politicians' choices, i.e.,

$$y^*(x_1,x_2) \quad = \quad \frac{x_1+x_2}{2}.$$

Indeed, this characterization as the midpoint of policy choices extends to any density that is symmetric around zero.

The preceding observations allow us to graphically depict a pure strategy electoral equilibrium for the case of two types. In Figure 7, we draw the indifference curves of U_1 and U_2 through the unique optimal policies, x_1^* and x_2^*, of the politician types given the constraint set determined by the cutoff y^*. This is reflected in the tangency condition at each optimal policy. Moreover, the voters' indifference condition implies that the likelihood of outcome y^* is equal given either optimal policy, as reflected in the equal slopes of the two tangent lines, or equivalently, the fact that y^* is the midpoint between x_1^* and x_2^*. Note that when the office benefit β increases, ceteris paribus, the indifference curves of the politician types become flatter, and optimal policies will move to the right, suggesting that higher office benefit leads to greater policy responsiveness. In equilibrium, however, the cutoff itself moves to the right, which shifts the politicians' constraint out to the right. The overall effect on policy choices is ambiguous, but the policy choice of at least one politician type must increase.

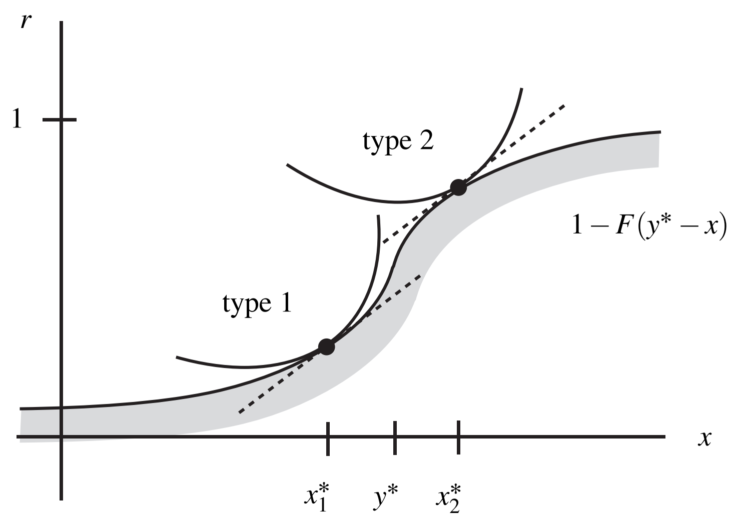

Figure 7: Electoral equilibrium

The next result establishes existence of electoral equilibrium in the two-period moral hazard model, along with a minimal characterization of equilibria. We see that even in the two-period model, where second-period policies are pinned down by end-game effects, electoral equilibria must solve a complicated fixed point problem: optimal policy choices of politician types depend on the cutoff used by voters, and the cutoff used by voters depends, via Bayes rule, on the policy choices of politician types. The proof relies on supermodularity of politician strategies, but the domain of mixed policy strategies with supports that are ordered by type is not convex. Thus, we use Proposition 3.4 and a reformulation of the domain to obtain convexity; the cost of doing so is that the fixed point correspondence is not defined simply by mixtures over optimal policy choices for the politicians, and the correspondence does not have convex values. Duggan and Martinelli (2017) circumvent this problem by establishing that it has contractible values and by applying the Eilenberg-Montomery fixed point theorem.[16]

Proposition 3.7 *In the two-period model of moral hazard with rent-seeking, assume (C1)–(C4). Then there is an electoral equilibrium, and in every electoral*

[16]See Duggan (2016) for a discussion of convexity issues that arise with respect to the domain of mixed policies with supports ordered according to type, and for a less direct proof of existence of equilibrium using Glicksberg's fixed point theorem in the infinite-horizon model with a two-period term limit. See Reny (2011) for another application of the Eilenberg-Montgomery fixed point theorem in the economics literature on auctions.

equilibrium, there exist mixed policy strategies $\pi_1^*,\ldots,\pi_n^*$ and a finite cutoff y^* such that:

(i) each type j politician mixes over policies using π_j^*, which places positive probability on at most two policies, say x_j^* and $x_{*,j}$, where $\hat{x}_j < x_{*,j} \leqslant x_j^*$,

(ii) the supports of policy strategies are strictly ordered by type, i.e., for all $j<n$, we have $x_j^* < x_{*,j+1}$,

(iii) voters re-elect the incumbent if and only if $y \geqslant y^*$, where the cutoff lies between the extreme policies shifted by the mode of the outcome density, i.e., $x_{*,1}+\hat{z} \leqslant y^* \leqslant x_n^* + \hat{z}$.

In proving the proposition, we must address three technical subtleties. The first is that when supports of mixed policy choices are only weakly ordered, the left-hand side of (7) is only weakly increasing, so that the equality has a closed, convex (not necessarily singleton) set of solutions. In fact, if all politician types choose the same policy with probability one, then updating does not occur and incumbents are always re-elected, so that the voter's cutoff is negatively infinite. As policy choices of politician types converge to the same policy, this means that the cutoff either jumps discontinuously (from a bounded, finite level) or diverges to negative infinity. We circumvent this problem by deriving a positive lower bound on the distance between optimal policy choices of the different types. Indeed, we first observe that equilibrium policy choices are bounded above by any choice $\overline{x}$ such that $\mathbb{E}[u(y)|\hat{x}_n] > w_n(\overline{x})+\beta$, i.e., $-w_n(\overline{x}) > \beta - \mathbb{E}[u(y)|\hat{x}_n]$. That is, if the type n politician prefers to choose her ideal policy with no chance of re-election rather than choose $\overline{x}$ and win with certainty, then no policy above $\overline{x}$ can be optimal for any type given any cutoff.

Next, given any cutoff $\overline{y}$ and any type j politician, there are at most two optimal policies, by Proposition 3.4, and each satisfies the first order condition (6). Note that $f(\overline{y}-x) \to 0$ uniformly on $[0,\overline{x}]$ as $|\overline{y}| \to \infty$, and from the first order condition, this implies that the optimal policies of the type j politician converge to the ideal policy, i.e., $x_j^*(\overline{y}) \to \hat{x}_j$ and $x_{*,j}(\overline{y}) \to \hat{x}_j$. Thus, we can choose a sufficiently large interval $[y_L,y_H]$ and $\varepsilon'>0$ such that for all $\overline{y}$ outside the interval, optimal policies differ across types by at least ε', i.e., for all $j<n$, we have $|x_{*,j+1}(\overline{y}) - x_j^*(\overline{y})| > \varepsilon'$. By upper semi-continuity of $x_j^*(\cdot)$ and lower semi-continuity of $x_{*,j+1}(\cdot)$, the function $|x_{*,j+1}(\overline{y}) - x_j^*(\overline{y})|$ attains its minimum on $[y_L,y_H]$, and this minimum is positive. Thus, there exists $\varepsilon''>0$ such that for all $\overline{y} \in [y_L,y_H]$, optimal policies differ by at least ε''. Finally, we set $\varepsilon = \min\{\varepsilon',\varepsilon''\}$ to establish the desired lower bound.

We are interested in the profiles $(\pi_1,\ldots,\pi_n)$ such that for all politician types j, π_j places positive probability on at most two alternatives, and the supports of mixed policy strategies are strictly ordered by type and separated by a distance of at least ε, i.e., for all $j < n$ and all policies x_j with $\pi_j(x_j) > 0$ and x_{j+1} with $\pi_{j+1}(x_{j+1}) > 0$, we have $x_j + \varepsilon \leqslant x_{j+1}$. It is convenient to represent such a profile by a $3n$-tuple (x,z,r), where $x = (x_1,\ldots,x_n) \in [0,\overline{x}]^n$, $z = (z_1,\ldots,z_n) \in [0,\overline{x}]^n$, and $r = (r_1,\ldots,r_n) \in [0,1]^n$. In addition, we require that for all j, we have $x_j \leqslant z_j$, and that for all $j < n$, we have $z_j + \varepsilon \leqslant x_{j+1}$. We then associate (x,z,r) with the profile of mixed policy strategies such that the type j politician places probability r_j on x_j and the remaining probability $1-r_j$ on z_j. Letting D^ε consist of all such $3n$-tuples (x,z,r), we see that D^ε is nonempty, convex, and compact. Using this representation, we can define (abusing notation slightly) the induced cutoff $y^*(x,z,r)$, which is continuous as a function of its arguments.

The second difficulty is that the set Y of policy outcomes is not compact, so that the voter's cutoff is, in principle, unbounded. To circumvent this problem, we note that by continuity of the function $y^*(\cdot)$ the image $y^*(D^\varepsilon)$ is compact, and we can let $\overline{Y}$ be a convex, compact set containing this image. The existence proof then proceeds by applying a generalization of Kakutani's fixed point theorem to an appropriately defined correspondence. We define the correspondence $\Phi\colon D^\varepsilon \times \overline{Y} \rightrightarrows D^\varepsilon \times \overline{Y}$ so that for each $(x,z,r,\overline{y})$, the value of Φ consists of $(3n+1)$-tuples $(\tilde{x},\tilde{z},\tilde{r},\tilde{y})$ such that for every politician type j, the mixed policy strategy represented by $(\tilde{x}_j,\tilde{z}_j,\tilde{r}_j)$ is optimal and $\tilde{y}$ is the unique cutoff induced by the indifference condition:

$$\Phi(x,z,r,\overline{y}) \quad = \quad \left\{ (\tilde{x},\tilde{z},\tilde{r},\tilde{y}) \in D^\varepsilon \times \overline{Y} \;\middle|\; \begin{array}{c} \text{for all } j, \tilde{x}_j \leqslant \tilde{z}_j, \\ \tilde{r}_j > 0 \Rightarrow \tilde{x}_j \in \{x_{*,j}(\overline{y}), x_j^*(\overline{y})\}, \\ \tilde{r}_j < 1 \Rightarrow \tilde{z}_j \in \{x_{*,j}(\overline{y}), x_j^*(\overline{y})\}, \\ \text{and } \tilde{y} = y^*(x,z,r) \end{array} \right\}.$$

Of note, we require that the first policy coordinate $\tilde{x}_j$ is less than or equal to the second, $\tilde{z}_j$, and we require that these are optimal when chosen with positive probability.

To deduce the existence of a fixed point of Φ, we first verify that the correspondence is upper hemi-continuous with closed values, i.e., it has closed graph. This property is not immediately obvious, because optimal policies are not unique, and the functions $x_j^*(\cdot)$ and $x_{*,j}(\cdot)$ are not continuous. It is important that we allow for the possibility that $\tilde{x}_j = \tilde{z}_j$, in which case both policies are equal to either the least optimal policy $x_{*,j}(\overline{y})$ or to the greatest optimal policy $x_j^*(\overline{y})$. Of course, these policies can coincide as well. Let $\{(x^m,z^m,r^m,\overline{y}^m)\}$ be any sequence converging to $(x,z,r,\overline{y})$ in $D^\varepsilon \times \overline{Y}$, and consider a corresponding sequence $\{(\tilde{x}^m,\tilde{z}^m,\tilde{r}^m,\tilde{y}^m)\}$ such that $(\tilde{x}^m,\tilde{z}^m,\tilde{r}^m,\tilde{y}^m)$ belongs to $\Phi(x^m,z^m,r^m,\overline{y}^m)$ for all m and $(\tilde{x}^m,\tilde{z}^m,\tilde{r}^m,\tilde{y}^m) \to (\tilde{x},\tilde{z},\tilde{r},\tilde{y})$. We must show that $(\tilde{x},\tilde{z},\tilde{r},\tilde{y}) \in \Phi(x,z,r,\overline{y})$. Since limits preserve weak inequalities, it is immediate that for all j, we have $\tilde{x}_j \leqslant \tilde{z}_j$, and continuity of $y^*(\cdot)$

implies $\tilde{y} = y^*(x,z,r)$. It remains to establish optimality of policies, and we consider the first restriction, as the argument for the second is analogous. To this end, suppose $\tilde{r}_j > 0$, so that for sufficiently high m, we also have $\tilde{r}^m_j > 0$, implying $\tilde{x}^m_j \in \{x_{*,j}(\overline{y}^m), x^*_j(\overline{y}^m)\}$. By the theorem of the maximum (see Border (1985), Theorem 12.1), $\tilde{x}_j$ is an optimal policy for the type j politician given cutoff $\overline{y}$. If $\tilde{x}_j \notin \{x_{*,j}(\overline{y}), x^*_j(\overline{y})\}$, then this implies the politician has at least three optimal policies, contradicting Proposition 3.4. Thus, $\tilde{x}_j$ is either the least or greatest optimal policy given $\overline{y}$, as desired.

This formulation yields a correspondence that is defined on a convex and compact domain and that is upper hemi-continuous and has nonempty, closed values. The typical application of Kakutani's fixed point theorem also proceeds by verifying convex values of the correspondence, and this leads to the third difficulty: Φ does not have this property. Nevertheless, the values of the correspondence are contractible, and this is sufficient for existence of a fixed point. A subset $C \subseteq \Re^d$ of Euclidean space is *contractible* if there is an element $\overline{c} \in C$ and a continuous mapping $h\colon C \times [0,1] \to C$ such that for all $c \in C$, $h(c,0) = c$ and $h(c,1) = \overline{c}$. That is, the set can be continuously deformed to a single element. Convex sets are contractible, but convexity is not necessary for contractibility. It is straightforward to see that $\Phi(x,z,r,\overline{y})$ is contractible to the element $(\hat{x},\hat{z},\hat{r},\hat{y})$ such that: for all j,

- $\hat{x}_j = x_{*,j}(\overline{y})$,
- $\hat{z}_j = x^*_j(\overline{y})$,
- $\hat{r}_j = 1$,

where of course $\hat{y} = y^*(x,z,r)$ is fixed by construction. To reduce notation, we provide an informal description of the mapping h, breaking the unit interval into five components. Consider any $(\tilde{x},\tilde{z},\tilde{r},\tilde{y}) \in \Phi(x,z,r,\overline{y})$. For $t \in [0,.2]$, we continuously adjust each $\tilde{r}_j$ by dropping these values to zero. For $t \in (.2,.4]$, we continuously adjust each $\tilde{x}_j$ to $x_{*,j}(\overline{y})$. This adjustment requires that $\tilde{x}_j$ take sub-optimal values, but because the probability on $\tilde{x}_j$ is zero, we remain in the value of the correspondence. For $t \in (.4,.6]$, we continuously adjust each $\tilde{r}_j$ to one. For $t \in (.6,.8]$, we continuously adjust each $\tilde{z}_j$ to $x^*_j(\overline{y})$. And for $t \in (.8,1]$, we continuously adjust each $\tilde{r}_j$ to one. This completes the construction, and we conclude that the values of Φ are contractible.

The correspondence Φ is upper hemi-continuous with nonempty, closed, contractible values, and the domain $D^\varepsilon \times \overline{Y}$ is nonempty, compact, and convex. Therefore, the Eilenberg-Montgomery fixed point theorem (see Border (1985), Theorem 15.9) implies that Φ has a fixed point,[17] (x^*,z^*,r^*,y^*), which yields an electoral

[17] The Eilenberg-Montgomery fixed point theorem holds for a domain that is a general acyclic

equilibrium. Finally, the characterization results in (i)–(iii) follow directly from Propositions 3.4–3.6, completing the proof.

We have not yet touched on the possibility for responsive democracy in the two-period moral hazard model with rent-seeking, where in the present context, we interpret responsive democracy to mean that office holders choose high levels of policy, despite short run incentives to choose their ideal policy. Given the short time horizon (and limited ability of voters to sanction politicians), and given the divergence in preferences between voters and politicians, the prospects for a well-functioning political system may seem dim. Nevertheless, when β is large, so that politicians are substantially office-motivated, we obtain a form of the responsive democracy result. We make use of an additional Inada-type condition: for all j,

$$\text{(C5)} \qquad \lim_{x\to\infty} w'_j(x) \;=\; -\infty.$$

Let $G = \{j : \mathbb{E}[u(y)|\hat{x}_j] > V^C\}$ denote the set of above average types, which are such that the expected utility from their ideal policy exceeds the expected utility from a challenger. Let $\ell = \min G$ be the smallest above average type.

The next result, due to Duggan and Martinelli (2017), provides a characterization of equilibria when office benefit is high. We find that voters become arbitrarily demanding, in the sense that the equilibrium cutoff diverges to infinity, that the policy choices of all politician types become close to their ideal policy or arbitrarily large, and that all above average types exert unbounded effort. An immediate implication, since type n is above average and $p_n > 0$, is that if voter utility is not bounded above, i.e., $\lim_{y\to\infty} u(y) = \infty$, then the voters' expected utility from politicians' choices in the first period increases without bound as office benefit becomes large, i.e.,

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j(x) \quad\to\quad \infty.$$

It is possible that some politician type mixes between a policy that is close to the ideal policy and another that becomes arbitrarily large, but because policy choices are ordered by type, an implication of the proposition is that this can obtain for at most one politician type; choices of lower types will converge to their ideal policies, while choices of higher types will diverge to infinity. Note that the Inada condition (C5) is used only to prove part (iii) of the result.

Proposition 3.8 *In the two-period model of moral hazard with rent-seeking and moral hazard, assume (C1)–(C5), and let the office benefit* β *become arbitrarily*

absolute neighborhood retract. Every compact, convex subset of Euclidean space is an acyclic ANR, so this assumption is satisfied automatically.

large. Then for every selection of electoral equilibria σ, the voters' cutoff diverges to infinity; for each politician type j, the policy choices of all above average types increase without bound; and the greatest policy choice of other types either increases without bound or accumulates at the ideal policy:

(i) $y^* \to \infty$,

(ii) for all j, all $\varepsilon > 0$, and sufficiently large β, we have $\{x_{*,j}, x_j^*\} \subseteq (\hat{x}_j, \hat{x}_j + \varepsilon) \cup (\frac{1}{\varepsilon}, \infty)$,

(iii) $x^*_{\ell-1} \to \infty$, *and thus for all $j \geqslant \ell$, we have $x_{*,j} \to \infty$.*

To prove the result, let β be large, and let σ be an electoral equilibrium. By Proposition 3.7, each politician type j mixes between two policies, x_j^* and $x_{*,j}$, and voters use a cutoff y^*. Suppose there is a subsequence such that y^* is bounded above, say $y^* \leqslant \bar{y}$. By Proposition 3.7, the equilibrium cutoff lies in the compact set $[\hat{x}_1, \bar{y}]$. Then the first order condition for the type 1 politician in (6) implies that $x_{*,1} \to \infty$, and in particular, we have $\bar{y} < x_{*,1}$ for large enough β, but this contradicts $x_{*,1} + \hat{z} \leqslant y^* \leqslant x_n^* + \hat{z}$. We conclude that y^* diverges to infinity, which proves (i).

To prove (ii), suppose there is a type j, an $\varepsilon > 0$, and a subsequence of office benefit levels such that $\hat{x}_j + \varepsilon \leqslant x_j^* \leqslant \frac{1}{\varepsilon}$. Going to a subsequence, we can assume $x_j^* \to \tilde{x}_j$ such that $\hat{x}_j < \tilde{x}_j < \infty$. Then for sufficiently large β, we have $\hat{x}_j < x_j^*$. For these parameters, the current gain to the type j politician from choosing $\hat{x}_j$ instead of x_j^* is non-positive, and thus we note that

$$(F(y^* - \hat{x}_j) - F(y^* - x_j^*))[w_j(\hat{x}_j) + \beta - V^C] \quad \geqslant \quad w_j(\hat{x}_j) - w_j(x_j^*).$$

That is, the current gains from choosing the ideal policy are offset by future losses. Since $y^* \to \infty$, the limit of

$$\frac{F(y^* - x_j^*) - F(y^* - \tilde{x}_j - 1)}{F(y^* - \hat{x}_j) - F(y^* - x_j^*)}$$

as β becomes large is indeterminate, and by L'Hôpital's rule, the limit is equal to

$$\lim \frac{f(y^* - x_j^*) - f(y^* - \tilde{x}_j - 1)}{f(y^* - \hat{x}_j) - f(y^* - x_j^*)} \; = \; \lim \frac{f(y^* - \tilde{x}_j - 1)\left(\frac{f(y^* - x_j^*)}{f(y^* - \tilde{x}_j - 1)} - 1\right)}{f(y^* - x_j^*)\left(\frac{f(y^* - \hat{x}_j)}{f(y^* - x_j^*)} - 1\right)} \; = \; \infty,$$

where we use (C1) and (C2). Then, however, the future gain from choosing $\tilde{x}_j + 1$ instead of x_j^* strictly exceeds current losses, i.e.,

$$(F(y^* - x_j^*) - F(y^* - \tilde{x}_j - 1))[w_j(\hat{x}_j) + \beta - V^C] \quad > \quad w_j(x_j^*) - w_j(\tilde{x}_j + 1), \quad (8)$$

for high enough β. To be specific, let

$$\begin{aligned} A &= w_j(\hat{x}_j) + \beta - V^C \\ B &= w_j(\hat{x}_j) - w_j(x_j^*) \\ C &= w_j(x_j^*) - w_j(\tilde{x}_j + 1), \end{aligned}$$

where A is evaluated at sufficiently large β. Note that since $\hat{x}_j < \tilde{x}_j < \infty$, we have $\lim B > 0$ and $\lim C < \infty$. We have noted that $(F(y^* - \hat{x}_j) - F(y^* - x_j^*))A \geqslant B$ for sufficiently large β, and we have shown that as β becomes large, we have

$$\frac{F(y^* - x_j^*) - F(y^* - \tilde{x} - 1)}{F(y^* - \hat{x}_j) - F(y^* - x_j^*)} \quad > \quad \frac{C}{B}.$$

Combining these facts, we have

$$(F(y^* - \hat{x}_j) - F(y^* - x_j^*))A\left(\frac{F(y^* - x_j^*) - F(y^* - \tilde{x} - 1)}{F(y^* - \hat{x}_j) - F(y^* - x_j^*)}\right) \quad > \quad B\left(\frac{C}{B}\right),$$

which yields (8) for large β. This gives the type j politician a profitable deviation from x_j^*, a contradiction. A similar argument holds for $x_{*,j}$, and (ii) follows.

Finally, to prove (iii), suppose that $x_{\ell-1}^*$ does not diverge to infinity. By (ii), there is a subsequence such that $x_{\ell-1}^* \to \hat{x}_{\ell-1}$. Now fix politician type $j \leqslant \ell - 1$, and note that since equilibrium policy choices are ordered by type, we have $x_j^* \to \hat{x}_j$. Using the expression for Bayes rule, the posterior probability of type j conditional on observing y^* satisfies

$$\mu_T(j|y^*) \quad = \quad \frac{p_j \sum_x f(y^* - x)\pi_j(x)}{\sum_k p_k \sum_x f(y^* - x)\pi_k(x)} \quad \leqslant \quad \frac{p_j f(y^* - x_j^*)}{\sum_{k \geqslant \ell} p_k \sum_x f(y^* - x)\pi_k(x)},$$

where the inequality uses $y^* \to \infty$ and single-peakedness of $f(\cdot)$. Note that

$$\sum_{k \geqslant \ell} p_k \sum_x f(y^* - x)\pi_k(x) \quad = \quad \sum_{k \geqslant \ell} p_k[f(y^* - x_k^*)\pi_k(x_k^*) + f(y^* - x_{*,k})\pi_k(x_{*,k})].$$

Dividing by $f(y^* - x_j^*)$, we obtain the expression

$$\sum_{k \geqslant \ell} p_k \left[\frac{f(y^* - x_k^*)}{f(y^* - x_j^*)}\pi_k(x_k^*) + \frac{f(y^* - x_{*,k})}{f(y^* - x_j^*)}\pi_k(x_{*,k})\right].$$

By the MLRP, we have $\frac{f(y^* - x_{*,k})}{f(y^* - x_j^*)} \to \infty$ for all $k \geqslant \ell$. Similarly, if $x_k^* \to \hat{x}_k$, then we have $\frac{f(y^* - x_k^*)}{f(y^* - x_j^*)} \to \infty$. By (ii), the remaining case is $x_k^* \to \infty$. Note that in this case, (C5) implies $w'_k(x^*_k) \to -\infty$, and thus the first order condition in (6) implies that $f(y^* - x^*_k)\beta \to \infty$. The first order condition for type j implies $f(y^* - x^*_j)\beta \to 0$, and we infer that $\frac{f(y^*-x^*_k)}{f(y^*-x^*_j)} \to \infty$. Thus, we have

$$\mu_T(j|y^*) \;\leqslant\; \frac{p_j}{\sum_{k \geqslant \ell} p_k \sum_x \frac{f(y^*-x)}{f(y^*-x^*_j)} \pi_k(x)} \;\to\; 0.$$

We conclude that the voters' posterior beliefs conditional on y^* place probability arbitrarily close to one on above average types $j \geqslant \ell$, contradicting the indifference condition in (7). Therefore, we have $x^*_{\ell-1} \to \infty$, and since policy choices are ordered by type, this implies that $x_{*,j} \to \infty$ for all $j \geqslant \ell$. This establishes (iii), as desired.

Thus, high office benefit sets up a rat race between the above average types, who chase after an increasingly demanding reelection threshold; in fact, Duggan and Martinelli (2017) show that under stronger preference restrictions, still generalizing quadratic utility, *all* politician types compete in the rate race, with the type 1 politician shirking with probability approaching zero in order to maintain the voters' indifference conditional on the cutoff being realized. Moreover, the probability that the incumbent is re-elected converges to one, demonstrating a stark form of incumbency advantage. With high office benefit, it is as if voters were facing two types of politicians; a "good" type interested in reelection, and a "bad" type mixing between shirking and participating in the rat race, but the probabilities of these two types are of course endogenous.

We formulate the rent-seeking environment so that the policy space $X = \mathbb{R}_+$ is unbounded above, and this has facilitated the development by removing the possibility of corner solutions and permitting a first order analysis. Of course, given an equilibrium with policy choices $x^*_1, \ldots, x^*_n$, the equilibrium survives if we modify the model by imposing an upper bound $\overline{x} \geqslant x^*_n$ on feasible policies. But imposing such a bound a priori, and independent of the office benefit, creates serious technical difficulties stemming from the possibility that all types pool at the upper bound, which implies that voters do not revise their prior beliefs after observing the outcome y. Specifically, this difficulty arises when the office benefit is large and office holders have strong incentives to exert higher effort to improve their chances of re-election. The problem is illustrated in Figure 8, where we suppose there are just two types. When the voters' cutoff $\overline{y}$ is large, the optimal policies of the politicians will be close to their ideal policies. Then the induced cutoff defined by the voters' indifference condition will be roughly the midpoint between the optimal policies. As we decrease the cutoff $\overline{y}$, the optimal policies increase, and so does the induced cutoff $\overline{y}'$, and an equilibrium is obtained when $\overline{y} = \overline{y}'$; the mapping from $\overline{y}$ to $\overline{y}'$ is represented by the kinked line in the figure.

At some point, the type 2 politicians' optimal policy may hit the upper bound $\overline{x}$, as indicated by the kink; it is possible, moreover, that the optimal policy of the type 1 politician also hits the upper bound *before the induced cutoff crosses the* 45° *line.* At that point, both types are at a corner solution, and voters don't update—this is indicated by the dark vertical line in the figure. Then induced cutoff jumps to negative infinity, and deferential voting implies that the incumbent is always re-elected, but then politicians have no incentive to exert effort in the first period. In this case, the assumption of deferential voting must be dropped, and the voters' cutoff $\overline{y}$ must be set to induce both politician types to choose $\overline{x}$. That is, in equilibrium, politician types pool at the upper bound $\overline{x}$, and voters, being indifferent between the incumbent and challenger, uses a finite cutoff $\overline{y}$ consistent with optimality of $\overline{x}$.

A problematic feature of the above construction is that there will be a continuum of cutoffs that can be used to support pooling at the upper bound. For example, given any predetermined value $\overline{y}$, we can choose β sufficiently large that the first order condition

$$w'_j(\overline{x}) \quad > \quad -f(\overline{y}-\overline{x})[w_j(\hat{x}_j)+\beta-V^C]$$

holds for both types $j=1,2$. Of course, this inequality holds in an open set around $\overline{y}$. Since voters do not revise their prior beliefs about the incumbent, they are indifferent between the candidates, and all cutoffs in that open set are optimal. Thus, this construction generates a continuum of pooling equilibria. We find it natural, especially when considering the case of highly office motivated politicians, to maintain the concept of electoral equilibrium and to allow unbounded policy choices, thereby avoiding these difficulties.

The main starting point of our analysis has been the statement of equilibrium existence in Proposition 3.7, which must address challenges due to non-convexities in the payoffs of politicians. This is a problem in previous work too, but one that has been neutralized by different modeling assumptions or focussing on necessary first order conditions. The existence problem carries over, in a simplified form, to the symmetric learning environment of Ashworth (2005), who assumes that office benefit is not too large in order to guarantee existence in pure strategies. In the models of Besley (2006), the politician's choice is either explicitly between two possible policies, or it reduces to a finite number of policies, so that equilibria in mixed strategies exist. In the informational setup of Fearon (1999), voters observe a stochastic level $u(x)+\varepsilon$ of utility, rather than a noisy policy outcome $y=x+\varepsilon$, allowing him to solve first order conditions explicitly and to verify second order conditions for a local equilibrium; when office benefit is high, however, these local equilibria admit non-local deviations that are profitable. In contrast, we assume

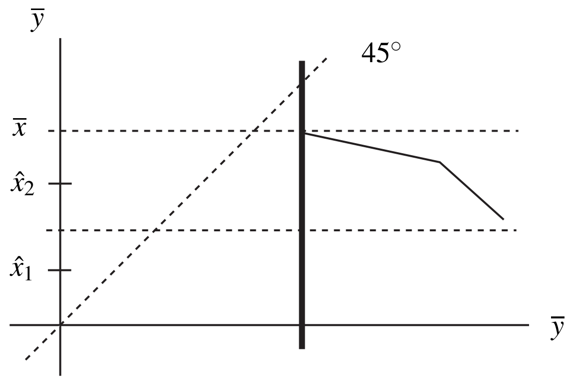

Figure 8: Existence problems with bounded policy

that voters observe a policy outcome y drawn from the conditional density $f(\cdot|x)$ and that they accrue utility $u(y)$ from that outcome; we allow politicians to choose from a continuum of possible policies; and we capture an arbitrarily large office benefit by allowing politicians to mix but imposing (C4) to limit mixing to at most two policy choices.

4 Dynamic framework

We now imbed the electoral model in an infinite-horizon dynamic setting such that in each period $t = 1, 2, \ldots$, an incumbent politician makes an unobservable policy choice x_t, this choice determines a publicly observed policy outcome y_t, a challenger is randomly drawn, and an election is held. Consistent with the citizen-candidate approach, we assume that the incumbent's choice is unrestricted and that the challenger cannot make binding campaign promises, and we therefore suppress political campaigns. Voters do not directly observe the type of the incumbent or challenger, but rather they update their beliefs about the incumbent on the basis of observed policy outcomes. Thus, the framework is the direct extension of the two-period model summarized in Figure 3, with the fundamental difference that there is no longer a last period.

As in the two-period model, there is a nonempty, convex action set $X \subseteq \mathbb{R}$ and a continuum N of citizens who are partitioned into a finite set of types $T = \{1, \ldots, n\}$, $n \geqslant 2$. In period 1, a politician is randomly drawn from the population of citizens, with each type j having probability p_j, the politician makes policy choice $x_1 \in X$, and voters observe policy outcome y_1 realized from the distribution $F(\cdot|x_1)$. Every period t thereafter, timing is as follows:

- the period t incumbent, i.e., the winner of the period $t-1$ election, chooses policy $x_t \in X$,
- voters observe policy outcome y_t drawn from the distribution $F(\cdot|x_t)$,
- a challenger is drawn from the electorate with each type j having probability p_j, and an election is held,
- and we move to period $t+1$ and repeat the above process.

Furthermore, we assume that challengers are drawn without replacement (so that once rejected, a politician does not run for office again) and independently from previous candidates. If an incumbent first assumes office in period t and is in term ℓ, then we refer to policy outcomes $\mathbf{y}^{\ell-1} = (y_t, y_{t+1}, \ldots, y_{t+\ell-1})$ generated by her choices as the politician's *personal history.* We let $\mathbf{y}^0$ be the "empty" history indicating a newly elected politician. In the following sections, we consider two specifications of the outcome distribution. First, we let $F(\cdot|x_t)$ be degenerate on $y_t = x_t$, so that politicians choose policy directly and these choices are essentially observable, i.e., monitoring is perfect; second, we let the distribution be continuous with full support, so that monitoring is imperfect.

Type j citizens receive a payoff $u_j(y_t)$ from policy y_t in period t if they are not in office and a payoff of $w_j(x_t) + \beta$ in period t if they hold office. Citizens have a common rate of time discounting, which is represented by the discount factor $\delta \in [0,1)$. Given a sequence $x_1, x_2, \ldots$ of actions and a sequence $y_1, y_2, \ldots$ of policies, the total payoff of a type j citizen is the discounted sum of per period payoffs,

$$\sum_{t=1}^{\infty} \delta^{t-1}[I_t(w_j(x_t) + \beta) + (1 - I_t)u_j(y_t)],$$

where $I_t \in \{0,1\}$ is an indicator variable that takes a value of one if the citizen holds office in period t and takes a value of zero otherwise.

We continue to focus on the *spatial preferences* environment in the context of pure adverse selection, i.e., perfect monitoring, and we focus on the *rent-seeking* environment in the presence of moral hazard, i.e., imperfect monitoring. Strategies are now potentially highly complex, as policy choices and votes could conceivably depend arbitrarily on observed histories of policy and electoral outcomes. To reduce the multiplicity of perfect Bayesian equilibria of the model, we must impose refinements that strengthen sequential rationality and Bayesian updating to preclude implausible behavior by voters and politicians. We extend the concept of electoral equilibrium from the two-period model to "stationary electoral equilibria" of the infinite-horizon model. These equilibria appear simple enough to be behaviorally plausible, and the refinement adequately reduces the equilibrium set to produce predictive power and permit comparative statics; and they can often be characterized by a finite system of equations, facilitating analytical and numerical computation. Importantly, the concept of stationary electoral equilibrium synthesizes many approaches taken in the existing literature on dynamic elections.

In contrast to the two-period model, which features a single election, strategies must at a minimum now be conditioned on information generated by the incumbent politician's past choices in office, which may reveal relevant information to voters. A *strategy for a type j politician* is a sequence $(\pi_j^\ell)_{\ell=1}^\infty$, where π_j^ℓ specifies the politician's mixture over policy choices in term ℓ of office as a function of personal history; we write $\pi_j^\ell(\mathbf{y}^{\ell-1}) \in \Delta(X)$ for the mixture over policy choices in term ℓ given personal history $\mathbf{y}^{\ell-1}$ over the first $\ell-1$ terms of office, so that that the type j politician's policy choice in term ℓ after policy outcomes $\mathbf{y}^{\ell-1} = (y_t,\dots,y_{t+\ell-1})$ is realized from the distribution $\pi_j^\ell(\mathbf{y}^{\ell-1})$. We write $\pi_j^\ell(Z|\mathbf{y}^{\ell-1})$ for the probability that a type j politician chooses a policy in the (measurable) set Z, and for tractability, we assume $\pi_j^\ell(\mathbf{y}^{\ell-1})$ has finite support. An alternative interpretation, which we adopt in the analysis of elections with perfect monitoring, is that $\pi_j^\ell(Z|\mathbf{y}^{\ell-1})$ is the *proportion* of type j politicians who choose a policy in the set Z; then we understand that each politician uses a pure strategy, but we allow politicians of the same type to choose different policies. A *strategy for a type j voter* is a vector $(\rho_j^\ell)_{\ell=1}^\infty$, where $\rho^\ell(\mathbf{y}^{\ell-1}, y_{t+\ell}) \in [0,1]$ determines the probability that the voter votes to re-elect the incumbent as a function of the personal history $\mathbf{y}^{\ell-1}$ of the incumbent in prior terms of office and the current policy outcome $y_{t+\ell}$. And a *belief system* is a sequence $(\mu^\ell)_{\ell=1}^\infty$, where $\mu^\ell(\cdot|\mathbf{y}^{\ell-1}, y_{t+\ell})$ is a probability distribution on $T \times X$ as a function of the personal history of the incumbent and current policy outcome.[18] If the incumbent is re-elected, then the marginal of this distribution on T determines the voters' *updated prior belief* regarding the incumbent's type at the beginning of the next period. We omit notation for more complex strategies in which citizens condition on histories of previous office holders, as they seem less compelling on behavioral grounds.

A strategy profile $\sigma = ((\pi_j^\ell)_{\ell=1}^\infty, (\rho_j)_{\ell=1}^\infty)_{j\in T}$ is *sequentially rational* given belief system μ if for every term of office ℓ and every personal history $\mathbf{y}^{\ell-1}$, no politician can gain by deviating to a different policy choice, and for all policy outcomes $y_{t+\ell}$, voters of each type vote for the candidate that offers the highest expected discounted payoff conditional on her information; and beliefs μ are *consistent* with σ if for every term of office ℓ and every personal history $\mathbf{y}^{\ell-1}$ and outcome $y_{t+\ell}$ on the

[18] Even more general history dependence could be allowed, as policy choices and votes could conceivably depend on personal histories of previous politicians; this will not be relevant for the subsequent analysis.

path of play, $\mu^{\ell}(j,x|\mathbf{y}^{\ell-1},y_{t+\ell})$ is derived using Bayes rule.[19] A *perfect Bayesian equilibrium* is a pair $\psi = (\sigma,\mu)$ such that σ is sequentially rational given μ and such that μ is consistent with σ.

We focus on pairs $\psi = (\sigma,\mu)$ that are *stationary,* in the sense that (i) the choices of a politician depend only on her type and the voters' updated prior beliefs at the beginning of the current period, (ii) votes of voters depend only on the updated priors and the current policy outcome, (iii) the belief system depends on personal history only through the voters' updated priors and the current policy outcome, and (iv) these functional relationships are constant over time. This implies that the continuation value of a challenger $V_j^C(\psi)$ for a type j voter is constant over time, but it implies more. Consider a type j politician with personal history $\mathbf{y}^{\ell-1}$ such that the voters' updated prior beliefs are b, and consider another type j politician with personal history $\tilde{\mathbf{y}}^{\ell'-1}$ leading to the same updated prior beliefs b. An action x by either politician leads to the same distribution of policy outcomes on which voters condition their posterior beliefs; and in either scenario, if a policy outcome y on the path of play is observed, then Bayesian updating leads to the same posterior beliefs about the incumbent's type. Thus, the situations faced by the politicians with the same updated prior are strategically isomorphic, as are the situations of voters with the same updated prior and observed policy outcome, and we assume that the behavior of citizens reflects this isomorphism.

We let $b \in \Delta(T)$ denote the prior beliefs of the voters at the beginning of a period, and given stationary strategy profile σ, we abuse notation slightly by writing $\pi_j(b)$ for the mixture over actions of a type j politician given updated priors b (alternatively, $\pi_j(Z|b)$ is the fraction of type j politicians who choose a policy in Z given beliefs b); we write $\rho_j(b,y)$ for the vote of a type j voter with prior b after observing policy outcome y; and we write $\mu(b,y)$ for the updated beliefs conditional on observing y given prior beliefs b. We henceforth write $V_j^I(b|\psi)$ for the expected discounted payoff of a type j voter from re-electing an incumbent given updated prior beliefs b, and (again abusing notation) we write $V_j^I(b,y|\psi) = V_j^I(\mu_T(b,y)|\psi)$ for the expected discounted payoff from the incumbent given prior b and observed outcome y.

A pair $\psi = (\sigma,\mu)$ is *deferential* if voters favor the incumbent when indifferent, or more formally, given any beliefs b and policy outcome y on the path of play, a type j voter votes for the incumbent if and only if $V_j^I(b,y|\psi) \geqslant V_j^C(\psi)$. As in the dynamic Hotelling-Downs model of Subsection 2.4, these continuation values can be written as expected utilities with respect to two lotteries, L and L', over policies, and it follows that the median type m is pivotal in the election; thus, the incumbent

[19] Recall that we only consider equilibria in which the mixtures $\pi_j^{\ell}(\mathbf{y}^{\ell-1})$ have finite support, so that Bayesian updating is well-defined.

wins if and only if $V_m^I(b,y|\psi) \geqslant V_m^C(\psi)$. We say ψ is *monotonic* if for all voter types j and all updated priors b, there is some utility cutoff $\underline{u}_j(b)$ such that for all policy outcomes y, the type j voters vote to re-elect the incumbent if and only if the utility from y meets or exceeds that cutoff, i.e.,

$$\rho_j(b,y) \quad = \quad \begin{cases} 1 & \text{if } u_j(y) \geqslant \underline{u}_j(b), \\ 0 & \text{else.} \end{cases}$$

Using decisiveness of the median voter, we can define the *acceptance set* of policy outcomes that lead to re-election given updated prior b as

$$A(b|\psi) \quad = \quad \{y \in Y : V_m^I(b,y|\psi) \geqslant V_m^C(\psi)\},$$

and by monotonicity, this will be a closed, convex subset.

Our main equilibrium concept for the infinite-horizon model is defined as follows: we say $\psi = (\sigma,\mu)$ is a *stationary electoral equilibrium* if it is a perfect Bayesian equilibrium that is stationary, deferential, and monotonic.[20] We emphasize that this is a refinement of perfect Bayesian equilibrium, so that after all histories, no citizen can increase her expected discounted payoff by deviating to another strategy (stationary or non-stationary). And although we allow in principle for behavior as a general function of updated priors, the restrictions we impose capture some intuitive ideas. The assumption of deferential strategies is a form of *prospective voting,* in which a voter casts her vote as though pivotal in the election, and the assumption of monotonicity formalizes *retrospective voting,* in which a voter asks, "What have you done for me lately?" and votes to re-elect the incumbent if the policy outcome delivered by the politician satisfies a certain threshold. Thus, in a stationary electoral equilibrium, prospective and retrospective voting are compatible and both describe the behavior of voters, and the choices of office holders are optimal given these voting strategies. Note that although our equilibrium concept is stationary with respect to voters' beliefs about the incumbent's type, stationary electoral equilibria allow non-trivial dynamics, for once an incumbent is re-elected and voters update their beliefs, it is possible that that the voters' acceptance set and the politicians' optimal policy choices change; and if the incumbent is again re-elected, then updating may continue and play may continue to evolve.

In the rent-seeking environment, all voters share the same preferences, and so every voter type can serve as a representative voter. In the spatial preferences environment, voter preferences are heterogeneous, but given a stationary electoral equilibrium, the median voter type is always pivotal in the sense that the incumbent

[20] We also consider a further refinement of "simple electoral equilibrium" in the perfect monitoring case.

wins if and only if the expected discounted payoff from re-electing her weakly exceeds that from electing a challenger. The following *representative voter theorem for dynamic elections* records this observation formally and is used throughout the subsequent analysis.

Proposition 4.1 *In the infinite-horizon electoral model with spatial preferences, if $\psi = (\sigma, \mu)$ is a stationary electoral equilibrium, then the median voter type is a representative voter, i.e., the incumbent wins if and only if $V_m^I(b, y|\psi) \geqslant V_m^C(\psi)$.*

Before proceeding to impose specific informational assumptions on the model, we note a useful and general principle that unpacks the logical implications of the representative voter theorem. Formally, we simply assume that each voter casts her ballot as though pivotal in an election, calculating the expected discounted payoffs from the incumbent and challenger in a sophisticated way but voting sincerely. The representative voter result in Proposition 4.1 shows that the median voter type actually is pivotal in elections, i.e., that the calculations of a majority of voters will always agree with the median voter type. In this sense, the median voter type is "representative," but only in a passive way: the result is that given future behavior of politicians *and voters,* the median voter type prefers one candidate to another if and only if a majority of voters do. This is distinct from the assumption that there is a single, unitary voter, for in that case, the unitary voter should not take as given her own future behavior and optimize only between the current candidates; rather, a unitary voter would rationally optimize over all (possibly non-stationary) voting plans as a function of histories. In other words, a unitary voter faces an *optimal retention problem* with value function

$$\overline{V}(b) = \max_{\nu \in [0,1]} \nu \sum_j b_j \sum_x \left[\int_y [u_m(y) + \delta \overline{V}(\mu_T(x|b,y))] F(dy|x) \right] \pi_j(x|b) + (1-\nu)\overline{V}(p),$$

where the maximization is with respect to the probability ν of retaining the incumbent given the voter's beliefs b after observing the outcome of the incumbent's policy choice. Here, $\overline{V}(b)$ is the optimum value of electing a politician given the voter's beliefs b about the politician's type. We use the fact that the problem is stationary with respect to the voter's beliefs about the incumbent's type, so that without loss of generality, we can write this value as a function of the voter's beliefs alone.

The next result states an *optimality principle for dynamic elections,* which carries over the insight from Bellman's optimality principle for dynamic programming to the electoral framework. It is well-known that in a standard dynamic programming problem, a sufficient condition for a plan to be optimal is that in every state—given the choices determined by the plan in the future—the choice dictated at the current state maximizes the expected discounted payoff of the decision maker. We apply this insight to the electoral model as follows: the outcome of each election is the representative voter's preferred choice, and the voter calculates the expected discounted payoffs from the incumbent and challenger taking her future choices as given; therefore, the choices determined by her equilibrium voting strategy constitute an optimal plan in the hypothetical optimal retention problem. That is, the representative voter is passive in our framework and does not assume she can control future electoral outcomes, but *even if the representative voter could control the outcomes of future elections,* her expected discounted payoff would not increase above its equilibrium level. This observation is made by Duggan and Forand (2014) in a related model of dynamic elections with complete information and an evolving state variable.

Proposition 4.2 *In the infinite-horizon electoral model, in either the spatial preferences or rent-seeking environment, given any stationary electoral equilibrium* $\psi = (\sigma, \mu)$*, the voting strategy of the representative voter type solves the optimal retention problem in the associated model with a unitary voter.*

Finally, we follow our observation for the two-period model by noting a straightforward *anti-folk theorem* for the infinite-horizon model in which voters observe the type of the incumbent. We modify the above framework so that voters observe the type of a politician once she takes office, and we consider subgame perfect equilibria that are deferential and stationary in the sense that voters and politicians do not condition on actions prior to the current period.

Proposition 4.3 *In the infinite-horizon electoral model in which voters observe the incumbent's type, every deferential, stationary subgame perfect equilibrium is such that office holders always choose their ideal policies.*

Indeed, given such an equilibrium, we can write the expected discounted utility of the median voter from a type j incumbent as V_m^j and from an unknown challenger as V_m^C; importantly, by stationarity, the expected payoff V_m^j from the incumbent is a constant and does not depend on the history of play. In a deferential equilibrium, the median voter votes to re-elect the incumbent if and only if $V_m^j \geqslant V_m^C$, and so ρ_m is constant. Then the type j office holder solves

$$\max_{x \in X} w_j(x) + \beta + \delta \left[\rho_m [w_j(\hat{x}_j) + \beta] + (1 - \rho_m) V_j^C \right],$$

which obviously has the unique solution $x = \hat{x}^j$. As we saw in the two-period model of Subsection 3.3, the absence of uncertainty about the incumbent's type removes all reputational concerns of the politician, electoral incentives lose all disciplining power, and the only possible equilibrium behavior replicates myopic play. This observation holds regardless of whether actions are perfectly observed and regardless of the preference environment.

5 Pure adverse selection

In this section, we consider the dynamic elections framework with spatial preferences and perfect monitoring. The seminal paper studying political accountability in an infinite horizon model with perfect monitoring is Barro (1973). As opposed to Barro, we consider spatial preferences rather than rent-seeking; more importantly, and consistent with the emphasis on reputation incentives, we consider several types of politicians, so that we present a model of adverse selection. Throughout this section, we assume that the policy choice x of an office holder determines the policy outcome $y = x$ with no noise, or in the terminology of the previous section, that the distribution $F(\cdot|x)$ over policies given policy choice x puts probability one on x. Thus, we drop the distinction between policy choices and outcomes, and we assume voter preferences are defined over x directly. A symmetric version of this model in which types are continuously distributed is investigated by Duggan (2000), and Bernhardt, Dubey, and Hughson (2004) consider the model with an arbitrary finite term limit. Banks and Duggan (2008) provide theorems on existence of a class of simple equilibria in the multidimensional model, and they give conditions on the one-dimensional model under which equilibrium policies converge to the median voter's ideal policy.

We begin with the analysis of existence and uniqueness of a special class of "simple electoral equilibria," in which updating occurs once when an incumbent is initially re-elected but then ceases: the acceptance set is fixed through time and the optimal policy of the office holder remains unchanged. We then show that such equilibria have the partitional form familiar from Subsection 3.3, and we establish a *strong form of responsive democracy* in the model without term limits: if either citizens are sufficiently patient or the office benefit is sufficiently high, then in equilibrium, all politician types always choose the median policy. This result does not carry over to the model with a two-period term limit, however, as the *commitment problem of voters* curtails the electoral incentives of politicians. We end the section with a discussion of several extensions of the model to allow for more realistic assumptions on partisanship, voter preferences, and political campaigns.

5.1 Existence of simple equilibria

The literature has focussed on a particularly simple form of equilibrium, such that the acceptance set remains unchanged when an incumbent chooses an acceptable policy and is re-elected. An implication is that a politician always chooses the same policy while in office, so that along the personal path of play of a politician, voters update after the first term of office but do not update after subsequent terms of the politician. In such an equilibrium, for all beliefs $b \in \Delta(T)$ and all acceptable policies $x \in A(b|\psi)$, we have

$$A(\mu_T(b,x)|\psi) \quad = \quad A(b|\psi), \tag{9}$$

i.e., after an acceptable policy is chosen and voters update beliefs, the acceptance set remains unchanged. In a politician's first term, the acceptance set in such an equilibrium is $A(p|\psi)$, where the voters' beliefs are given by the prior p. If the politician is type j and her ideal policy $\hat{x}_j$ belongs to the acceptance set $A(p|\psi)$, then she can secure re-election by choosing $\hat{x}_j$, and since the acceptance set remains the same, she can continue to choose her ideal policy and gain re-election in every period. Such a politician type is a "winner," and their optimization problem is trivial. Otherwise, if the politician's ideal policy does not belong to the acceptance set, then the office holder faces a trade off: either (i) shirk by choosing her ideal policy $\hat{x}_j$ today, foregoing re-election, or (ii) compromise by choosing a policy that is in the acceptance set, though not ideal, in order to gain re-election. Thus, the optimization problem is

$$\max\left\{u_j(\hat{x}_j)+\beta+\delta V_j^C(\psi),\ \max_{x\in A(p|\psi)}\frac{u_j(x)+\beta}{1-\delta}\right\}. \tag{10}$$

Since $A(p|\psi)$ is a closed interval and u_j is strictly quasi-concave, this means that if the politician's ideal policy does not belong to the acceptance set, then she either shirks or chooses the endpoint of $A(p|\psi)$ closest to her ideal policy.

From the perspective of voters, because the median is a representative voter, we know $A(p|\psi)$ consists of every policy x such that $V_m^I(p,x|\psi) \geqslant V_m^C(\psi)$. And since an incumbent who is re-elected continues to choose the same policy, this means that

$$A(p|\psi) \subseteq \left\{x \in X : \frac{u_m(x)}{1-\delta} \geqslant V_m^C(\psi)\right\},$$

so a policy can be acceptable to the median voter only if the utility from that policy is at least equal to the continuation value of a challenger. The maximally permissive acceptance set that is consistent with this criterion is

$$A(p|\psi) \quad = \quad \left\{x \in X : \frac{u_m(x)}{1-\delta} \geqslant V_m^C(\psi)\right\}, \tag{11}$$

so that a policy choice gains re-election if and only if the median voter weakly prefers that policy to a challenger. Because the continuation value of a challenger is stationary, the above acceptance set is independent of the voters' beliefs p, and so the acceptance set under the maximally permissive criterion is constant, and we write it simply as $A(\psi)$. If a stationary electoral equilibrium $\psi = (\sigma, \mu)$ is such that all type j politicians solve (10) and acceptance sets satisfy (9) and (11), then we say it is a *simple electoral equilibrium.*

Banks and Duggan (2008) establish existence of equilibria in this class.

Proposition 5.1 *In the infinite-horizon model of adverse selection of pure adverse selection, there exists a simple electoral equilibrium.*

Details of the proof are omitted, but we note that it relies on a structural similarity with infinite-horizon bargaining models based on the protocol of Baron and Ferejohn (1989); in particular, it follows along the lines of the existence proof in the spatial bargaining model of Banks and Duggan (2000). To provide some insight into the parallels, an office holder's choice of policy in the electoral model is similar to a proposer making a proposal in the bargaining model; the election is similar to a vote over the proposal; and the random draw of a challenger is similar to the random selection of a new proposer. The main difference between the two frameworks is that in the electoral model, the policy "proposed" by an office holder goes into effect for one period before it is voted on; if the proposal passes (i.e., the politician is re-elected), then it remains in place forever, and if it fails, then a new politician makes a policy choice in the next period.

The above proposition does not address the question of uniqueness of equilibrium, which is proved in symmetric models with continuous types by Duggan (2000) and Bernhardt, Campuzano, Squintaini, and Camara (2009). An analogue is available in a symmetric version of the current framework, but it relies on our maintained assumption that a challenger is the median type with positive probability, i.e., $p_m > 0$. When this assumption is violated, it is easy to obtain multiple equilibria. In Figure 9, for example, we assume there are three types, $T = \{1, 2, 3\}$, where the median voter type is $m = 2$ but challengers are drawn from the extreme types with equal probability, i.e., $p_1 = p_3 = \frac{1}{2}$. Assuming quadratic utility and placing the extreme types at equal distance from the median, it is trivially a simple electoral equilibrium for all politicians to choose their ideal policies in every period and for voters to re-elect an incumbent following any choice in the interval $[\hat{x}_1, \hat{x}_3]$ of ideal points. Adding the assumption that $\delta > 0$, we can modify these strategies for small $\varepsilon > 0$ so that the extreme types compromise by choosing $\hat{x}_1 + \varepsilon$ and $\hat{x}_3 - \varepsilon$, and voters re-elect following any choice in the interval $[\hat{x}_1 + \varepsilon, \hat{x}_3 - \varepsilon]$. In particular, since the discount factor is positive and ε is small, each politician

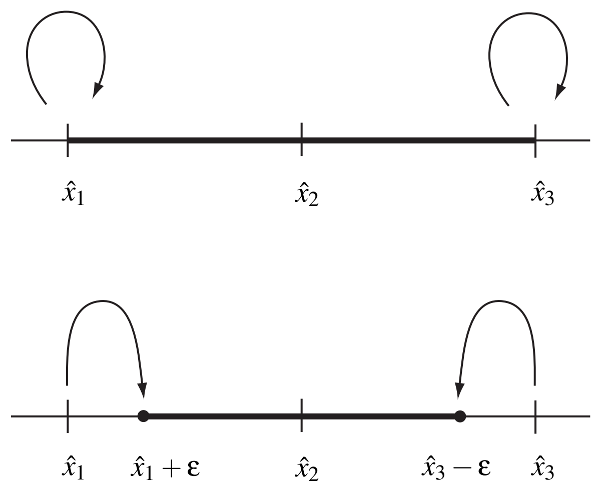

Figure 9: Multiple equilibria

prefers to compromise rather than shirk. Varying ε, we then obtain a continuum of non-payoff equivalent equilibria.

5.2 Partitional characterization

The simple electoral equilibria established in Proposition 5.1 involve a separation of politician types that has the partitional structure familiar from Proposition 3.2 in the two-period model. Given a simple electoral equilibrium $\psi = (\sigma, \mu)$ with acceptance set $A(\psi)$, let

$$\begin{aligned} W &= \{j \in T : \hat{x}_j \in A(\psi)\}, \\ C &= \left\{ j \in T \setminus W : \frac{1}{1-\delta}\left(\max_{x \in A(\psi)} u_j(x) + \beta \right) \geqslant u_j(\hat{x}_j) + \beta + \delta V_j^C(\psi) \right\}, \\ L &= T \setminus (W \cup C). \end{aligned}$$

Note that the expected discounted payoff to a politician type $j \in C$ from choosing the best acceptable policy in the current and all future terms of office is just

$$\frac{1}{1-\delta} \max_{x \in A(\psi)} u_j(x) + \beta,$$

while the expected discounted payoff from shirking by choosing the ideal policy $\hat{x}_j$ and foregoing re-election is

$$u_j(\hat{x}_j) + \beta + \delta V_j^C(\psi),$$

so such politicians weakly prefer to compromise to retain office. In contrast, types $j \in L$ strictly prefer to shirk at the cost of losing the election. Thus, we refer to politicians in these sets, respectively, as "winners," "compromisers," and "losers." Let $\ell = \min W$ and $r = \max W$ denote the smallest and largest winning types, respectively.

Clearly, the winning types have centrally located ideal policies in the interval $A(\psi)$ around the median voter's ideal policy. Intuitively, a politician whose ideal policy is outside but close to the acceptance set should also compromise, as the cost of doing so is small. The cost of compromise may be prohibitive for some extreme types, which are thus losing, but it is possible in principle that even more extreme types may have incentives to compromise in order to avoid electing a challenger who chooses policy at the opposite extreme of the policy space. Such types have more to lose than moderate politicians by inserting a challenger in office. The next proposition establishes that under our assumptions, this phenomenon does not arise in equilibrium, and that the set of compromising types consists of two "connected" sets on either side of the acceptance set. Note that it is possible that the compromise set is empty.

Proposition 5.2 *In the infinite-horizon model of pure adverse selection, let $\psi = (\sigma,\mu)$ be a simple electoral equilibrium. Then there exist integers k' and k'' such that $k' \leqslant \ell \leqslant m \leqslant r \leqslant k''$ and*

$$\begin{aligned} W &= \{\ell, \ell+1, \ldots, r-1, r\} \\ C &= \{k', \ldots, \ell-1\} \cup \{r+1, \ldots, k''\}. \end{aligned}$$

We must argue that for a simple electoral equilibrium $\psi = (\sigma,\mu)$, the compromise set has the form above. It suffices to show, letting k'' be the highest compromising type, that the compromise set contains all types j with $r < j < k''$. Let $A(p|\psi) = [x', x'']$. As in Subsection 2.4, we can write the normalized expected discounted payoff from electing a challenger for a type j voter as the expected payoff from a lottery L on X, so that

$$(1-\delta)V_j^C(\psi) \;=\; \mathbb{E}_L[u_j(x)] \;=\; \sum_x L(x)u_j(x).$$

For a general value $\theta > \theta_m$, we can use compactness of X and strict concavity of $u(x|\theta) = \theta v(x) - c(x)$ in x to define the unique ideal policy

$$\hat{x}(\theta) \quad = \quad \arg\max_{x \in X} u(x|\theta).$$

A hypothetical politician with parameter θ is then willing to compromise at x' if and only if

$$\frac{1}{1-\delta}\left(u(x'|\theta) + \beta\right) \quad \geqslant \quad u(\hat{x}(\theta)|\theta) + \beta + \frac{\delta}{1-\delta}\mathbb{E}_L[u_j(x)],$$

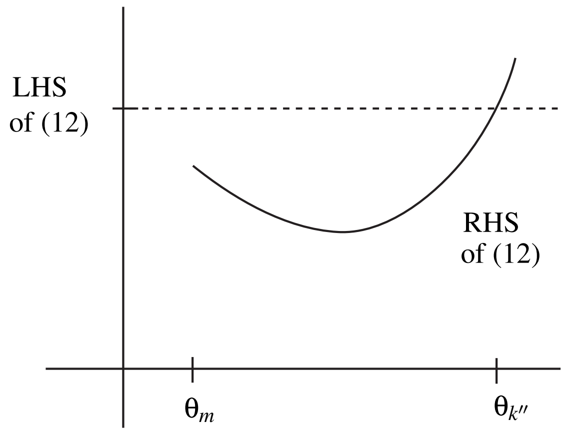

Figure 10: Single crossing for compromisers

or equivalently,

$$\delta\left(\beta+\sum_x L(x)c(x))\right)-c(x')\geqslant(1-\delta)u(\hat{x}(\theta)|\theta)-\theta\left(v(x')-\delta\sum_x L(x)v(x)\right). \quad (12)$$

Noting that the left-hand side of (12) is constant in θ, the envelope theorem implies that the first derivative of the right-hand side with respect to θ is

$$(1-\delta)v(\hat{x}(\theta))-v(x')+\delta\sum_x L(x)v(x).$$

Since $\hat{x}(\theta)$ is increasing in θ and v is increasing, this derivative is increasing in θ, and it follows that the right-hand side of (12) is convex in θ. Clearly, the inequality holds when $\theta=\theta_m$ and $\theta=\theta_{k''}$, and thus it holds for all types j with $\theta_m<\theta_j<\theta_{k''}$; see Figure 10. Thus, the compromising types $j>m$ form a connected set $\{r+1,\ldots,k''\}$, and a symmetric argument addresses compromising types less than the median type.

The above proposition delivers a partitioning of types similar to that of Figure 4. A difference between the current equilibrium analysis and the backward induction construction of Subsection 3.3 is that now mixed policy strategies are needed to ensure existence; in particular, it is necessary to allow some types of politician to mix between compromising and shirking. The single-crossing argument for Proposition 5.2 implies, however, that the need for mixing is limited: in a simple electoral equilibrium, there is at most one type on each side of the median voter that mixes.

This is seen immediately from the indifference condition

$$v(x') - \delta \sum_x L(x)v(x) + \beta \quad = \quad (1-\delta)u(\hat{x}(\theta)|\theta) + \theta\delta \sum_x L(x)(c(x') - c(x)),$$

which is satisfied by at most one $\theta > \theta_m$. If there is no type j such that $\theta_j = \theta$, then there is no politician type $j < m$ that mixes in the first term of office; otherwise, there may be exactly one type to the left of the median that mixes, and similarly to the right; see Figure 11, where types $k' = 2$ and $k'' = 5$ mix in equilibrium.

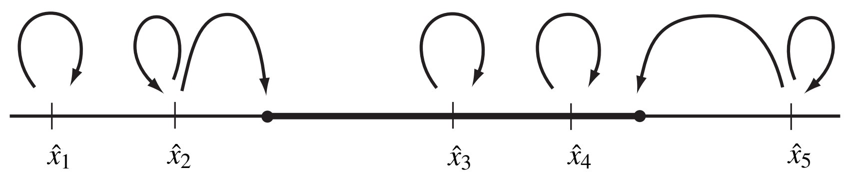

Figure 11: Partitional structure

The preceding observation implies that a simple electoral equilibrium can be characterized as a solution to a system of equations; in fact, equilibria may solve any of a finite number of systems of equations, depending on the cutoff types k', ℓ, r, k''. The system corresponding to a set of cutoffs involves nine equations in nine unknowns, $\pi', \pi'', \xi', \xi'', \alpha', \alpha'', V, V', V''$, where: $\pi', \pi'' \in [0,1]$ represent the probability of compromise by types k' and k'', $\xi' \in [\underline{x}, \hat{x}_m]$ and $\xi'' \in [\hat{x}_m, \overline{x}]$ are the endpoints of the acceptance set, $\alpha', \alpha'' \geqslant 0$ are the net normalized discounted payoffs from compromising over shirking for types k' and k'', and V, V', V'' are the normalized continuation values of a challenger for the median voter and the type k' and k'' voters. For the case in which k' and k'' are indifferent between compromise and shirking, the first seven equations are

$$V \quad = \quad u_m(\xi') \tag{13}$$

$$V \quad = \quad u_m(\xi'') \tag{14}$$

$$\alpha' \quad = \quad u_{k'}(\xi') + \beta - (1-\delta)(u_{k'}(\hat{x}_{k'}) + \beta) - \delta V' \tag{15}$$

$$\alpha'' \quad = \quad u_{k''}(\xi'') + \beta - (1-\delta)(u_{k''}(\hat{x}_{k''}) + \beta) - \delta V'' \tag{16}$$

$$0 \quad = \quad (1-\pi')\alpha' \tag{17}$$

$$0 \quad = \quad (1-\pi'')\alpha'' \tag{18}$$

$$V \quad = \quad \sum_{j=1}^{k'-1} p_j[(1-\delta)u_m(\hat{x}_j) + \delta V] \tag{19}$$

$$
\begin{aligned}
&+p_{k'}[\pi' u_m(\xi') + (1-\pi')((1-\delta)u_m(\hat{x}_{k'}) + \delta V)] \\
&+ \sum_{j=k'+1}^{\ell-1} p_j u_m(\xi') + \sum_{j=\ell}^{r} u_m(\hat{x}_j) + \sum_{j=r+1}^{k''-1} p_j u_m(\xi'') \\
&+p_{k''}[\pi'' u_m(\xi'') + (1-\pi'')((1-\delta)u_m(\hat{x}_{k''}) + \delta V] \\
&+ \sum_{k''+1}^{n} p_j[(1-\delta)u_m(\hat{x}_j) + \delta V],
\end{aligned}
$$

with equations for V' and V'' defined analogously. Here, (13) and (14) are indifference conditions for the median voter; (15) and (16) give the net advantage of compromise for types k' and k''; (17) and (18) allow mixing between compromise and shirking by k' and k'' when indifferent; and (19) gives the continuation value of a challenger for the median voter. This characterization, while notationally complex, implies that the model is computationally tractable, as the unique simple electoral equilibrium can be found as the solution to a relatively small system of equations.

5.3 Responsive democracy

To further characterize stationary electoral equilibria, we first note that in any such equilibrium (σ,μ), given any beliefs b and acceptable policy x, the acceptance set following the choice of x and updating by voters is nonempty.

Proposition 5.3 *In the infinite-horizon model of pure adverse selection, let $\psi = (\sigma,\mu)$ be a stationary electoral equilibrium. Then for all $b \in \Delta(T)$ and all $x \in A(b|\psi)$, we have $A(\mu_T(b,x)|\psi) \neq \emptyset$.*

Let $b' = \mu_T(b,x)$ denote the voters' beliefs given prior b and observed policy $x \in A(b|\psi)$. If $A(b'|\psi) = \emptyset$, then each type j with $b'_j > 0$ chooses her ideal policy $\hat{x}_j$. Then each type is revealed to voters, and we have

$$
V_m^I(b,x|\psi) \;=\; \sum_j \rho_j[u_m(\hat{x}_j) + \delta V_m^C(\psi)] \;\geqslant\; V_m^C(\psi),
$$

where the inequality follows because x is acceptable. This implies that for some type k, we have $u_m(\hat{x}_k) \geqslant (1-\delta)V_m^C(\psi)$. Conditional on observing $\hat{x}_k$, voters update that the politician is type k with probability one. If the incumbent is re-elected, then she either chooses her ideal policy and is replaced, or she chooses an acceptable policy, and so the expected payoff to the median voter from re-electing the politician satisfies

$$
V_m^I(b',\hat{x}_k|\psi) \;\geqslant\; \min\{u_m(\hat{x}_k) + \delta V_m^C(\psi), V_m^C(\psi)\} \;\geqslant\; V_m^C(\psi).
$$

Since (σ,μ) is deferential, we conclude that $\hat{x}_k \in A(b'|\psi)$, a contradiction.

The previous proposition assumes existence of an acceptable policy, but it is clear that the acceptance set is nonempty in an office holder's first term, when the voters' beliefs are given by the prior. Indeed, suppose otherwise. Then all politician types choose their ideal policies in equilibrium, and in particular, the median type m politician chooses $\hat{x}_m$ and reveals her type. If the median voter re-elects the type m incumbent, then the politician either shirks by choosing the median policy or chooses an acceptable policy, and in both cases the expected discounted payoff to the median voter is at least that of electing a challenger. Thus, the incumbent is re-elected, and we conclude that $\hat{x}_m$ is acceptable, a contradiction. We have argued for the following result.

Proposition 5.4 *In the infinite-horizon model of pure adverse selection, let* $\psi = (\sigma,\mu)$ *be a stationary electoral equilibrium. Then* $A(p|\psi) \neq \emptyset$.

An implication of Propositions 5.3 and 5.4 is that the type m politician always chooses the median ideal policy in her first term: the acceptance set is nonempty, so by monotonicity, $\hat{x}_m$ is acceptable, and by the previous proposition, the acceptance set continues to be nonempty when the politician chooses $\hat{x}_m$, yielding the maximum possible payoff to the politician.

Next, we investigate the implications for equilibrium outcomes when politicians are highly office motivated, and in particular the possibility of obtaining a dynamic version of the Downsian median voter result. Note by Propositions 5.3 and 5.4, that the acceptance set is nonempty in a politician's first term of office, and that by choosing an acceptable policy, a politician can secure re-election with probability one in every period. One strategy that achieves this is simply choosing the median $\hat{x}_m$ in every period, but we have not ruled out other strategies that deliver the same outcome. The payoff to a type j politician from choosing the median in the first and all future terms of office is

$$\frac{u_j(\hat{x}_m)+\beta}{1-\delta},$$

and the payoff from shirking is no more than

$$u_j(\hat{x}_j)+\beta+\frac{\delta u_j(\hat{x}_j)}{1-\delta}.$$

Thus, assuming the discount factor is positive, a sufficient condition for all types to compromise in all terms of office is

$$\beta \;>\; \frac{1}{\delta}\max_j[u_j(\hat{x}_j)-u_j(\hat{x}_m)] \;\equiv\; \overline{\beta}.$$

Denoting the right-hand side of the above inequality by $\overline{\beta}$, it follows that when the benefit of office exceeds this level, all politician types will secure re-election in all periods.

Proposition 5.5 *In the infinite-horizon model of pure adverse selection, assume $\beta > \overline{\beta}$, and let $\psi = (\sigma, \mu)$ be a stationary electoral equilibrium. Then for all types $j \in T$, all ℓ, and all personal histories $\mathbf{y}^\ell$ along the path of play, the type j politician chooses an acceptable policy with probability one, i.e., $\pi_j(A(\mu_T(\mathbf{y}^\ell)|\psi)|\mathbf{y}^\ell) = 1$.*

The preceding proposition gives a stability result for dynamic elections with office-motivated politicians: in equilibrium, office holders use their incumbency advantage to ensure continual re-election, so there is no turnover in office. But the policies chosen to achieve this outcome may be undesirable to the median voter, contrary to a responsive democracy result. The next result establishes that the median voter's ex ante payoff converges to the ideal payoff if politicians are sufficiently office-motivated and citizens are sufficiently patient. This means that all stationary electoral equilibria determine policies close to the median with high probability, delivering a *strong responsive democracy result* for the class of stationary electoral equilibria. Note that we normalize continuation values in the following result, so that we can compare them with expected per period payoffs.

Proposition 5.6 *In the infinite-horizon model of pure adverse selection, assume that $\beta > u_j(\hat{x}_j) - u_j(\hat{x}_m)$ for all $j \in T$, and let the discount factor δ approach one. Then for every selection of stationary electoral equilibria $\psi = (\sigma, \mu)$ given discount factor δ, the median voter's (normalized) continuation value of a challenger converges to the ideal payoff, i.e.,*

$$\lim_{\delta \to 1} (1-\delta) V_m^C(\psi) \;=\; u_m(\hat{x}_m).$$

The argument for this result is facilitated by the observation that the type m politician simply chooses $\hat{x}_m$ in every period and is complicated by the possibility that less desirable types pool with the median type in early terms of office. We begin by considering a fixed δ close enough to one such that $\beta > \overline{\beta}$ holds and an equilibrium ψ given δ, so that by Proposition 5.5 all politician types choose acceptable policies in all terms of office. Note that

$$V_m^C(\psi) \;=\; \sum_j p_j \sum_x [u_m(x) + \delta V_m^I(p,x|\psi)] \pi_j(x|p), \tag{20}$$

where we use the fact that incumbents are re-elected with probability one, from Proposition 5.5. Recall that Propositions 5.3 and 5.4 imply that the type m politician chooses the median policy in every term of office. Then, since $V_m^I(p,x|\psi) \geqslant$

$V_m^C(\psi)$ along the path of play and utilities are non-negative, it follows that

$$V_m^C(\psi) \geqslant \delta p_m V_m^I(p,\hat{x}_m|\psi) + \delta(1-p_m)V_m^C(\psi),$$

or equivalently,

$$V_m^I(p,\hat{x}_m|\psi) \leqslant \left(1+\frac{1-\delta}{\delta p_m}\right)V_m^C(\psi). \tag{21}$$

We are interested in personal histories along the path of play in which the median policy is chosen initially. Let $\mathbf{z}^k = (\hat{x}_m,\ldots,\hat{x}_m,y)$ denote a personal history in which the median policy is chosen for the first k terms of office followed by a choice $y \neq \hat{x}_m$; let Z^k be the set of such personal histories that occur with positive probability along the path of play determined by ψ; and let $\mathbf{z}^\infty = (\hat{x}_m,\hat{x}_m,\ldots)$ be the infinite sequence of median policies. Let $V_m^I(\mathbf{z}^k|\psi)$ denote the expected discounted payoff to the median voter from re-electing an incumbent with personal history $\mathbf{z}^k$ in equilibrium ψ. Finally, we denote by $\Pr(\mathbf{z}^k)$ the probability of $\mathbf{z}^k$ determined by ψ, conditional on choice $\hat{x}_m$ in the first term. Then

$$V_m^I(p,\hat{x}_m|\psi) = \Pr(\mathbf{z}^\infty)\frac{u_m(\hat{x})}{1-\delta} + \sum_{k=1}^{\infty}\sum_{\mathbf{z}^k\in Z^k}\Pr(\mathbf{z}^k)(\delta)^{k-1}V_m^I(\mathbf{z}^k|\psi),$$

where $\Pr(\mathbf{z}^\infty) \geqslant p_m > 0$. For simplicity, assume that the infimum of $V_m^I(\mathbf{z}^k|\psi)$ is attained over k and $\mathbf{z}^k \in Z^k$ by $\mathbf{z}^{k*}$.[21]

We claim that

$$V_m^I(\mathbf{z}^{k*}|\psi) \geqslant \frac{u_m(\hat{x}_m)}{1-\delta} - \frac{(1-\delta)u_m(\hat{x}_m)}{\delta p_m^2},$$

for suppose otherwise. Then we have

$$\begin{aligned} V_m^I(\mathbf{z}^{k*}|\psi) &< \Pr(\mathbf{z}^\infty)\left(\frac{u_m(\hat{x})}{1-\delta} - \frac{(1-\delta)u_m(\hat{x}_m)}{\delta p_m^2}\right) + \sum_{k=1}^{\infty}\sum_{\mathbf{z}^k\in Z^k}\Pr(\mathbf{z}^k)\delta^{k-1}V_m^I(\mathbf{z}^k|\psi) \\ &\leqslant V_m^I(p,\hat{x}_m|\psi) - p_m\left(\frac{(1-\delta)u_m(\hat{x}_m)}{\delta p_m^2}\right) \\ &\leqslant V_m^C(\psi), \end{aligned}$$

where the last inequality uses (21). But then an incumbent with personal history $\mathbf{z}^{k*}$ is not re-elected, a contradiction. This establishes the claim.

[21] If the infimum is not attained, then the argument is easily modified to choose a $\mathbf{z}^{k*}$ that yields a payoff close to the infimum.

Note that a lower bound for $V_m^I(p,\hat{x}_m|\psi)$ is obtained if we suppose that conditional on choice $\hat{x}_m$ in the first term, all types other than the median choose the worst possible policy in the second term. By the claim, this policy minimizes the median voter's utility in the current period (and this minimum payoff is non-negative) but entails a payoff close to the ideal thereafter; thus,

$$V_m^I(p,\hat{x}_m|\psi) \;\geqslant\; p_m\left(\frac{u_m(\hat{x}_m)}{1-\delta}\right) + (1-p_m)\delta\left(\frac{u_m(\hat{x}_m)}{1-\delta} - \frac{(1-\delta)u_m(\hat{x}_m)}{\delta p_m^2}\right),$$

which delivers the desired result. Combining the above inequality with (21), we have

$$(1-\delta)V_m^C(\psi) \;\geqslant\; \left(\frac{\delta p_m}{1+\delta p_m-\delta}\right)\left(p_m+\delta-\delta p_m - \frac{(1-p_m)\delta(1-\delta)^2}{\delta p_m^2}\right)u_m(\hat{x}_m).$$

Taking the limit as $\delta \to 1$, we have $(1-\delta)V_m^C(\psi) \to u_m(\hat{x}_m)$, as desired.

Focussing on simple electoral equilibria, Banks and Duggan (2008) provide a tighter responsive democracy result. For a simple equilibrium $\psi = (\sigma,\mu)$, since the acceptance set $A(\psi)$ is fixed, a type j politician can simply choose the best acceptable policy to secure re-election in every period, so the payoff from compromising in every term of office is $\max_{x\in A(\psi)} u_j(x)/(1-\delta)$, and all types will compromise in equilibrium if

$$\frac{1}{1-\delta}\max_{x\in A(\psi)} u_j(x) + \frac{\beta}{1-\delta} \;\;>\;\; u_j(\hat{x}_j) + \beta + \delta V_j^C(\psi).$$

A risk aversion argument implies that the first term on the left-hand side is greater than the last term on the right-hand side, and so a sufficient condition for all politicians to compromise is that for all types $j \in T$,

$$\frac{\beta\delta}{1-\delta} \;\;\geqslant\;\; u_j(\hat{x}_j). \tag{22}$$

Because all politician types compromise, under this condition, it follows that every type is "above average," and therefore the median voter's expected payoff conditional on each $x \in A(\psi)$ on the path of play is equal to $V_m^C(\psi)$. In particular, the expected payoff from re-electing a type m politician equals the expected payoff from a challenger, and this is only possible if all types choose the median. This logic delivers the following result, a sharp responsive democracy result for simple electoral equilibria.

Proposition 5.7 *In the infinite-horizon model of pure adverse selection, assume inequality (22) holds for all types, and let $\psi = (\sigma,\mu)$ be a simple electoral equilibrium. Then each type j politician chooses the median policy in the first term, i.e., $\pi_j(\{\hat{x}_m\}|p) = 1$, and in all future terms of office.*

Note that inequality (22) holds if $\delta > 0$ and β is sufficiently large, and it holds if $\beta > 0$ and δ is sufficiently close to one. This yields the following corollary, which extends Proposition 5.6 by establishing a median voter theorem even when the office benefit is small.

Corollary 5.1 *In the infinite-horizon model of pure adverse selection, assume* $\delta\beta > 0$*. If either* β *is sufficiently large, or* δ *is sufficiently close to one, then in every simple electoral equilibrium, each type* j *politician chooses the median policy in every term of office.*

Although the result is silent on the case of policy-motivated candidates, Banks and Duggan (2008) show that when there is no office benefit, i.e., $\beta = 0$, if the discount factor approaches one, then the acceptance set (and the policy choices of all politician types) converge to the median policy. Again, elections facilitate commitment, and electoral incentives lead to responsive democracy if politicians are patient or office benefit is large.

5.4 Term limits

Bernhardt, Dubey, and Hughson (2004) consider a version of the model with a continuum of types and an arbitrary, finite term limit, but for tractability we discuss just the two-term-limit version of the model. Clearly, as in Section 3.3, second-term politicians simply choose their ideal policies in equilibrium, a feature that qualitatively changes the equilibrium dynamics of the infinite-horizon model, and in contrast to Section 3.3, we no longer obtain the responsive democracy result when politicians are office motivated and citizens are patient. Even equilibrium existence becomes problematic, and we must relax our restriction of deferential voting and allow for mixed electoral outcomes (i.e., endogenous uncertainty about an incumbent's prospects for re-election). We extend the concept of stationary electoral equilibrium to allow politicians' strategies to depend on the term of office, while still imposing optimality of these choices. We assume that a first-term incumbent who chooses x is re-elected if that is the strict preference of the median voter, i.e., $V_m^I(p,x|\psi) > V_m^C(\psi)$, and that the challenger is elected if the reverse inequality holds; if the median is indifferent, then the probability of re-election, $\rho_m(x)$, may now be between zero and one. We continue to assume voters' strategies are monotonic, but instead of simply expressing this condition in terms of a utility cutoff for the median voter, we require that the probability of re-election is weakly increasing in the utility of the median voter from a first-term office holder's policy choice.

To see why mixed voting is needed, we consider the case of highly office-motivated candidates and suppose instead that voting is pure. We first note that the acceptance set for first-term office holders is nonempty (otherwise, the type m politician chooses $\hat{x}_m$, revealing her type, and so she is re-elected), so $A(p|\psi) \neq \emptyset$. The payoff to a type j politician from compromising is then at least equal to

$$u_j(\hat{x}_m) + \beta + \delta[u_j(\hat{x}_j) + \beta] + \delta^2 V_j^C(\psi),$$

and the payoff from shirking is no more than

$$u_j(\hat{x}_j) + \beta + \delta V_j^C(\psi),$$

so that all politician types compromise in the first term if

$$u_j(\hat{x}_m) + \beta \quad > \quad (1-\delta)\left[u_j(\hat{x}_j) + \beta + \delta \frac{u_j(\hat{x}_j)}{1-\delta}\right].$$

This holds if $\beta > \overline{\beta}$, which we again assume, so that every type of first-term office holder will compromise by choosing the best element of $A(p|\psi)$.

We demonstrate a contradiction to existence of equilibrium in pure voting strategies in two exhaustive and mutually exclusive cases, assuming without loss of generality that the median voter weakly prefers $\hat{x}_1$ to $\hat{x}_n$. Note that

$$V_m^C(\psi) \quad = \quad \sum_j p_j \sum_x [u_m(x) + \delta u_m(\hat{x}_j) + \delta^2 V_m^C(\psi)]\pi_j^1(x),$$

where π_j^1 is the proposal strategy of the type j politician in her first term of office. In the first case, the ideal policy of the type n politician belongs to the acceptance set, so all politician types simply choose their ideal policies, revealing their types to voters, and they are continually re-elected. It follows that voters place probability one on type $j = n$ after observing $\hat{x}_n$, and then

$$\begin{aligned} V_m^C(\psi) \quad &\geqslant \quad p_m(1+\delta)u_m(\hat{x}_m) + (1-p_m)(1+\delta)u_m(\hat{x}_n) + \delta^2 V_m^C(\psi) \\ &> \quad (1+\delta)u_m(\hat{x}_n) + \delta^2 V_m^C(\psi). \end{aligned}$$

This implies that

$$V_m^C(\psi) \;>\; u_m(\hat{x}_n) + \delta V_m^C(\psi) \;=\; V_m^I(p, \hat{x}_n|\psi),$$

and the median voter strictly prefers to elect a challenger over re-electing the type n incumbent, a contradiction.

In the second case, the acceptance set excludes $\hat{x}_n$. Then since type n politicians (and all other types whose ideal policies are not in the acceptance set) compromise by choosing an acceptable policy in their first term, the expected payoff to the median voter from policies chosen in the first term exceeds the lottery over ideal policies, i.e.,

$$\sum_j p_j \sum_x u_m(x)\pi_j^1(x) \quad > \quad \sum_j p_j u_m(\hat{x}_j).$$

Let y minimize the median voter's expected payoff $V_m^I(p,x|\psi)$ from re-electing the incumbent over the equilibrium policy choices x of first-term office holders. Then

$$\begin{aligned} V_m^I(p,y|\psi) &\leqslant \sum_j p_j \sum_x V_m^I(p,x|\psi)\pi_j^1(x) \\ &= \sum_j p_j u_m(\hat{x}_j) + \delta V_m^C(\psi) \\ &< \sum_j p_j \sum_x u_m(x)\pi_j^1(x) + \delta \sum_j p_j u_m(\hat{x}_j) + \delta^2 V_m^C(\psi) \\ &= V_m^C(\psi), \end{aligned}$$

and the median voter strictly prefers to elect a challenger over re-electing the incumbent following the choice y, again producing a contradiction. We conclude that equilibria in pure voting strategies do not generally exist, so that uncertainty about electoral outcomes arises as a necessity in the model with term limits.

Driving the failure of existence is the fact that given the opportunity to secure re-election, office-motivated politicians will do so, but this creates a *commitment problem for voters,* who then have an incentive to replace a first-term incumbent with a fresh challenger. To balance these incentives, mixing is needed not just in policy choices of office holders but in electoral outcomes. For example, if an office holder gains a chance of re-election by choosing the median policy and the probability of success is less than one, then some politician types can be dissuaded from compromising, and then the incentive to insert a challenger in place of a first-term incumbent decreases, ameliorating the incentive problem of voters. Thus, the probability of electoral success must be set in equilibrium to obtain the correct separation of politician types and to generate indifference needed for mixed policy choices, and mixing over policies is set to maintain the median voter's indifference between incumbent and challenger over the relevant range.

The full responsiveness result cannot be obtained in the term limit model, for suppose there is an equilibrium $\psi = (\sigma, \mu)$ such that in every period, all types of office holder choose the median ideal policy. Since second-term politicians choose their ideal policies, it must be that in each period, the incumbent is replaced by a challenger with probability one, but then the first term of office is a politician's last, and non-median type $j \neq m$ politicians have no incentive to compromise, so they would simply choose their ideal policies. The most that can be hoped for is that first-term politicians choose the median policy, with slack to allow non-median politicians to be re-elected with positive probability and choose their ideal policies in the second term. But the incentives of the term limit model preclude even that form of responsiveness, even when politicians are highly office motivated and citizens are patient, and even if we only ask that first-term office holders choose policies close to the median with high probability: the next result establishes that the expected payoff to the median voter from equilibrium policy choices of first-term office holders is bounded strictly below the ideal payoff.

Proposition 5.8 *In the infinite-horizon model of pure adverse selection with two-period term limit, there is a bound* $\overline{u} < u_m(\hat{x}_m)$ *such that for all levels of office benefit* $\beta \geqslant 0$ *and all discount factors* $\delta \in [0,1)$*, in every stationary electoral equilibrium* $\psi = (\sigma, \mu)$ *for parameters* (β, δ)*, the expected utility to the median voter from policies chosen by first-term office holders is below this bound, i.e.,*

$$\sum_j p_j \sum_x u_m(x)\pi_j^1(x) \quad \leqslant \quad \overline{u}.$$

The logic for the above result is complicated by the possibility that for high β and δ close to one, first-term office holders are incentivized to choose policies close to the median by a positive but small probability of re-election. The continuation value of a challenger now takes the form

$$\begin{aligned} V_m^C(\psi) \quad = \quad & \sum_j p_j \Bigg[\sum_x [u_m(x) + \delta(\rho_m(x)[u_m(\hat{x}_j) + \delta V_m^C(\psi)] \\ & \qquad + (1 - \rho_m(x))V_m^C(\psi))]\pi_j^1(x) \Bigg], \end{aligned}$$

reflecting the fact that the politician who chooses x is re-elected with probability $\rho_m(x)$. We can break the right-hand side of the above expression into the sum $A + B + C$ of three terms:

$$\begin{aligned} A \quad &= \quad \sum_j p_j \sum_x u_m(x)\pi_j^1(x) \\ B \quad &= \quad \delta \sum_x \rho_m(x) \left(\sum_j p_j [u_m(\hat{x}_j) + \delta V_m^C(\psi)]\pi_j^1(x) \right) \\ C \quad &= \quad \delta \left(\sum_j p_j \sum_x (1 - \rho_m(x))\pi_j^1(x) \right) V_m^C(\psi), \end{aligned}$$

where the term in parentheses in C is the ex ante probability that a first-term office holder fails to be re-elected; denote this quantity by e. We can rewrite B as

$$B = \delta \sum_x \rho_m(x) \left(\sum_k p_k \pi_k^1(x) \right) V_m^I(p,x|\psi),$$

where the term in parentheses is the probability that a first-term office holder chooses policy x; denote this by $d(x)$. Since $\rho_m(x) > 0$ implies that $V_m^I(p,x|\psi) \geqslant V_m^C(\psi)$, we therefore have

$$V_m^C(\psi) \geqslant \sum_j p_j \sum_x u_m(x)\pi_j^1(x) + \delta \sum_x \rho_m(x) d(x) V_m^C(\psi) + \delta e V_m^C(\psi).$$

Using $(\sum_x \rho_m(x) d(x)) + e = 1$, this implies

$$(1-\delta)V_m^C(\psi) \geqslant \sum_j p_j \sum_x u_m(x)\pi_j^1(x).$$

Now suppose that there exist office benefit β and discount factor δ and a stationary electoral equilibrium $\psi = (\sigma, \mu)$ such that the expected payoff from policy choices of first-term office holders approaches the ideal payoff of the median voter, i.e.,

$$\sum_j p_j \sum_x u_m(x)\pi_j^1(x) \to u_m(\hat{x}_m).$$

By the foregoing argument, it follows that the normalized continuation value of a challenger in fact approaches the ideal payoff of median voter, i.e., $(1-\delta)V_m^C(\psi) \to u_m(\hat{x}_m)$. Note that for all politician types j, if π_j^1 puts positive probability on policy x such that $\rho_m(x) = 0$, then $x = \hat{x}_j$, and therefore $\sum_x \text{sign}(\rho_m(x))\pi_j^1(x) \to 1$ for all j. Furthermore, this implies $\sum_x (1 - \text{sign}(\rho_m(x)))d(x) \to 0$. Then we have

$$\begin{aligned} V_m^C(\psi) &\leqslant \sum_x (1-\text{sign}(\rho_m(x)))d(x)V_m^C(\psi) + \sum_x \text{sign}(\rho_m(x))d(x)V_m^I(p,x|\psi) \\ &= \sum_x (1-\text{sign}(\rho_m(x)))d(x)V_m^C(\psi) \\ &\quad + \sum_j \sum_x \text{sign}(\rho_m(x)) p_j \pi_j^1(x)[u_m(\hat{x}_j) + \delta V_m^C(\psi)], \end{aligned}$$

where we use the fact that

$$V_m^I(p,x|\psi) = \sum_j \frac{p_j \pi_j(x)}{d(x)}[u_m(\hat{x}_j) + \delta V_m^C(\psi)].$$

Therefore,

$$V^C(\psi)(1-\delta)\sum_x \text{sign}(\rho_m(x))d(x) \quad \leqslant \quad \sum_j p_j \sum_x \text{sign}(\rho_m(x))\pi_j(x)u_m(\hat{x}_j).$$

Taking limits, we have $u_m(\hat{x}_m) \leqslant \sum_{j \neq m} p_j u_m(\hat{x}_j)$, a contradiction. We conclude that the expected payoff from equilibrium policy choices of first-term office holders is bounded strictly below the median voter's ideal payoff, as desired.

The root cause of this difficulty is again the *commitment problem of voters,* who prefer to replace an incumbent who is expected to shirk in her second term with a challenger who offers close to the ideal payoff to the median voter—even though voters might prefer to commit to re-election contingent on the choice of the median policy. This unraveling does not occur in the two-period model, precisely because there is no third period, so challengers are expected to shirk if elected, and there is no temptation to replace an incumbent who chooses the median.

5.5 Extensions and variations

The pure adverse selection model has been used to study the effects of various types of structure on institutions and policies. Bernhardt, Campuzano, Squintani, and Camara (2009) analyze the effects of parties, i.e., drawing challengers from the side of the spectrum opposite the incumbent. This strengthens the threat of an outside challenger and provides greater discipline of incumbent politicians: more substantial competition leads to greater moderation of policy choices. Put differently, elections provide a stronger form of commitment in the partisan model, because voters know the cost of foregoing re-election is higher for an incumbent, making the prospect of choosing policies closer to the median more credible.

Bernhardt, Camara, and Squintani (2011) add valence to the model, and assume that voters observe an incumbent's valence but not her policy preferences. They show that if restrictions on risk aversion, office benefit, and valence heterogeneity hold, then there is a unique stationary equilibrium, and equilibria possess a partitional form, where now the cutoffs defining the win set and compromise set depend on the valence of the office holder. Furthermore, they show that equilibrium voter welfare increases as the distribution of valence increases in the sense of first order stochastic dominance. Camara (2012) considers dynamic elections in the context of a general equilibrium model of public good provision, where politicians are distinguished by a vector of productivity parameters and choose a tax rate while in office, and voters are distinguished by their productivity of labor and preference for public good. He establishes existence of a stationary equilibrium and shows that equilibria have a partitional form in which, again, the cutoffs defining the win set and compromise set vary with the productivity parameters of the office holder.

Kang (2005) introduces a signaling model of electoral campaigns, in which a challenger can signal her quality by a costly campaign activity. She characterizes an equilibrium in which only high quality challengers run costly campaigns, and whether this occurs is determined by the incumbent's attractiveness to voters: if the incumbent is very strong, then the challenger never signals her quality and the incumbent is always re-elected; and if the incumbent is very weak, then the challenger is automatically elected and does not signal. In the complementary case, however, high quality challengers do signal their quality and defeat incumbents of intermediate strength. In work related to the pure adverse selection model, Meirowitz (2007) considers a model in which politicians have private information about their budget constraints, and Casamatta and De Paoli (2007) assume politicians have private information about the cost of public good production. Kalandrakis (2009) studies electoral dynamics in a two-party system in which each party's type may change stochastically over time.

A literature on dynamic elections with complete information, in which voters observe the types of elected politicians, has received attention and has connections to the framework proposed above. A modeling challenge present in this approach is posed by the anti-folk theorem for dynamic elections, Proposition 4.3, which states that when voters observe the incumbent's type, the only deferential, stationary subgame perfect equilibrium is that in which office holders always choose their ideal policies. To avoid this degenerate prediction, different authors have employed different analytical tactics. Barro (1973) considers a model of public good provision in which voters and politicians are identical (except that politicians receive rents from office), and voters re-elect an incumbent depending on whether the politician's public good production satisfies a threshold. This threshold is chosen optimally by the voters, meaning that after some histories, voters elect a challenger over an incumbent despite being indifferent between the two candidates; formally, these strategies are not deferential. This creates a multiplicity of equilibria, including the degenerate equilibrium in which all politicians shirk and are removed from office, but Barro (1973) essentially selects the optimal stationary equilibrium for the voters. He examines the optimization problems of voters and politicians, and he establishes the responsive democracy result that public good levels converge to the voters' ideal when the rewards of office are large and when the discount factor is close to one.

Van Weelden (2013) considers a model under similar assumptions, but he assumes that office benefit is endogenously determined as an amount of rent seeking chosen by the politician, and that this is desirable to politicians and costly to voters. Thus, he effectively assumes a two-dimensional policy space, and in contrast to our framework, it is not possible to elect a politician whose policy preferences align with the median voter's. Moreover, in his baseline model, Van Weelden (2013) allows a representative voter to directly choose the challenger's type. Like Barro (1973), he avoids the anti-folk theorem by dropping the restriction to deferential voting strategies, and he selects from the ensuing multiplicity of equilibria by analyzing the optimal stationary equilibrium for the voter. In particular, Van Weelden (2013) shows that it is better for the voter to alternately select from two non-median types (using one as a threat for the other) than to always elect the median politician, and he establishes the responsive democracy result that as citizens become patient, the policy and rent-seeking choices of office holders converge to the voter's ideal in the optimal equilibrium. Van Weelden (2015) examines a variant of the model with two possible candidates and three types of voter, and he shows that when citizens are sufficiently patient, the optimal amount of polarization between the candidates' ideal policies is positive.

Aragones, Palfrey, and Postlewaite (2007) use history-dependent strategies to generate equilibria that mirror the simple electoral equilibria of the pure adverse selection model. In the latter model, an office holder may be induced to compromise her policy choices by the incentive to pool with desirable politician types, i.e., to appear like a more moderate politician, who will choose moderate policies in the future and thus be re-elected. In the model of Aragones, Palfrey, and Postlewaite (2007), an office holder who has compromised in the past continues to compromise in the future, but once she deviates, she chooses her ideal policy in all future periods, leading voters to elect a challenger. Thus, an office holder may be induced to compromise to maintain her "reputation" for policy moderation, generating the familiar partitional form of equilibrium in the pure adverse selection model. Formally, the authors use non-stationary strategies to escape the anti-folk theorem.

Duggan and Forand (2014) assume that voters observe an office holder's type, and they avoid the anti-folk theorem by assuming an office holder is committed to a platform once she chooses it, until a variable describing the state of the economy changes; the single-state version of this model leads to equilibria with a partitional form that correspond to the simple electoral equilibria of the pure adverse selection model. They establish strong responsive democracy results in two cases: when politicians are purely policy motivated, i.e., $\beta = 0$, the median voter's equilibrium expected discounted payoff converges to her ideal payoff as citizens become patient; and when the office benefit is sufficiently large, all politician types choose the median voter's ideal policy in equilibrium. The authors extend these results to the multi-state model, and they show that a weaker form of responsive democracy holds even if the median voter's type depends on the state. In this setting, we can imagine that policies are chosen by median voters in all states directly, removing politicians from the equation, in a hypothetical "representative voting game." The authors show that for every stationary Markov perfect equilibrium in the representative voting game, there is an equivalent stationary electoral equilibrium in the dynamic electoral model.

The preceding analysis of the pure adverse selection model has focussed on stationary or simple electoral equilibria, and it is instructive to consider the restrictiveness of this concept by characterizing equilibrium outcomes when more general equilibria are allowed, where voters and politicians can condition their choices on the history of play. It is clear that if policy choices are observable, if there is a positive benefit to holding office, and if politicians are sufficiently farsighted, then almost any path of policies can be supported as the path of play of some perfect Bayesian equilibrium. Driving this simple observation is the fact that for a given office benefit $\beta > 0$, when δ is close enough to one, the discounted sum $\frac{\beta}{1-\delta}$ of potential office benefits can be used to induce an office holder to choose the worst policy in the interval X. Indeed, let $\mathbf{x} = (x_1, x_2, \ldots)$ be an arbitrary sequence of policies in X, and consider a strategy profile such that in each period t, every type of politician chooses x_t; voters re-elect an incumbent as long as she has chosen x_t in every period for which she held office, and otherwise they elect a challenger; and beliefs may be specified arbitrarily. Since all politician types choose the same policies, voters are indifferent between re-electing an incumbent and electing a challenger. Given a type j office holder, the discounted payoff to choosing x_t in period t, being re-elected, and continuing to follow the above strategy is is

$$w_j(x_t) + \beta + \left(\sum_{t'=t+1}^{\infty} \delta^{t'-1} w_j(x_{t'}) \right) + \frac{\delta\beta}{1-\delta},$$

and the best one-shot deviation is to choose her ideal policy $\hat{x}_j$ and forego re-election, which yields a discounted payoff of

$$w_j(\hat{x}_j) + \beta + \sum_{t'=t+1}^{\infty} \delta^{t'-1} w_j(x_{t'}).$$

Note that the latter is less than the former if

$$\frac{\delta\beta}{1-\delta} \quad > \quad w_j(\hat{x}_j) - w_j(x_t),$$

which holds when δ is close enough to one. Thus, it is optimal for each type of office holder to follow the prescribed strategy, thereby supporting the arbitrary path $\mathbf{x}$ of policies.

The above argument is quite general and holds across the complete information and pure adverse selection models, but it relies on increasing δ given a fixed, positive level of office benefit. Duggan (2014a) shows that in the model of pure adverse selection, a more complex equilibrium construction can be used to support arbitrary policy paths when citizens are patient, even if politicians are purely policy motivated. Thus, the restriction to stationarity (or some other restriction on history dependence) is needed to obtain predictive power in the pure adverse selection model.

A related literature departs from the dynamic electoral framework considered above by assuming commitment on the part of the challenger, or both the challenger and the incumbent. The latter models extend the traditional Downsian model, in which both candidates take positions, and includes work by Alesina (1988) and Duggan and Fey (2006). The former class includes papers by Wittman (1977), Kramer (1977), and Forand (2014). In particular, Forand (2014) establishes a form of responsive democracy as politicians become patient, showing that equilibrium dynamics lead to alternation between two policies, and that these policies converge to the median voter's ideal policy as the discount factor approaches one. All of these papers assume complete information. As discussed above, Duggan and Forand (2014) allow for an evolving state variable and assume that the incumbent can make "ex post commitments," i.e., an elected politician's first policy choice in a state of the world is binding until the state changes, but the politician cannot commit to policy in advance of transitioning to a new state.

6 Moral hazard and adverse Selection

We now extend the pure adverse selection model of Section 5 to the model with imperfect monitoring. Throughout this section, we consider the rent-seeking environment with moral hazard, and with the exception of Subsection 6.2, we maintain assumptions (C1)–(C4). Thus, as in Subsection 3.4, a policy choice x stochastically determines an outcome y, which is realized from the density $f(y-x)$. In the rent-seeking environment, we may interpret y as a public good level and x as an effort choice, so that voters prefer higher outcomes, whereas politicians internalize the effort choice and, depending on their types, prefer lower exertion of effort. By the MLRP, higher outcomes are evidence of higher effort choices, and so equilibria are characterized by a cutoff outcome $\overline{y}$ that is necessary and sufficient for re-election of the incumbent, and the question of responsive democracy reduces to inducing the greatest effort possible.

We begin by examining the pure moral hazard model, in which there is a single type, or what is the same, the voters' prior beliefs are concentrated on a single type. In this setting, which has played an important role in the development of the literature, stationary electoral equilibria degenerate because of indifference among candidates, as deferential voting then implies that incumbents are re-elected, regardless of outcome; but then, of course, politicians exert zero effort in equilibrium. Thus, the analysis typically takes the perspective of setting an optimal cutoff for voters. We then proceed to the model with adverse selection and one-sided learning, where we discuss a simplified version of the model with no term limit, and we give a more complete analysis of the model with two-period term limit, after which we close by discussing the literature on symmetric learning, where politicians do not observe their own types and update their beliefs in the same way as voters do.

Throughout, we emphasize the success or limitations of electoral mechanisms in eliciting effort on the part of office holders. In the pure moral hazard model, when deferential voting is dropped, we show that high office benefit leads to a *strong form of responsive democracy:* as politicians become more office-motivated, policy choices increase without limit. In the model with both adverse selection and moral hazard, however, the *commitment problem of voters* implies that equilibrium policy choices are bounded above: the continuation value of a challenger cannot exceed the expected utility from the ideal policy of the highest type. This negative result establishes an inherent bound on responsive democracy, and it raises the question of whether and how this bound can be achieved. In the absence of a term limit, we deduce the *qualified responsive democracy result* that the continuation value of a challenger approaches this upper bound as citizens become patient. In the presence of a term limit, we deduce an upper bound strictly below this limit: the commitment problem of voters imposes further constraints on the effectiveness of the electoral mechanism, and patience does not produce the qualified responsiveness result. This problem does not arise in the two-period model of Subsection 3.4, where the game ends after the second period, so that voters are not subject to the temptation of replacing a hard-working incumbent with a fresh challenger.

6.1 Pure moral hazard

In the rent-seeking environment, voters prefer greater effort by politicians and are therefore modeled as a unitary actor, but we can still consider the issue of policy responsiveness: here, responsiveness corresponds to positive levels of effort chosen by politicians, as greater levels of effort generate higher expected utility for all voters. The literature on infinite-horizon models of pure moral hazard is occupied by Ferejohn's (1986) model of political agency, which differs from ours in that in his framework, an office holder observes an idiosyncratic productivity shock before choosing an unobservable action; moreover, voters are assumed to be risk neutral. He shows, among other things, that the highest equilibrium payoff of the voters is increasing in office benefit, establishing a responsiveness result, but Ferejohn does not consider the degree of responsiveness that can be achieved by large office benefits.

An issue that arises in the pure moral hazard model is that because all politicians are identical after all histories, the restriction to stationary electoral equilibria leaves only a trivial equilibrium: deferential voting implies that voters, being indifferent between the incumbent and challenger, always vote for the incumbent, so the office holder always shirks, i.e., chooses zero effort. This is just the problem posed by the anti-folk theorem, Proposition 4.3. Accordingly, in this setting, the focus on deferential strategies is relaxed to allow the voter to set an arbitrary cutoff in the space of outcomes, and then the value of the optimal cutoff is analyzed. Since the voter is in fact indifferent between the two candidates, every cutoff is time-consistent, and so this approach amounts to a selection of equilibria. Taking the larger view that voters differ in their preferences and that (by an order restriction argument) the representative voter is the median voter type, it may be unrealistic to assume that the electorate coordinates on the equilibrium that is optimal for the median voter. Nevertheless, it is a reasonable starting point for the analysis and, at any rate, establishes an important normative benchmark.

We drop Ferejohn's productivity shock and examine the median voter's optimization problem in the simple rent-seeking environment, effectively allowing the voter to commit to a cutoff before the beginning of the game; again, since the voter is always indifferent between the candidates, this commitment assumption does not rely on an outside enforcement mechanism. We find that the median voter can induce positive effort, and that *greater office-motivation leads to arbitrarily high effort levels, delivering responsive democracy* in the optimal equilibrium for the voter. The logic is simple: if the voter uses a fixed cutoff as the value of office increases, then politicians could and would win with probability approaching one by increasing their effort to arbitrarily high levels, and the optimal cutoff can do no worse than this.

The analysis of the model with a two-period term limit is qualitatively different: every finite cutoff in the space of outcomes is time-inconsistent, so that an outside enforcement mechanism is needed to generate positive levels of effort in equilibrium. The see the logic of this, suppose that politicians chose a positive level of effort in their first term in order to satisfy, in expectation, a finite cutoff of the voter. The voter's cutoff can be finite in equilibrium only if the expected payoff of re-electing the incumbent is at least equal to the expected payoff of a challenger. But the incumbent will always shirk in her second term, whereas the voter can select a challenger, who (by supposition) exerts a positive level of effort. In this case, the temptation of an untried challenger makes it impossible for the voter to commit to a finite cutoff, and therefore first-term politicians cannot be induced to exert positive effort. This logic extends to arbitrary, finite term limits.

We begin by departing from our maintained assumption that the distribution of the challenger's type has full support, and instead we assume that it is degenerate on a single type, $\theta = 1$, and we suppress this in payoffs and the ideal policy, writing $w(x)$ and $\hat{x}$, respectively. Since updating no longer occurs, we suppress the belief system μ in the description of equilibrium. Now, consider a strategy profile σ such that voters re-elect the incumbent if and only if $y \geqslant \bar{y}$, where y is the realized outcome and $\bar{y}$ is an arbitrary cutoff, and such that office holders mix over policy according to $\tilde{\pi}$ in each period. For most of the subsection, the cutoff will be fixed at $\bar{y}$. Then the continuation value of a challenger is simply the expected utility generated by mixtures over policy, i.e., $V^C(\sigma) = \frac{1}{1-\delta}\sum_x \mathbb{E}[u(y)|x]\tilde{\pi}(x)$, and the value of the office holder's optimization problem net of current office benefit, denoted $W(\tilde{x}, \bar{y})$, uniquely satisfies

$$W(\tilde{\pi},\bar{y}) = \max_{x\in X} w(x) + \delta\left[(1-F(\bar{y}-x))(W(\tilde{\pi},\bar{y})+\beta) + F(\bar{y}-x)\frac{\sum_{\tilde{x}}\mathbb{E}[u(y)|\tilde{x}]\tilde{\pi}(\tilde{x})}{1-\delta}\right],$$

where we drop the subscript from the utility function w, and we add the politician's office benefit for the current period as a separate term. The first order condition for this problem is of course

$$w'(x) + f(\bar{y}-x)\delta\left[W(\tilde{\pi},\bar{y}) + \beta - \frac{\sum_x \mathbb{E}[u(y)|x]\tilde{\pi}(x)}{1-\delta}\right] \quad = \quad 0. \tag{23}$$

Let $\hat{x}$ denote the ideal policy of the politician.

As in Subsection 3.4, we can consider the constrained version of the first-term politicians' optimization problem with objective function

$$U(x,r;\tilde{\pi}) \quad = \quad w(x) + r\delta\left(W(\tilde{\pi},\bar{y}) + \beta - \frac{\sum_{\tilde{x}}\mathbb{E}[u(y)|\tilde{x}]\tilde{\pi}(\tilde{x})}{1-\delta}\right),$$

which we now explicitly parameterize by the mixed policy strategy $\tilde{\pi}$. We encounter some difficulty in reformulating (C3) in the present setting, because the mixture $\tilde{\pi}$ is endogenous. We must therefore phrase the condition so that the payoff from holding office is positive for all possible values that $\tilde{\pi}$ might take, but policy is unbounded, and so the voter's expected utility from $\tilde{\pi}$ is, in principle, unbounded. Of course, if the payoff from holding office is negative, then politicians will not mix over high policies. It suffices for the equilibrium analysis to specify the condition for the case in which politicians choose their ideal policy with probability one,

(C3) $$\text{for all } \bar{y},\ W(\hat{x},\bar{y}) + \beta - \frac{\mathbb{E}[u(y)|\hat{x}]}{1-\delta} \quad > \quad 0,$$

which holds if β is sufficiently large. In particular, (C3) is satisfied if politician preferences are obtained from voter preferences by a cost term that is not too large, i.e., $v(x) = \mathbb{E}[u(y)|x]$ and $\beta > c(\hat{x})$.

Proposition 6.1 *In the infinite-horizon model of pure moral hazard, (C3) is satisfied if for all x,* $v(x) = \mathbb{E}[u(y)|x]$ *and* $\beta > c(\hat{x})$.

To see the result, suppose that (C3) does not hold. Note that given cutoff $\overline{y}$, the value of the politicians' problem is at least equal to the discounted expected payoff from choosing the ideal policy $\hat{x}$. Since (C3) does not hold, a lower bound for the latter is obtained by the discounted expected payoff if the politician is always re-elected after choosing $\hat{x}$. Then we have

$$\begin{aligned} W(\hat{x},\overline{y}) + \beta - \frac{\mathbb{E}[u(y)|\hat{x}]}{1-\delta} &\geqslant \frac{w(\hat{x})+\beta}{1-\delta} - \frac{\mathbb{E}[u(y)|\hat{x}]}{1-\delta} \\ &= \frac{\mathbb{E}[u(y)|\hat{x}] - c(\hat{x}) + \beta}{1-\delta} - \frac{\mathbb{E}[u(y)|\hat{x}]}{1-\delta} \\ &> 0, \end{aligned}$$

a contradiction. This establishes (C3), as desired.

The objective function $U(x,r;\tilde{\pi})$ is concave, but in the corresponding optimization problem,

$$\begin{gathered} \max_{(x,r)} U(x,r;\tilde{\pi}) \\ \text{s.t. } r - 1 + F(\overline{y} - x) \leqslant 0, \end{gathered}$$

the constraint inherits the natural non-convexity of the distribution function F, paralleling difficulties in the basic probabilistic voting model and the two-period moral hazard model; see Figures 1 or 5. This leads to the possibility of multiple optimal policies.

Because an office holder takes the mixture $\tilde{\pi}$ used by politicians in the future as given in her optimization problem, and because her payoff depends on $\tilde{\pi}$ through the continuation value of a challenger, politicians are engaged in a dynamic *game*—we cannot treat it simply as a dynamic programming problem—and non-convexities necessitate the analysis of equilibria in mixed strategies. This difficulty could be assumed away by setting the payoff of an out of office politician equal to zero, but we maintain the assumption that politicians return to the electorate after their political careers have ended, consistent with the citizen-candidate approach to elections.

We deal with the problem of mixing by again assuming (C4), which implies that for every cutoff $\overline{y}$ and every mixture $\tilde{\pi}$, if the net value of office is positive, i.e.,

$$W(\tilde{\pi},\overline{y}) + \beta - \frac{\sum_x \mathbb{E}[u(y)|x]\tilde{\pi}(x)}{1-\delta} \quad > \quad 0,$$

then the objective function of the politician has at most two local maximizers and, therefore, at most two maximizers; the proof proceeds exactly as that for Proposition 3.4. We let $x^*(\tilde{\pi},\overline{y})$ and $x_*(\tilde{\pi},\overline{y})$ denote the greatest and least optimal policy choices, respectively. By standard continuity arguments, the functions $x^*(\tilde{\pi},\overline{y})$ and $x_*(\tilde{\pi},\overline{y})$ are upper semi-continuous and lower semi-continuous, respectively. Of course, any mixture over these optimal policies is optimal, and since the optimal policies themselves depend on the expected mixture $\tilde{\pi}$, we see that even in the simple model of pure moral hazard, an equilibrium in the game among politicians must solve a fixed point problem. This, in turn, raises the issues of existence and uniqueness of equilibrium.

Fortunately, both issues can be resolved by elementary arguments. Our analysis is facilitated by the following result, which shows if the policy mixture used by politicians is degenerate on a policy $\tilde{x}$, then the value of office is decreasing in $\tilde{x}$. In the statement of the next proposition, we write the value function $W(\tilde{x},\overline{y})$ as a function of a policy choice $\tilde{x}$, rather than the mixed policy strategy that places probability on on that choice.

Proposition 6.2 *In the infinite-horizon model of pure moral hazard, given cutoff $\overline{y}\in\mathbb{R}$, the expression*

$$W(\tilde{x},\overline{y})-\frac{\mathbb{E}[u(y)|\tilde{x}]}{1-\delta}$$

is decreasing in $\tilde{x}$.

To prove the result, we differentiate the Bellman equation with respect to $\tilde{x}$ to obtain

$$\frac{\partial W}{\partial \tilde{x}}(\tilde{x},\overline{y}) \quad = \quad \delta(1-F(\overline{y}-x))\frac{\partial W}{\partial \tilde{x}}(\tilde{x},\overline{y})+\frac{\delta F(\overline{y}-x)}{1-\delta}\frac{d\mathbb{E}[u(y)|\tilde{x}]}{d\tilde{x}},$$

where x is any solution to the politician's problem. Solving, we obtain

$$\frac{\partial W}{\partial \tilde{x}}(\tilde{x},\overline{y}) \quad = \quad \left(\frac{\delta F(\overline{y}-x)}{1-\delta+\delta F(\overline{y}-x)}\right)\frac{d}{d\tilde{x}}\frac{\mathbb{E}[u(y)|\tilde{x}]}{1-\delta}.$$

The coefficient in parentheses is strictly between zero and one, which implies that the expression

$$W(\tilde{x},\overline{y})-\frac{\mathbb{E}[u(y)|\tilde{x}]}{1-\delta}$$

is strictly decreasing in $\tilde{x}$, as claimed.

The above proposition implies that if other politicians choose policy $\tilde{x}$, then an office holder's optimal policy choices decrease with $\tilde{x}$, for when $W(\tilde{x},\overline{y})-\frac{1}{1-\delta}\mathbb{E}[u(y)|\tilde{x}]$ decreases, the politicians' indifference curves become steeper, reflecting the relatively greater weight placed on current policy. More formally, if x is optimal given $\tilde{x}$ and x' is optimal given $\tilde{x}'>\tilde{x}$, then $x'>x$. If there are multiple optimal policy choices given $\tilde{x}$, then an increase in $\tilde{x}$ may lead to discontinuities

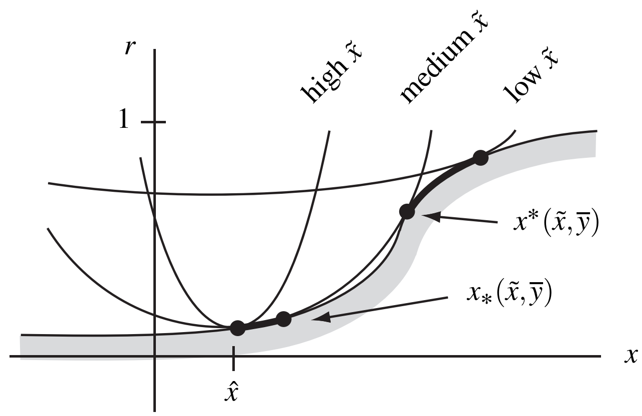

Figure 12: Multiple optimal policies

in optimal policies. As long as the net value of office is strictly positive, however, there are at most two optimal policy choices, and such a discontinuity can only occur if there is a unique optimal policy as we approach $\tilde{x}$ from below, and the optimal policies jump down at $\tilde{x}$. and decreasing in $\tilde{x}$; see Figure 12.

Now, by (C3), the net value of office is strictly positive when politicians choose the ideal policy $\hat{x}$. Thus, by inspection of the first order condition in (23), politicians optimally exert positive effort given $\hat{x}$, i.e., $x_*(\hat{x},\overline{y}) > \hat{x}$. Increasing the policy used by politicians to $\tilde{x} > \hat{x}$, we see that the "best response" policy either decreases continuously until it crosses the 45° line, in which case the unique equilibrium in the game among politicians is in pure strategies, or it jumps across the 45° line, in which case we rely on mixed strategies. Specifically, if $\tilde{x}$ is the value at which the best response policy jumps across the diagonal, then there is a least optimal policy, denoted x_*, and a greatest optimal policy, denoted x^*. We identify the mixed strategy equilibrium in the game among politicians as the mixture π^* over x^* and x_* such that

$$\pi^*(x^*)\mathbb{E}[u(y)|x^*(\tilde{x},\overline{y})] + \pi^*(x_*)\mathbb{E}[u(y)|x_*] \quad = \quad \mathbb{E}[u(y)|\tilde{x}],$$

so that the Bellman equation for the politicians is the same given π^* as it is given $\tilde{x}$; see Figure 13. If the net value of office is strictly positive at $\tilde{x}$, then by (C4) these are the only optimal policy choices, and the equilibrium is unique.

This argument proves existence of equilibrium, and in the following result, we establish that the net value of office is indeed strictly positive in every equilibrium, delivering uniqueness as well. In the sequel, we let $\pi^*(\overline{y})$ be the unique equilibrium mixture over policies; we may write $x^*(\overline{y})$ and $x_*(\overline{y})$ for the greatest and least optimal policy choices of the politicians, respectively; and we let $r^*(\overline{y})$ be the

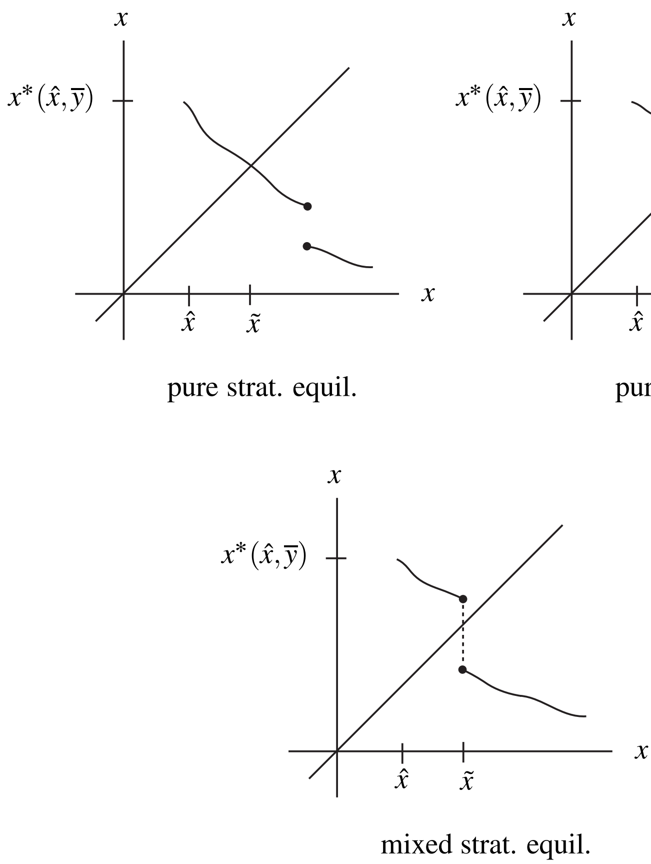

Figure 13: Existence and uniqueness

probability that a politician chooses the greatest optimal policy. Finally, we write $\mathbb{E}[u(y)|\pi]$ as shorthand for the expected utility $\sum_x \mathbb{E}[u(y)|x]\pi(x)$.

Proposition 6.3 *In the infinite-horizon model of pure moral hazard, assume (C1)–(C4). Then for every finite cutoff $\overline{y} \in \mathbb{R}$, there is a unique equilibrium mixed policy strategy $\pi^*(\overline{y})$ in the game among politicians, and in equilibrium, the net value of holding office is positive, i.e.,*

$$W(\pi^*(\overline{y}),\overline{y}) + \beta - \frac{\mathbb{E}[u(y)|\pi^*(\overline{y})]}{1-\delta} \quad > \quad 0,$$

and $\pi^*(\overline{x})$ places positive probability on at most two policies, say x^* and x_*, where $\hat{x} < x_* \leqslant x^*$. Moreover, $\pi^*(\overline{y})$ is a continuous function of the cutoff.

We have already discussed existence. Now let π^* be any equilibrium in the game among politicians. To prove that the net value of office is positive, suppose otherwise. By inspection of the first order condition in (23), it follows that politicians mix over policies less than or equal to the ideal policy, i.e., $x^*(\bar{y}) \leqslant \hat{x}$. Since an office holder can always choose the ideal policy, we have

$$W(\pi^*,\bar{y}) \quad \geqslant \quad w(\hat{x}) + \delta\left[(1-F(\bar{y}-\hat{x}))(W(\pi^*,\bar{y})+\beta) + F(\bar{y}-\hat{x})\frac{\mathbb{E}[u(y)|\pi^*]}{1-\delta}\right].$$

This implies that

$$\begin{aligned} & W(\pi^*,\bar{y}) + \beta - \frac{\mathbb{E}[u(y)|\pi^*]}{1-\delta} \\ & \quad \geqslant \quad w(\hat{x}) + \beta - \mathbb{E}[u(y)|\pi^*] + \delta(1-F(\bar{y}-\hat{x}))\left[W(\pi^*,\bar{y}) + \beta - \frac{\mathbb{E}[u(y)|\pi^*]}{1-\delta}\right], \end{aligned}$$

and thus we have

$$\begin{aligned} W(\pi^*,\bar{y}) + \beta - \frac{\mathbb{E}[u(y)|\pi^*]}{1-\delta} \quad & \geqslant \quad \frac{w(\hat{x}) + \beta - \mathbb{E}[u(y)|\pi^*]}{1-\delta(1-F(\bar{y}-\hat{x}))} \\ & \geqslant \quad \frac{w(\hat{x}) + \beta - \mathbb{E}[u(y)|\hat{x}]}{1-\delta(1-F(\bar{y}-\hat{x}))} \\ & > \quad 0, \end{aligned}$$

where the second inequality follows from $x^*(\bar{y}) \leqslant \hat{x}$, and the third from (C3). Thus, the net value of office is positive, a contradiction. We conclude that the net value of office is indeed positive in equilibrium, and the first order condition implies that optimal policies are strictly greater than the ideal policy $\hat{x}$. By (C4), this implies that given any equilibrium π^*, the politician has at most two optimal policies, and then uniqueness follows from above arguments. Continuity follows from standard arguments.

We next confirm the intuitive result that in equilibrium, politicians become worse off as voters become more demanding.

Proposition 6.4 *In the infinite-horizon model of pure moral hazard, assume (C1)–(C4). Then for any finite cutoff $\bar{y} \in \mathbb{R}$, the politicians' indirect utility, $W(\bar{y}) \equiv W(\pi^*(\bar{y}),\bar{y})$, is decreasing in $\bar{y}$.*

To prove the proposition, suppose that for $y' < y''$, we have $W(y') < W(y'')$. We write $\pi' = \pi^*(y')$ and $\pi'' = \pi^*(y'')$. Note that the expected utility from π' is less than the expected utility from π'', i.e., $\mathbb{E}[u(y)|\pi'] < \mathbb{E}[u(y)|\pi'']$; otherwise, the politicians' objective function for (π', y') exceeds that for (π'', y'') for all values of x, contradicting $W(y') < W(y'')$. It follows that

$$W(\pi'', y'') - \frac{\mathbb{E}[u(y)|\pi'']}{1-\delta} \quad > \quad W(\pi', y') - \frac{\mathbb{E}[u(y)|\pi']}{1-\delta},$$

and therefore

$$[W(\pi'', y'') - W(\pi'', y')] + [W(\pi'', y') - W(\pi', y')] \quad > \quad \frac{\mathbb{E}[u(y)|\pi'']}{1-\delta} - \frac{\mathbb{E}[u(y)|\pi']}{1-\delta}.$$

Since $W(\pi'', y'') < W(\pi'', y')$, this implies

$$W(\pi'', y') - W(\pi', y') \quad > \quad \frac{\mathbb{E}[u(y)|\pi'']}{1-\delta} - \frac{\mathbb{E}[u(y)|\pi']}{1-\delta}. \tag{24}$$

Now let $\tilde{x}$ satisfy the Bellman equation for (π'', y'). Rewriting (24) as

$$W(\pi'', y') - \frac{\mathbb{E}[u(y)|\pi'']}{1-\delta} \quad > \quad W(\pi', y') - \frac{\mathbb{E}[u(y)|\pi']}{1-\delta},$$

it follows that $\tilde{x} > x^*(y')$. Note that

$$\begin{aligned}
&W(\pi'', y') - W(\pi', y') \\
&\quad = \quad w(\tilde{x}) + \delta\left[(1 - F(y' - \tilde{x}))(W(\pi'', y') + \beta) + F(y' - \tilde{x})\frac{\mathbb{E}[u(y)|\pi'']}{1-\delta}\right] \\
&\qquad\quad -w(x^*(y')) - \delta\left[(1 - F(y' - x^*(y')))(W(\pi', y') + \beta)\right. \\
&\qquad\qquad\qquad\qquad \left. + F(y' - x^*(y'))\frac{\mathbb{E}[u(y)|\pi']}{1-\delta}\right].
\end{aligned}$$

Using $\hat{x} < x^*(y') < \tilde{x}$ and $W(\pi', y') + \beta > \frac{1}{1-\delta}\mathbb{E}[u(y)|\pi']$, this implies

$$\begin{aligned}
W(\pi'', y') - W(\pi', y') \quad < \quad & \delta(1 - F(y' - \tilde{x}))(W(\pi'', y') - W(\pi', y')) \\
& + \delta F(y' - \tilde{x})\left[\frac{\mathbb{E}[u(y)|\pi'']}{1-\delta} - \frac{\mathbb{E}[u(y)|\pi']}{1-\delta}\right].
\end{aligned}$$

Simplifying, this is

$$W(\pi'', y') - W(\pi', y') \quad < \quad \left(\frac{F(y' - \tilde{x})}{1 - \delta + \delta F(y' - \tilde{x})}\right)\left[\frac{\mathbb{E}[u(y)|\pi'']}{1-\delta} - \frac{\mathbb{E}[u(y)|\pi']}{1-\delta}\right],$$

and since the coefficient in parentheses is strictly between zero and one, this contradicts (24). This establishes the monotonicity result.

To this point, we have taken the cutoff $\overline{y}$ as exogenously fixed, but we can endogenize $\overline{y}$ by allowing voters to set this cutoff optimally, i.e., voters solve

$$\max_{\overline{y}} \mathbb{E}[u(y)|\pi^*(\overline{y})].$$

It is important to note that increasing the cutoff has two effects. First, there is a direct effect on the probability of re-election for any given policy choice x; diagrammatically, this has the effect of shifting the politicians' constraint to the right. Second, there is an indirect effect on the net value of office, $W(\pi^*(\overline{y}),\overline{y}) - \frac{1}{1-\delta}\mathbb{E}[u(y)|\pi^*(\overline{y})]$, which can be positive or negative, and this can is reflected in the marginal rate of substitution of the politician's objective function $U(x,r;\pi^*(\overline{y}))$. These effects can in turn lead to a change in the equilibrium policies, $x^*(\overline{y})$ and $x_*(\overline{y})$, and they can change the mixing probabilities over these policies.

Next, we note that the voters do indeed have an optimal cutoff, and we give a partial characterization in terms of first order conditions.

Proposition 6.5 *In the infinite-horizon model of pure moral hazard, assume (C1)–(C4). Then there is an optimal cutoff for the voters, which solves*

$$\max_{\overline{y}} \mathbb{E}[u(y)|\pi^*(\overline{y})],$$

and for every optimal cutoff y^*, if $x^*(\cdot)$ and $x_*(\cdot)$ are differentiable at y^* and satisfy the second order condition with strict inequality at y^*, if $r^*(\cdot)$ is differentiable at y^*, and if $W(\cdot)$ is differentiable at y^*, then

$$\frac{dr^*}{d\overline{y}}(y^*)[\mathbb{E}[u(y)|x^*] - \mathbb{E}[u(y)|x_*]] \;=\; \alpha^*\frac{d\mathbb{E}[u(y)|x^*]}{dx} + \alpha_*\frac{d\mathbb{E}[u(y)|x_*]}{dx},$$

where

$$\begin{aligned}
\alpha^* &= r^*(y^*)\cdot\frac{\delta f'(y^*-x^*)\Delta(y^*) + \delta f(y^*-x^*)\Delta'(y^*)}{w''(y^*) - \delta f'(y^*-x^*)\Delta(y^*)}\\
\alpha_* &= (1-r^*(y^*))\cdot\frac{\delta f'(y^*-x_*)\Delta(y^*) + \delta f(y^*-x_*)\Delta'(y^*)}{w''(y^*) - \delta f'(y^*-x_*)\Delta(y^*)}\\
\Delta(\overline{y}) &= W(\overline{y}) + \beta - \frac{\mathbb{E}[u(y)|\pi^*(\overline{y})]}{1-\delta}\\
x^* &= x^*(y^*)\\
x_* &= x_*(y^*).
\end{aligned}$$

Recall the first order condition for the politicians' optimization problem: $x^*(\overline{y})$ and $x_*(\overline{y})$ solve

$$w'(x) + f(\overline{y}-x)\delta\left[W(\overline{y}) + \beta - \frac{\mathbb{E}[u(y)|\pi^*(\overline{y})]}{1-\delta}\right] \;=\; 0.$$

For existence of an optimal cutoff, fix any $\overline{y}$. Recall that the term in brackets above is positive by Proposition 6.3, so that $\hat{x} < x_*(\overline{y})$. Note that, as in the proof of Proposition 3.7, equilibrium policies are bounded above by any policy $\overline{x}$ such that $w(\hat{x}) > w(\overline{x}) + \frac{\delta}{1-\delta}[w(\overline{x}) + \beta]$. Now consider arbitrarily low values of the cutoff, and note that $f(y - x^*(y)) \to 0$ as $|y| \to \infty$. From the first order condition, this implies that $w'(x^*(y)) \to 0$ as $|y| \to \infty$, which implies that $x^*(y) \to \hat{x}$. Thus, we can choose a sufficiently large interval $[y_L, y_H]$ such that for all $\tilde{y}$ outside the interval, we have $x^*(\tilde{y}) < x_*(\overline{y})$, which implies $\mathbb{E}[u(y)|\pi^*(\tilde{y})] < \mathbb{E}[u(y)|\pi^*(\overline{y})]$. Thus, we can restrict the optimal cutoff problem of the voter to the compact set $[y_L, y_H]$, and by continuity of the objective function, a solution exists.

To deduce the necessary condition, we insert $x^*(\overline{y})$ into the first order condition and differentiate at y^* to obtain

$$\begin{aligned} & w''(x^*)\frac{dx^*}{d\overline{y}}(y^*) + \delta f'(y^* - x^*)\left(1 - \frac{dx^*}{d\overline{y}}(y^*)\right)\left[W(y^*) + \beta - \frac{\mathbb{E}[u(y)|x^*]}{1-\delta}\right] \\ & \quad + \delta f(y^* - x^*)\left[W'(y^*) - \frac{1}{1-\delta}\frac{d}{dx}\mathbb{E}[u(y)|x^*]\frac{dx^*}{d\overline{y}}(y^*)\right] \\ = & \quad 0. \end{aligned}$$

Solving for $\frac{dx^*}{d\overline{y}}(y^*)$, we find that

$$\frac{dx^*}{d\overline{y}}(y^*) = \frac{-\delta f'(y^* - x^*)\Delta(y^*) - \delta f(y^* - x^*)\Delta'(y^*)}{w''(x^*) - \delta f'(y^* - x^*)\Delta(y^*)},$$

and similarly for $\frac{dx_*}{d\overline{y}}(y^*)$. Note that since the second order condition holds strictly, the denominator of the above expression is negative. Now turning to the voters' maximization problem,

$$\max_{\overline{y}} r^*(\overline{y})\mathbb{E}[u(y)|x^*(\overline{y})] + (1 - r^*(\overline{y}))\mathbb{E}[u(y)|x_*(\overline{y})],$$

the first order condition is

$$\begin{aligned} & \frac{dr^*}{d\overline{y}}(y^*)[\mathbb{E}[u(y)|x^*] - \mathbb{E}[u(y)|x_*]] + r^*(y^*)\frac{d\mathbb{E}[u(y)|x^*]}{dx}\frac{dx^*}{d\overline{y}}(y^*) \\ & \quad + (1 - r^*(y^*))\frac{d\mathbb{E}[u(y)|x_*]}{dx}\frac{dx_*}{d\overline{y}}(y^*) = 0. \end{aligned}$$

Substituting the above expressions for $\frac{dx^*}{d\overline{y}}$ and $\frac{dx_*}{d\overline{y}}$, we obtain the desired result.

The characterization in Proposition 6.5 has an implication for the possibility that the optimal cutoff induces the office holder to choose policies below y^*, shifted by the mode of $f(\cdot)$, which we denote $\hat{z}$. Suppose that in the equilibrium induced by an optimal cutoff y^*, the optimal policies of the politician both fall below $y^* - \hat{z}$, i.e., $y^* - x^*(y^*) > \hat{z}$. Because the second order condition holds strictly at $x^*(y^*)$, it follows that

$$w''(y^*) - \delta f'(y^* - x^*(y^*))\Delta(y^*) \quad < \quad 0,$$

and by the MRLP, $f(\cdot)$ is single-peaked, so $f'(y^* - x^*(y^*)) \leqslant 0$. Let us strengthen this slightly to assume that the derivative of $f(\cdot)$ is strictly negative when evaluated at outcomes greater than the mode. Moreover, we have shown that $\Delta(y^*) > 0$ and $W'(y^*) \leqslant 0$, and it follows that α^* and α_* are positive. Then the characterization implies

$$\frac{dr^*}{dy}(y^*)[\mathbb{E}[u(y)|x^*] - \mathbb{E}[u(y)|x_*]] \quad > \quad 0.$$

That is, an increase in the cutoff from the optimal level y^* decreases the effort exerted by politicians at both optimal policies, and so it must be that voters are compensated by a shift in probability from the lower optimal policy to the greater one. In particular, we must have $x^*(y^*) > x_*(y^*)$. We conclude that if the cutoff y^* induces an equilibrium in pure policy strategies, then it must be that the politicians' policy choice exceeds the mode of $f(\cdot)$. Intuitively, assuming $f(\cdot)$ is symmetric about zero, this implies that the voter optimally sets a relatively low bar, and politicians optimally respond by choosing policy that exceeds that bar: it is better to encourage the politician to jump a lower bar, rather than demoralize them by setting a bar that is difficult to achieve.

Geometrically, we can imagine the optimal cutoff problem of the voter by increasing the cutoff and sweeping out the policy choices of politicians; this is the dark locus of points in Figure 14. The maximum effort induced by choice of cutoff corresponds to a solution of the voters' problem. Note that as $\overline{y}$ varies, so does the value of office for politicians: when the value of office $W(\overline{y}) - \frac{1}{1-\delta}\mathbb{E}[u(y)|x^*(\overline{y})]$ decreases, the politicians' indifference curves become steeper. Here, we depict the case in which politicians choose a single policy x^* with probability one given the optimal cutoff, and consistent with the first order characterization, this means that x^* exceeds the cutoff, and the pair (x^*, y^*) lies on the constraint to the northeast of the inflection point.

We have not yet considered the possibility of responsive democracy in the infinite-horizon model of pure moral hazard: Proposition 6.5 establishes that the optimal cutoff (indeed, any cutoff) induces politicians to exert positive effort, but the result is silent on the level of the policy choices that can be attained. The next proposition establishes a *strong responsiveness result* when politicians are substantially office-motivated: when the office benefit β is large, the equilibrium policy

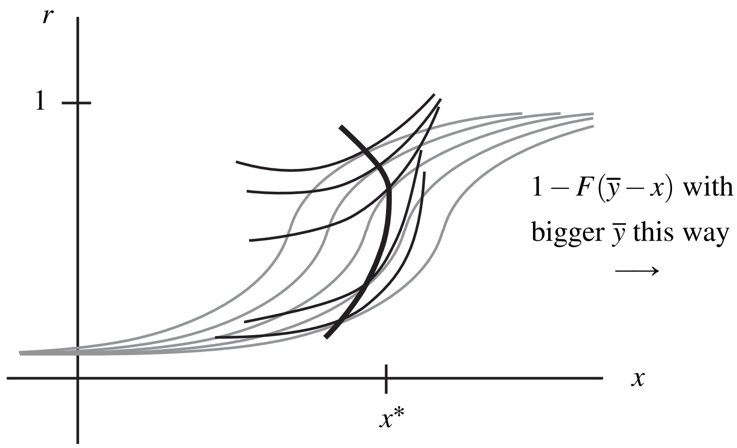

Figure 14: Locus of policy choices

choices of the politicians increase without bound for an arbitrarily fixed cutoff. Obviously, the result is then reinforced if the cutoff is set optimally.

Proposition 6.6 *In the infinite-horizon model of pure moral hazard, assume (C1)–(C4). Fix the finite cutoff $\overline{y} \in \mathbb{R}$, and let the office benefit β be arbitrarily large. Then the politicians' policy choice increases without bound, i.e.,*

$$\lim_{\beta\to\infty} x_*(\overline{y}) \quad = \quad \infty.$$

To prove the proposition, it suffices to show that there is no sequence of equilibria such that the least optimal policy $x_*(\overline{y})$ converges to a finite $\tilde{x} < \infty$ as office benefit becomes large. We simply consult the first order condition,

$$w'(x) \quad = \quad -\delta f(\overline{y}-x)\left[W(x,\overline{y}) + \beta - \frac{\mathbb{E}[u(y)|x]}{1-\delta}\right].$$

For such a sequence, the left-hand side of the equation converges to $w'(\tilde{x})$, while the right-hand side diverges to negative infinity, a contradiction.

In contrast to the positive result on responsiveness in the model with no term limit, the nature of equilibria are qualitatively different when office holders are subject to a term limit: voters cannot credibly commit to a finite cutoff, as an office holder in her final term of office is known to shirk, whereas a newly elected challenger would exert positive effort. This time consistency problem causes equilibria with finite cutoffs to unravel, leaving only the trivial equilibria in which voters always re-elect the incumbent or always elect the challenger, and thus office holders always shirk by choosing their ideal policy.

We now assume that politicians can hold office for at most K terms, and we allow voters to use cutoffs $\overline{y}_1,\ldots,\overline{y}_{K-1}$ that depend on the incumbent's term of office; that is, if the incumbent has completed their tth term, then she is re-elected if and only if the realized outcome that period satisfies $y \geqslant \overline{y}_t$. Let $V_t^I(\overline{y}_1,\ldots,\overline{y}_{K-1})$ denote the voters' expected discounted payoff from re-electing the incumbent after her tth term of office, and let $V^C(\overline{y}_1,\ldots,\overline{y}_{K-1})$ be the payoff of electing a challenger. We say $(\overline{y}_1,\ldots,\overline{y}_{K-1})$ is *time consistent* if for all $t = 1,\ldots,K-1$, the cutoff never dictates that the electorate votes against their preferences, i.e., $\overline{y}_t > -\infty$ implies

$$V^C(\overline{y}_1,\ldots,\overline{y}_{K-1}) \quad \geqslant \quad V_t^I(\overline{y}_1,\ldots,\overline{y}_{K-1}),$$

and $\overline{y}_t < \infty$ implies

$$V_t^I(\overline{y}_1,\ldots,\overline{y}_{K-1}) \quad \geqslant \quad V^C(\overline{y}_1,\ldots,\overline{y}_{K-1}).$$

Note that if $\overline{y}_t$ is finite, then time consistency implies that voters are indifferent between the incumbent and challenger, as the realized outcome can lead to a vote for either candidate.

Proposition 6.7 *In the infinite-horizon model of pure moral hazard with finite term limit, assume (C1)–(C4). If $(\overline{y}_1,\ldots,\overline{y}_{K-1})$ is time consistent, then for all t, $\overline{y}_t \in \{-\infty,\infty\}$, and in equilibrium politicians always choose $\hat{x}$.*

To see the proposition, suppose $(\overline{y}_1,\ldots,\overline{y}_{K-1})$ is time consistent but some cutoff is finite, and let $\overline{y}_t$ be the highest indexed finite cutoff. Since the office holder's policy choices in her tth or later terms of office do not affect her re-election chances, the politician simply chooses her ideal policy in all remaining terms. Letting $(\pi_1^*,\ldots,\pi_K^*)$ denote equilibrium mixed policy choices of politicians during their tenure in office, we then have $\pi_{t+1}^*(\hat{x}) = \cdots = \pi_K^*(\hat{x}) = 1$. Then we can write the payoff of re-electing the incumbent after their tth term as

$$V_t^I(\overline{y}_1,\ldots,\overline{y}_{K-1}) \quad = \quad \alpha\frac{\mathbb{E}[u(y)|\hat{x}]}{1-\delta} + (1-\alpha)V^C(\overline{y}_1,\ldots,\overline{y}_{K-1}),$$

where α is between zero and one. By time consistency and $\overline{y}_t$ finite, the left-hand side is equal to the continuation value of a challenger, and thus

$$V^C(\overline{y}_1,\ldots,\overline{y}_{K-1}) \quad = \quad \frac{1}{1-\delta}\mathbb{E}[u(y)|\hat{x}]. \tag{25}$$

By a first-order argument similar to that for Proposition 6.5, the politician in her tth term of office exerts positive effort, i.e., $\pi_t^*(\hat{x}) = 0$. Thus, $V_{t-1}^I(\bar{y}_1, \ldots, \bar{y}_{K-1}) > V^C(\bar{y}_1, \ldots, \bar{y}_{K-1})$, and time consistency implies that $\bar{y}_{t-1} = -\infty$. That is, the value of the incumbent in term $t-1$ strictly exceeds the payoff of a challenger, so voters always re-elect. Thus, in term $t-2$, the value of the incumbent again strictly exceeds the payoff of a challenger, so voters always re-elect. This argument carries back to the first term of office, and we conclude that voters always re-elect incumbents in the first $t-1$ terms of office and that politicians exert positive effort in term t. But then the payoff of a challenger strictly exceeds the discounted expected utility from the ideal policy $\hat{x}$, contradicting (25).

6.2 One-sided learning

Due to difficult theoretical issues related to updating of voter beliefs, the literature on infinite-horizon problems is small, and existence of stationary electoral equilibrium is problematic. Banks and Sundaram (1993) prove existence in history-dependent trigger strategies, which we discuss at the end of this subsection. Banks and Sundaram (1998) and Duggan (2016) establish existence in the infinite-horizon model with a two-period term limit, which we take up in the next subsection, but the question in the model with no term limits is open. Schwabe (2011) considers a simplified version of the model with no term limits in which there are two politician types and the behavior of the bad type is exogenously fixed; he shows existence of equilibria in reputation-dependent cutoffs, which allow for greater history dependence than stationarity.

We discuss some of the technical difficulties presented in a version of the model with two types, and we provide a qualified responsiveness result: as citizens become patient, the (normalized) continuation value of a challenger converges to the expected utility from the ideal policy of the high type, assuming equilibria exist. The form of this responsiveness result differs from Proposition 3.8 for the two-period model, where the strategic structure of the game implies that office-motivated politicians' effort levels increase without bound. The difference owes to the fact that whereas the voters' prior beliefs about the incumbent are fixed on the two-period model, beliefs b are used as a state variable in the infinite-horizon model with one-sided learning. In particular, given any parameters of the model, when beliefs b place probability close to one on the high politician type, an office holder will leverage this reputation by choosing near ideal policies for an arbitrarily long duration. The qualification is moot in the model of pure adverse selection, where the ideal policy of the "best type" of politician is the median ideal policy, and Proposition 5.6 delivers the responsive democracy result. We maintain assumptions (C1)–(C3), but we make only expositional use of (C4) in the present subsection.

The focus on monotonic, deferential equilibria implies that election outcomes are characterized by a cutoff, as in the two-period model of Subsection 3.4 and in the pure moral hazard model of Subsection 6.1. In contrast to previous analyses, we now explicitly write the equilibrium cutoff and policy strategies, $y^*(b)$ and $\pi^*_j(b)$, as functions of voter beliefs. We write $\mu_T(b,y)$ for the voters' beliefs about an office holder's type conditional on policy outcome y and given prior beliefs b. Since all policy outcomes are on the path of play, Bayes rule pins down the beliefs of voters given σ, and we therefore summarize a stationary electoral equilibrium by the strategy profile σ, leaving beliefs implicit.

Given strategies σ and beliefs b about the incumbent's type, we let $W_j(b|\sigma)$ denote the value of the type j politician's optimization problem, $V^C(\sigma)$ denote the continuation value of a challenger, and $V^I(b|\sigma)$ denote the continuation value of re-electing the incumbent. These uniquely satisfy the following functional equations: for all b,

$$\begin{aligned} W_j(b|\sigma) &= \max_{x\in X} w_j(x) + \delta\Bigg[\int_{y\geqslant y^*(b)} [W_j(\mu_T(b,y)|\sigma)+\beta]f(y-x)dy \\ &\quad + F(y^*(b)-x)V^C(\sigma)\Bigg] \\ V^C(\sigma) &= \sum_j p_j \sum_x \Bigg[\mathbb{E}[u(y)|x] + \delta[\int_{y\geqslant y^*(b)} [V^I(\mu_T(p,y)|\sigma)+\beta]f(y-x)dy \\ &\quad + F(y^*(b)-x)V^C(\sigma)]\Bigg]\pi^*_j(x|p) \\ V^I(b|\sigma) &= \sum_j b_j \sum_x \Bigg[\mathbb{E}[u(y)|x] + \delta[\int_{y\geqslant y^*(b)} [V^I(\mu_T(b,y)|\sigma)+\beta]f(y-x)dy \\ &\quad + F(y^*(b)-x)V^C(\sigma)]\Bigg]\pi^*_j(x|b). \end{aligned}$$

The right-hand side of the type j politicians' Bellman equation is evidently differentiable (and therefore continuous) in the policy choice x, and so an optimal choice exists and satisfies the first order condition

$$\begin{aligned} -w'_j(x) &= \delta\Bigg[\int_{y\geqslant y^*(b)} [W_j(\mu_T(b,y)|\sigma)+\beta]f'(y-x)dy \\ &\quad + f(y^*(b)-x)V^C(\sigma)\Bigg]. \end{aligned} \tag{26}$$

The indifference condition determining the voters' cutoff is

$$V^I(\mu_T(b,y)|\sigma) = V^C(\sigma), \tag{27}$$

given that the voters' prior beliefs at the beginning of the period are b.

We first note that in equilibrium, it is possible for a first-term office holder to be re-elected and to be rejected. In contrast to the subsequent analysis, this minimal step does not rely on the assumption of two types. Indeed, if this were not the case, then every politician type would choose their ideal policy in the first term of office, but then (C1) and (C2) imply that voters strictly prefer to re-elect the incumbent following sufficiently high policy outcomes.

Proposition 6.8 *In the infinite-horizon model of adverse selection and one-sided learning, assume (C1) and (C2) hold. In every stationary electoral equilibrium σ, we have $-\infty < y^*(p) < \infty$.*

Suppose that in equilibrium, $y^*(p) = \infty$. Because first-term office holders cannot be re-elected, every type chooses her ideal policy, so the (normalized) continuation value of a challenger is $(1-\delta)V^C(\sigma) = \sum_j p_j \mathbb{E}[u(y)|\hat{x}_j]$. Since $\hat{x}_n$ exceeds all other ideal policies, it follows that for sufficiently large outcomes y, the voters' posterior over the incumbent's type becomes arbitrarily close to degenerate on type n. That is, $\mu_T(n|p,y) \to 1$ as $y \to \infty$. Then the voters' expected discounted payoff from re-electing the incumbent satisfies

$$\begin{aligned} V^I(\mu_T(p,y)|\sigma) &\geqslant \mu_T(n|p,y)\mathbb{E}[u(y)|\hat{x}_n] + (1-\mu_T(n|p,y))\mathbb{E}[u(y)|\hat{x}_1] + \delta V^C(\sigma) \\ &\to \mathbb{E}[u(y)|\hat{x}_n] + \delta V^C(\sigma) \\ &> V^C(\sigma) \end{aligned}$$

as y becomes large, but then the incumbent is re-elected, i.e., $y \geqslant y^*(p)$, a contradiction. An analogous argument, focusing on low policy outcomes rather than large, rules out $y^*(p) = -\infty$.

Dependence of W_j and V^I on beliefs b introduces significant complications over the pure moral hazard model, reflecting the possibility that voters learn about an incumbent's type as the game evolves. Of note, voter beliefs become a state variable in the dynamic electoral game, in which politicians and voters condition their actions directly on b, and the transition on beliefs depends on the politicians' choices. Such dependence creates the possibility that strategies and continuation values are discontinuous in beliefs—it may be that $x_j^*(b)$ and $x_{*,j}(b)$ jump in response to a discontinuity in $y^*(b)$, and reciprocally, that the cutoff jumps in response to a discontinuity in policy choices—with ensuing difficulties for equilibrium existence. Furthermore, the updated belief $\mu_T(b,y)$ is a highly non-linear function, and the composition $W_j(\mu_T(b,y)|\sigma)$ is potentially badly behaved. Finally, the transition on the state variable b, which is given by the Bayesian posterior $\mu_T(b,y)$, depends implicitly on strategies σ, i.e., it is endogenous. These technical issues combine to present formidable challenges to the analysis of the general model.

Digression on existence

To facilitate the discussion of equilibrium existence, we temporarily simplify further by assuming a maximum feasible policy, $\overline{x}$, by assuming one type has zero cost, and by assuming (C4). This means that the "good" type has essentially the same preferences as voters and that the policy choice of such politicians is pinned down at the maximum, i.e., the unique optimal policy for all voter beliefs is $x_2^*(b) = \overline{x}$, so the main question concerns the policy choice of the lower type. Assuming for the sake of discussion that the payoff $V^I(b|\sigma)$ from an incumbent is increasing in the probability b_2 of the high type, the voters' cutoff from (27) is just the solution to $\mu_T(2,y) = p_2$, i.e., the cutoff is such that conditional on $y^*(b)$, the probability the incumbent is the high type is just equal to the prior probability. This means that $y^*(b)$ solves the equation

$$p_2 = \frac{b_2 f(y-\overline{x})}{b_1 \sum_x f(y-x)\pi_1(x|b) + b_2 f(y-\overline{x})}, \tag{28}$$

or after manipulating,

$$\frac{b_1 \sum_x f(y-x)\pi_1(x|b)}{b_2 f(y-\overline{x})} = \frac{p_1}{p_2}.$$

For the special case in which $\pi_1(b)$ is degenerate on a single policy $x_1^*(b) < \overline{x}$ and the density $f(\cdot)$ is normal with mean zero, this simplifies further to

$$y^*(b) = \frac{\ln(\frac{p_1 b_2}{p_2 b_1})}{2(x_1^*(b)-\overline{x})} + \frac{x_1^*(b)+\overline{x}}{2}.$$

We see that even in this very special case of the model, the cutoff is unbounded, non-linear in policy choices, and ostensibly non-monotonic in beliefs. The exception is in the first term of office, where $b = p$, in which case the cutoff reduces to the midpoint between the maximum policy and the choice of the type 1 politicians.

Assuming that π_1 places positive probability on policies strictly less than $\overline{x}$, the MLRP implies that equation (28) for the voters' cutoff has a unique solution, say $\zeta(b|\pi_1)$. Then we can rewrite the Bellman equation for the type 1 politicians as parameterized by the policy strategy π_1, as in

$$\begin{aligned} W_1(b|\pi_1) = \max_{x\in X} w_j(x) + \delta\Bigg[& \int_{y \geqslant \zeta(b|\pi_1)} [W_1(\mu_T(b,y)|\pi_1) + \beta] f(y-x)dy \\ & + F(\zeta(b|\pi_1)-x)V^C(\pi_1)\Bigg], \end{aligned}$$

where we use a similar convention in writing $V^C(\pi_1)$. The right-hand side of the Bellman equation is continuous in the policy choice x, so a maximizer exists for each b, but note that the Bellman equation itself depends on π_1 through the cutoff and (implicitly) through Bayesian updating. Then π_1 corresponds to an equilibrium if for all b, the distribution $\pi_1(b)$ places probability one on solutions to the politician's Bellman equation.

To discuss the equilibrium existence issue in more detail, we focus on the case in which the type 1 politicians use a pure strategy, which we represent by a policy function $x_1(\cdot)$ that specifies a policy choice $x_1(b)$ as a function of the voters' beliefs; the technical issues are equally germane to the case of mixed strategy equilibria. In line with the above observations, we pursue the following route to existence: given a policy strategy $x_1(\cdot)$, we derive a new policy strategy $\tilde{x}_1(\cdot)$ from the solutions to the Bellman equation; *if* these functions can be restricted a priori to a compact space, and *if* the mapping $x_1(\cdot) \mapsto \tilde{x}_1(\cdot)$ is continuous, then it has a fixed point; and this fixed point yields an equilibrium.

A number of remarks are in order. First, our ability to restrict policy strategies to a compact space in a strong topology hinges on deriving an a priori limitation on the variation of these policy functions. In the literature on dynamic games, this is normally done by imposing structure on the exogenously given transition probability on the state variable (in this case, b), but this course is not available in the current setting: the transition on beliefs is dictated by Bayes rule and therefore itself depends (implicitly) on the strategy used by type 1 politicians. Second, continuity of the fixed point mapping hinges critically on joint continuity of the right-hand side of the Bellman equation on the policy choice x and the policy strategy $x_1(\cdot)$ itself. The Bellman equation is continuous in x alone, but the joint continuity condition is much more demanding and depends on the notion of convergence applied to policy strategies. Third, we do not have conditions under which optimal policies are ensured to be unique, so the mapping suggested above relies on taking a selection from the correspondence of optimal policies. This correspondence may not have convex values, and it might not even be possible to find a continuous selection. Because voters will not be indifferent between the politicians' optimal policies, this multiplicity would normally necessitate mixing by the type 1 politicians. Then the domain of the fixed point argument will not be functions $x_j(\cdot)$ from beliefs to policies, but rather mappings $\pi_j(\cdot)$ from beliefs to probability distributions over policies, i.e., transition probabilities of the form $\pi_j\colon \Delta(T) \to \Delta(X)$. We consider in this discussion the optimistic case in which this problem does not arise.[22]

[22]Fourth, the standard topology on the space of such transition probabilities, called the "narrow topology," delivers compactness but is too weak to ensure the joint continuity condition needed for the fixed point argument.

In the current setting, we can circumvent some of these issues by accepting discontinuities in the strategy of type 1 politicians, sacrificing uniform convergence, and giving the space of policy functions the weak* topology as a subset of $L^\infty([0,1])$, where we now view $x_1^*(b_2)$ as a function of the probability of type 2 alone.[23] This space of functions is indeed compact. To avoid discontinuities in voter payoffs, we simplify the model further so that a politician receives a zero payoff when removed from office, whereby the type 1 politicians' Bellman equation reduces to

$$W_1(b|x_1(\cdot)) \quad = \quad \max_{x \in X} w_1(x) + \delta \int_{y \geqslant \zeta(b|x_1(\cdot))} [W_1(\mu_T(b,y)|x_1(\cdot)) + \beta] f(y-x) dy.$$

Then, given a policy strategy $x_1(\cdot)$, we can simply choose the greatest optimal policy for each belief b, giving us a new policy strategy $\tilde{x}_1(\cdot)$, which will belong to $L^\infty([0,1])$. The last stumbling block to existence—which appears fundamental—is the potential discontinuity of the Bellman equation with respect to the mapping from $x_1(\cdot)$ to $\tilde{x}_1(\cdot)$, stemming from the poor pointwise properties of weak* convergence.

To expand on this point, let x' and x'' be any two policies, and let $\{x_1^k(\cdot)\}$ be a sequence of policy strategies that alternates between x' and x'' at an increasing rate; Figure 15 depicts the initial strategies in the sequence, which is known in analysis as the Rademacher sequence. This sequence converges in the weak* topology to the strategy $x_1(\cdot)$ that always chooses the midpoint $(x'+x'')/2$. For almost all beliefs b, the politicians' choice switches between x' and x'' infinitely often, and so the current period payoffs $w_1(x_1^k(b))$ along the sequence will not converge to the payoff from $(x'+x'')/2$. Moreover, the voters' cutoff switches between $y'(b)$ and $y''(b)$ infinitely often along the sequence, where in the normal case these cutoffs solve

$$\frac{b_1 f(y-x')}{b_2 f(y-\overline{x})} = \frac{p_1}{p_2} \quad \text{and} \quad \frac{b_1 f(y-x'')}{b_2 f(y-\overline{x})} = \frac{p_1}{p_2},$$

respectively. Clearly, these cutoffs will not generally converge to the cutoff for $x_1(\cdot)$. Finally, dependence of the Bayesian posterior $\mu_T(b,y)$ on the politicians' policy strategy is also problematic. Note that conditional on an outcome y, the voters' updated beliefs switch infinitely often between

$$b_2' = \frac{b_2 f(y-\overline{x})}{b_1 f(y-x') + b_2 f(y-\overline{x})} \quad \text{and} \quad b_2'' = \frac{b_2 f(y-\overline{x})}{b_1 f(y-x'') + b_2 f(y-\overline{x})}$$

along the sequence, and these updated beliefs will not converge to the updated belief for $x_1(\cdot)$. These considerations appear to lead inevitably to a failure of joint

[23] This means that a sequence $\{x_1^k(\cdot)\}$ of policy strategies converges to $x_1(\cdot)$ if and only if for every integrable function $g\colon [0,1] \to \mathbb{R}$, the integrals $\int g(b_2) x_1^k(b_2) db_2$ converge to $\int g(b_2) x_1(b_2) db_2$.

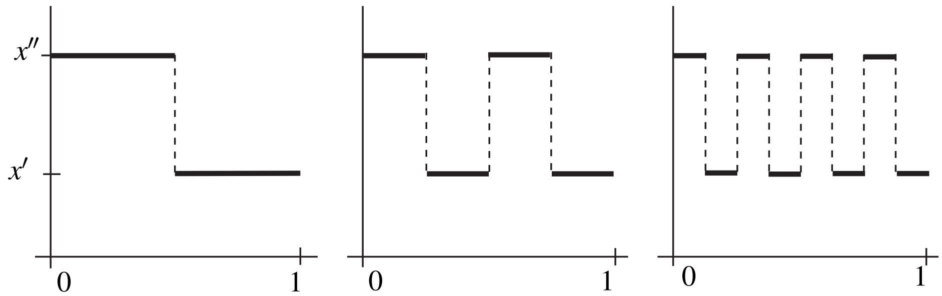

Figure 15: Problematic policy strategies

continuity and the impracticality of the fixed point approach. In sum, existence of stationary electoral equilibrium is a thorny issue.

End of digression

To provide some insight into the model, we specialize to the model with two types, which we refer to as the *infinite horizon model of moral hazard and adverse selection with two types,* and our characterization results will sidestep the technical challenges alluded to above, and we implicitly *assume equilibria exist.* To avoid perverse incentives of office holders in the subsequent analysis, we reformulate (C3) as follows,

$$\text{(C3)} \qquad w_1(\hat{x}_1)+\beta \quad > \quad \mathbb{E}[u(y)|\hat{x}_2],$$

which means that a first term politician prefers to remain in office, even if she can return to the electorate and in all future periods, outcomes are determined by the ideal policy of the high type. We assume provisionally that $(1-\delta)V^C(\sigma) \leqslant \mathbb{E}[u(y)|\hat{x}_2]$, and we will see that $\mathbb{E}[u(y)|\hat{x}_2]$ does indeed bound the (normalized) continuation value of a challenger in equilibrium, so the new (C3) means that all politicians are in principle interested in re-election. Even under (C3), however, characterization of equilibrium is challenging. At a minimum, we might expect that in equilibrium, the high type (low cost) politician chooses a higher policy than the low type. This will be true if the Bellman equation $W_j(b|\sigma)$ is increasing in the probability b_2 that the incumbent is the high type. This, in turn, seems natural: when a politician is perceived as more likely to be the higher type, she has greater "political capital," and the politician can decrease effort while maintaining a good reputation. We will say that a stationary equilibrium σ is *reputation monotonic* if for all y, $W_j(\mu_T(b,y)|\sigma)$ is weakly increasing in b_2 for both types.

Now, under (C3), we can show that the right-hand side of the first order condition in (26) is negative, and it follows that given any reputation-monotonic equilibrium of the two-type model, politicians exert positive effort. These claims are established in the next proposition.

Proposition 6.9 *In the infinite-horizon model of adverse selection and one-sided learning with two types, assume (C1)–(C3). Then for all types $j = 1,2$, all beliefs b, and every reputation-monotonic stationary electoral equilibrium σ such that $(1-\delta)V^C(\sigma) \leqslant \mathbb{E}[u(y)|\hat{x}_n]$, we have*

$$\int_{y \geqslant y^*(b)} [W_j(\mu_T(b,y)|\sigma) + \beta] f'(y-x)dy + f(y^*(b)-x)V^C(\sigma) \quad > \quad 0,$$

and thus $\pi_j((\hat{x}_j,\infty)) = 1$.

To prove the proposition, note that $f(y^*(b)-x) = -\int_{y^*(b)}^{\infty} f'(y-x)dy$. Thus, we can write the desired inequality as

$$\begin{aligned} & \frac{d}{dx}\int I_{y^*(b)}(y)[W_j(\mu_T(b,y)|\sigma) + \beta - V^C(\sigma)] f(y-x)dy \\ = \quad & \int I_{y^*(b)}(y)[W_j(\mu_T(b,y)|\sigma) + \beta - V^C(\sigma)] f'(y-x)dy \\ > \quad & 0, \end{aligned}$$

where $I_{y^*(b)}(y)$ is an indicator function taking a value of one if $y \geqslant y^*(b)$ and zero otherwise. By (C3) and the assumption that σ is reputation monotonic, the expression $W_j(\mu_T(b,y)|\sigma) + \beta - V^C(\sigma)$ is weakly increasing on $y \geqslant y^*(b)$. If we can show that $W_j(b|\sigma) + \beta V^C(\sigma) > 0$ for all beliefs b, then the integrand

$$I_{y^*(b)}(y)[W_j(\mu_T(b,y)|\sigma) + \beta - V^C(\sigma)]$$

is weakly increasing and not constant. Moreover, an increase in the policy choice x determines a first order stochastic increase in the distribution of policy outcomes, and this distribution has full support, so the inequality must hold.

Thus, it suffices to show that for all b, we have $W_j(b|\sigma) + \beta > V^C(\sigma)$. To simplify the proof, assume there is some $\underline{b}$ that minimizes $W_j(\underline{b}|\sigma)$.[24] Then because a type j politician can always choose the ideal policy $\hat{x}_j$, we have

$$W_j(\underline{b}|\sigma) + \beta \quad \geqslant \quad w_j(\hat{x}_j) + \beta + \delta\left[\int_{y \geqslant y^*(\underline{b})} [W_j(\mu_T(\underline{b},y)|\sigma) + \beta] f(y-\hat{x}_j)dy\right.$$

[24] In general, we can work with beliefs for which the infimum of the value function is approximated.

$$
\begin{aligned}
& +F(y^*(\underline{b})-\hat{x}_j)V^C(\sigma)\Big] \\
\geqslant\ & w_j(\hat{x}_j)+\beta+\delta\Big[(1-F(y^*(\underline{b})-\hat{x}_j)[W_j(\underline{b}|\sigma)+\beta] \\
& +F(y^*(\underline{b})-\hat{x}_j)V^C(\sigma)\Big].
\end{aligned}
$$

This implies that

$$
\begin{aligned}
W_j(\underline{b}|\sigma)+\beta &\geqslant \frac{w_j(\hat{x}_j)+\beta+\delta F(y^*(\underline{b})-\hat{x}_j)V^C(\sigma)}{1-\delta(1-F(y^*(\underline{b})-\hat{x}_j))} \\
&> \frac{(1-\delta)V^C(\sigma)+\delta F(y^*(\underline{b})-\hat{x}_j)V^C(\sigma)}{1-\delta(1-F(y^*(\underline{b})-\hat{x}_j))} \\
&= V^C(\sigma),
\end{aligned}
$$

where the strict inequality uses (C3) and $(1-\delta)V^C(\sigma) \leqslant \mathbb{E}[u(y)|\hat{x}_n]$. Since the net value of office is minimized at $\underline{b}$, this establishes the desired inequality. We have shown that the right-hand side of the politicians' first order condition (26) is positive, and we conclude that each politician type exerts positive effort in equilibrium. This completes the proof.

As another small step in understanding the two-type model, we exploit supermodularity of the politicians' payoffs to deduce that the policy choices of office holders are always ordered strictly by type.

Proposition 6.10 *In the infinite-horizon model of adverse selection and one-sided learning with two types, assume (C1) and (C2) hold. In every reputation-monotonic stationary electoral equilibrium* σ*, the type 1 politicians' policy choices are strictly less than the type 2 politicians', i.e., for all b and all x and z, if* $\pi_1^*(x|b) > 0$ *and* $\pi_2^*(z|b) > 0$*, then* $x < z$*.*

To prove the result, write the right-hand side of the first order condition for the type j office holder's optimization problem as

$$
\frac{d}{dx}\int I_{y^*(b)}(y)[W_j(\mu_T(b,y)|\sigma)+\beta-V^C(\sigma)]f(y-x)dy \quad > \quad 0,
$$

following the proof of Proposition 6.9. Then the fact that $W_2(b|\sigma) > W_1(b|\sigma)$ for all beliefs implies that the right-hand side is higher for the type 2 politician than for the type 1, which implies $x < z$.

The next result establishes that in the context of the two-type model, the voters' continuation value of a challenger is bounded above by the discounted expected utility from the ideal policy of the highest type. A similar result is proved by Schwabe (2011) under the assumptions that the bad type of politician shirks and that the increment to the voter's utility from a good politician (independent of effort) is sufficiently small. This result has no parallel in the model of pure adverse selection in the spatial environment, where the ideal policy of the "best type" coincides with the ideal policy of the median voter. It also reveals a fundamental difference between the two-period and infinite-horizon models of moral hazard and adverse selection, as in the former model, the continuation value of a challenger increases without bound as politicians become office motivated. It is tempting to suppose that the strong responsiveness result carries over to the infinite-horizon model, but we find that the removal of the terminal period imposes constraints on the effectiveness of electoral incentives.

The proof of the bound on responsiveness relies on the fact that the voters' beliefs about the incumbent's type are now viewed as a state variable, and when voters place probability close to one on the high politician type, the office holder will leverage this reputation by choosing near ideal policies for an arbitrarily long duration. The voters' discounted expected payoff at such a state is close to the discounted expected payoff from the politician's ideal policy. Finally, an office holder who is the high type with probability close to one must be re-elected, which gives us the bound. An implication of the bound, with Proposition 6.9, is that the net value of office is indeed positive in equilibrium.

Proposition 6.11 *In the infinite-horizon model of adverse selection and one-sided learning with two types, assume (C1)–(C3) hold. For all levels of office benefit* $\beta \geqslant 0$ *and all discount factors* $\delta \in [0,1)$*, in every reputation-monotonic stationary electoral equilibrium* σ*, the voters' continuation value of a challenger is no more than the expected discounted utility from the ideal policy of the type* 2 *politician, i.e.,*

$$V^C(\sigma) \quad \leqslant \quad \frac{\mathbb{E}[u(y)|\hat{x}_2]}{1-\delta}.$$

To prove the result, fix $\beta \geqslant 0$ and $\delta < 1$, so that optimal policy choices can be bounded by an interval $[0,\overline{x}]$. By Proposition 6.8, we have $y^*(p) < \infty$, and by Proposition 6.10, we have $x < z$ for all x and z with $\pi_1^*(x|b) > 0$ and $\pi_2^*(z) > 0$ and for all b. We claim that as b_2 approaches one, the voters' cutoff decreases without bound, i.e., $y^*(b) \to -\infty$. To see this, let $\tilde{y}(b)$ be the solution to

$$p_2 \quad = \quad \frac{b_2 \sum_x f(y-x)\pi_2(x|b)}{b_1 \sum_x f(y-x)\pi_1(x|b) + b_2 \sum_x f(y-x)\pi_2(x|b)},$$

which, with (C1) and (C2), is uniquely defined by Proposition 6.10. Therefore, $\mu_T(b,\tilde{y}(b)) = p$, which implies that $V^I(b,\tilde{y}(b)|\sigma) = V^I(p|\sigma) = V^C(\sigma)$. Because σ

is deferential and monotonic, we then have $y^*(b) \leqslant \tilde{y}(b)$, so it suffices to show that $\tilde{y}(b) \to -\infty$ as $b_2 \to 1$. By construction of $\tilde{y}(b)$, we have

$$\frac{b_2 \sum_x f(\tilde{y}(b)-x)\pi_2(x|b)}{b_1 \sum_x f(\tilde{y}(b)-x)\pi_1(x|b)} = \frac{p_2}{p_1},$$

and thus the likelihood ratio $\frac{\sum_x f(\tilde{y}(b)-x)\pi_2(x|b)}{\sum_x f(\tilde{y}(b)-x)\pi_1(x|b)}$ must converge to zero. Because policy choices belong to a compact set, it follows that the values $|\tilde{y}(b)|$ must be divergent. And if there is a subsequence such that $\tilde{y}(b) \to \infty$, then because $f(\cdot)$ is single-peaked, using Proposition 6.10, we have

$$\frac{\sum_x f(\tilde{y}(b)-x)\pi_2(x|b)}{\sum_x f(\tilde{y}(b)-x)\pi_1(x|b)} \geqslant \frac{\sum_x f(\tilde{y}(b)-x)\pi_1(x|b)}{\sum_x f(\tilde{y}(b)-x)\pi_1(x|b)} = 1,$$

a contradiction. Thus, $\tilde{y}(b) \to -\infty$, as claimed.

Note that Bayesian updating $\mu_T(b,y)$ becomes insensitive to the outcome y when the belief b is close to degenerate; in particular, as $b_2 \to 1$, the function $W_j(\mu_T(b,y)|\sigma)$ becomes arbitrarily close to constant in y. Since $y^*(b) \to -\infty$, by the previous claim, this implies that the right-hand side of the type 2 politicians' first order condition in (26) converges to zero uniformly on $[0,\overline{x}]$. It follows that the policy choice of type 2 politicians converges to their ideal policy, i.e., $x_2^*(b) \to \hat{x}_2$ as $b_2 \to 1$. Moreover, the updated beliefs of the voters in the next term of office will be arbitrarily close to the initial b_2 for all outcomes outside a set of arbitrarily small measure. By the same argument, the office holder's policy choice in her second term also converges to the ideal policy, and so on for an arbitrarily long horizon. Since voters are not perfectly patient ($\delta < 1$), this implies that the expected payoff from re-electing the incumbent, conditional on the politician being type 2, converges to the discounted expected utility from the ideal policy, i.e., $V^I(b,2|\sigma) \to \frac{1}{1-\delta}\mathbb{E}[u(y)|\hat{x}_2]$ as $b_2 \to 1$. We therefore have

$$V^C(\sigma) \leqslant V^I(b|\sigma) \leqslant \frac{b_1\mathbb{E}[u(y)|\overline{x}]}{1-\delta} + b_2 V^I(b,2|\sigma) \to \frac{\mathbb{E}[u(y)|\hat{x}_2]}{1-\delta}$$

as $b_2 \to 1$. We conclude that $V^C(\sigma) \leqslant \frac{1}{1-\delta}\mathbb{E}[u(y)|\hat{x}_2]$, as desired.

We can give a fairly loose lower bound on the continuation value of an incumbent for arbitrary parameters of the model. Given stationary electoral equilibrium policy strategies, suppose voter beliefs are $b = (b_1, b_2)$, and consider the discounted expected payoff from the voting strategy that simply re-elects the incumbent always, i.e., $y(b) \equiv -\infty$. By Proposition 6.9, the incumbent's policy choice always exceeds her ideal policy, and thus the voters' payoff from this voting strategy exceeds $b_1\mathbb{E}[u(y)|\hat{x}_1]/(1-\delta) + b_2\mathbb{E}[u(y)|\hat{x}_2]/(1-\delta)$. By the principle of optimality for dynamic elections, Proposition 4.2, the voters' equilibrium expected payoff cannot fall below this bound, which delivers the following result.

Proposition 6.12 *In the infinite-horizon model of adverse selection and one-sided learning with two types, assume (C1)–(C3) hold. In every reputation-monotonic stationary electoral equilibrium* σ*, the voters' expected discounted payoff from re-electing an incumbent given beliefs* b *satisfies*

$$V^I(b|\sigma) \quad \geqslant \quad \frac{b_1\mathbb{E}[u(y)|\hat{x}_1]}{1-\delta} + \frac{b_2\mathbb{E}[u(y)|\hat{x}_2]}{1-\delta}.$$

The next proposition shows that as voters become patient, the continuation value from a challenger converges to the upper bound established in Proposition 6.11, giving us a *qualified responsive democracy result* in the infinite-horizon model with no term limit. Note that the extent of responsiveness differs from the result stated in Proposition 3.8 for the two-period model. There, we assume no discounting and let office benefit become large; here, we fix office benefit and let citizens become patient. There, policy choices of above average types increase without bound, and the continuation value of a challenger becomes arbitrarily; here, responsiveness is constrained by the preferences of the potential candidates for election. Like Proposition 6.12, the proof uses an optimality principle argument. As citizens become patient, a possible voting strategy for the representative voter is to remove an office holder after her first term, unless the posterior probability that she is the high type is very close to one; in the latter case, the incumbent is retained thereafter, regardless of the observed outcome. Conditional on being re-elected, an incumbent will then, with probability close to one, choose policies above $\hat{x}_2$. The cost of this strategy is the potentially long delay until the outcome realized following a first-term office holder's policy choice is sufficiently high, but as voters become patient, this cost goes to zero, and the limit is achieved.

Proposition 6.13 *In the infinite-horizon model of adverse selection and one-sided learning with two types, assume (C1)–(C3) hold, and let the discount factor approach one. For every selection of reputation-monotonic stationary electoral equilibria* σ*, the voters' (normalized) continuation value of a challenger converges to the expected utility from the ideal policy of the type 2 politician, i.e.,*

$$\lim_{\delta\to 1}(1-\delta)V^C(\sigma) \quad = \quad \mathbb{E}[u(y)|\hat{x}_2].$$

By Proposition 6.11, it suffices to show that there is no subsequence such that the limit of (normalized) continuation values is strictly less than the expected utility from $\hat{x}_2$. To this end, suppose that $\lim_{\delta\to 1}(1-\delta)V^C(\sigma) < \mathbb{E}[u(y)|\hat{x}_2]$. Let $\eta \in (0,1)$ be small enough that

$$\eta\mathbb{E}[u(y)|\hat{x}_1] + (1-\eta)\mathbb{E}[u(y)|\hat{x}_2] \quad > \quad \lim_{\delta\to 1}(1-\delta)V^C(\sigma).$$

We show that voters can improve their expected payoff by waiting to re-elect a first-term incumbent until the realized policy outcome after the first term is such that the voters' updated belief that the incumbent is type 2 exceeds $1-\eta$, and then re-electing the politician regardless of realized outcomes thereafter. Since policies are strictly ordered by type, (C1) and (C2) yield a cutoff $\overline{y}(\delta)$ such that $\mu(2|b,\overline{y}) = 1-\eta$ in the equilibrium corresponding to discount factor δ. A technical issue that must be addressed is the possibility that $y(\delta)$ diverges to infinity; we must show that it is uniformly bounded as $\delta \to 1$. Proposition 6.10 implies that the supports of $\pi_1(\cdot|b)$ and $\pi_2(\cdot|b)$ are ordered by type, and thus the voters' updated belief is

$$\mu_T(2|p,y) \quad = \quad \frac{p_2 \sum_x f(y-x)\pi_2(x|p)}{p_1 \sum_x f(y-x)\pi_1(x|p) + p_2 \sum_x f(y-x)\pi_2(x|p)}$$

conditional on observing y in a politician's first term of office. A potential problem that arises is that the equilibrium policy choices of the politicians converge to the same level, in which updating diminishes and higher outcomes are needed to move the voters' beliefs, or the policy choices diverge to infinity.

Since β is fixed, the type j politician prefers to choose $\hat{x}_j$ than x in her first term of office if

$$w_j(\hat{x}_j) + \frac{\delta \mathbb{E}[u(y)|\hat{x}_1]}{1-\delta} \quad > \quad \frac{w_j(x)+\beta}{1-\delta},$$

which holds uniformly across δ if $w_j(x) < \min\{w_j(\hat{x}_j), \mathbb{E}[u(y)|\hat{x}_1]\} - \beta$. This inequality indeed holds for sufficiently high x, and we conclude that equilibrium policy choices will belong to some interval $[0,\overline{x}]$ for all δ.

Furthermore, we claim that as $\delta \to 1$, the optimal policy choices of the type 2 politician are greater than those of the type 1 politician by an increment that does not go to zero. Let $x_{*,2}(\delta)$ denote the least optimal policy for the type 2 politician during her first term of office in the equilibrium corresponding to δ, and let $x_1^*(\delta)$ denote the greatest optimal policy for the type 1 politician. Suppose toward a contradiction that $|x_{*,2}(\delta) - x_1^*(\delta)| \to 0$, and going to a subsequence if necessary, assume $\lim_{\delta\to 1} x_{*,2}(\delta) = \lim_{\delta \to 1} x_1^*(\delta) = \tilde{x}$. Then $\tilde{x} \geqslant \hat{x}_2 > \hat{x}_1$, and it follows that the right-hand side of the type 1 politician's first order condition in (26) converges to a positive number, say $\alpha > 0$. But the left-hand side is smaller in magnitude for the type 2 politician, as $c'(\tilde{x})/\theta_1 > c'(\tilde{x})/\theta_2$, and the right-hand side is larger, contradicting optimality for the type 2 politician. This establishes the claim. Let $\varepsilon > 0$ be such that for δ sufficiently close to one, the equilibrium policy choices of the type 2 politician exceed those of the type 1 politician by at least ε.

Define Π to consist of all pairs (π_1, π_2) with supports ordered according to type and separated by an increment of at least ε. This is a compact set.[25] For each

[25]Here, and in the remainder of this argument, we impose the topology induced by the weak*

$(\pi_1, \pi_2) \in \Pi$, define the function

$$\mu(y|\pi_1,\pi_2) = \frac{p_2 \sum_x f(y-x)\pi_2(x|p)}{p_1 \sum_x f(y-x)\pi_1(x|p) + p_2 \sum_x f(y-x)\pi_2(x|p)}.$$

This is a jointly continuous function of (y, π_1, π_2). By (C1) and (C2), the problem

$$\begin{array}{c} \min_y y \\ \text{s.t. } \mu(y|\pi_1,\pi_2) \geqslant 1-\eta \end{array}$$

has a solution, say $y(\pi_1, \pi_2)$, and by the maximum theorem (see Border (1985), Theorem 12.1), the solution varies continuously as a function of (π_1, π_2). Since Π is compact, the function $y(\pi_1, \pi_2)$ attains a maximum on Π, which we denote $\hat{y}$. We conclude that for δ sufficiently close to one, the updated beliefs of the voters that the incumbent is type 2, conditional on observing $y \geqslant \hat{y}$ in the first term of office, is at least $1-\eta$, i.e., $\mu_T(2|p,y) \geqslant 1-\eta$.

Now consider the voters' payoff, given a newly elected office holder, from the following plan: replace the incumbent with a challenger unless the realized outcome exceeds $\hat{y}$, in which case re-elect the incumbent thereafter. For an arbitrary pair $(\pi_1, \pi_2) \in \Pi$ of mixed policy strategies, the probability that the realized outcome exceeds this threshold in a politician's first term of office is

$$\rho(\pi_1,\pi_2) = p_1 \sum_x (1-F(\hat{y}-x))\pi_1(x|p) + p_2(1-F(\hat{y}-x))\pi_2(x|p) > 0.$$

This function is continuous, and thus it achieves a positive minimum on Π, say $\hat{\rho} > 0$. In particular, the voter's equilibrium expected discounted payoff from following this plan, normalized by $1-\delta$, is at least

$$\sum_{k=1}^{\infty} (1-\hat{\rho})^{k-1}\hat{\rho}\left[(1-\delta^{k-1})\mathbb{E}[u(y)|\hat{x}_1] + \delta^{k-1}[\eta\mathbb{E}[u(y)|\hat{x}_1] + (1-\eta)\mathbb{E}[u(y)|\hat{x}_2]]\right].$$

In the above expression, the sum is over strings of length k of outcomes below $\hat{y}$. The expected discounted payoff to the voters is at least the payoff from the type 1 politician's ideal policy for the first $k-1$ period; and conditional on realizing $y > \hat{y}$, the politician is type 2 with probability at least $1-\eta$, and the policy choice of the incumbent is at least equal to $\hat{x}_2$ thereafter.

---

topology on the space of signed Borel measures on $[0,\bar{x}]$; we impose the product topology on the space of pairs (π_1, π_2) of Borel probability measures on $[0,\bar{x}]$; and we give Π the relative topology induced by the product topology.

For each k, the term in large brackets converges to the quantity $\eta\mathbb{E}[u(y)|\hat{x}_1] + (1-\eta)\mathbb{E}[u(y)|\hat{x}_2]$, and the weights $(1-\hat{\rho})^{k-1}\hat{\rho}$ sum to one. Therefore, the voters' (normalized) expected discounted payoff converges to

$$\eta\mathbb{E}[u(y)|\hat{x}_1] + (1-\eta)\mathbb{E}[u(y)|\hat{x}_2] \quad > \quad \lim_{\delta\to 1}(1-\delta)V^C(\sigma).$$

But by the optimality principle for dynamic elections, Proposition 4.2, the equilibrium continuation value of a challenger, $V^C(\sigma)$, must be at least equal to the expected discounted payoff from following the above plan, a contradiction. This completes the proof.

Departing from our restriction to stationary electoral equilibria, Banks and Sundaram (1993) show existence of an equilibrium in the class of trigger strategies, in which voters and politicians use history-dependent strategies that condition on past outcomes generated by an incumbent (which are always on her personal path of play) and not only on the voters' posterior beliefs. In particular, if the realized policy outcome falls below a given cutoff level during a politician's term, the politician shirks (i.e., chooses zero effort) thereafter, and the voter removes the incumbent from office. This approach is not without its shortcomings. First, even if the incumbent is a good type with arbitrarily high probability, there is always a positive probability that a bad outcome will be realized and the voter will replace the incumbent. Second, the exact value of the trigger is not pinned down in the model, and in fact a continuum of values can be supported in equilibrium. Third, the analysis relies on the assumption that all politician types are equivalent when they shirk; without this assumption, the trigger strategy construction breaks down, as voters may have an incentive to re-elect an incumbent who is a good type with high probability, even if it is known that she will shirk in the future.

6.3 One-sided learning with term limits

In the infinite-horizon model with a two-period term limit, Banks and Sundaram (1998) extend the two-period model of moral hazard and adverse selection by imposing a two-period term limit. We review the existence question and state an existence result from Duggan (2016), and under (C4), we again obtain equilibria in which each politician type mixes over at most two policies, analogous to Proposition 3.3. Equilibria in the presence of a two-period term limit are, however, qualitatively different than those in the simple version of the model with no term limit and those in the two-period model: because voters cannot commit to decline the option of an untried challenger, high levels of effort cannot be supported in equilibrium. Duggan (2016) shows that as politicians become office motivated, the voters' expected utility from the policy choices of first-term office holders (and therefore the expected payoff of a challenger) is bounded above by a level that is below the expected utility when the highest type, $j = n$, chooses her ideal policy, $\hat{x}_n$. This upper bound on responsiveness is similar to the result in Proposition 6.11 for the two-type model of moral hazard and adverse selection with no term limit, but the qualified responsiveness demonstrated in Proposition 6.13 for the model without term limits fails in the model with term limits—a negative conclusion that parallels Proposition 5.8 for the pure adverse selection model with term limit. Thus, the commitment problem of voters has a significant impact on the effectiveness of electoral incentives, and this impact appears to be magnified in the presence of term limits.

In the model with a two-period term limit, we maintain (C1)–(C4) and extend our definition of stationary strategy profile to allow for politicians to condition their choices on the term of office, as obviously, a second term politician will simply choose her ideal policy; we let π_j^1 denote the type j politician's mixed policy choice in her first term of office. With this modification, a profile σ that is deferential and monotonic determines an acceptance set of the form $A(\sigma) = [\overline{y}, \infty)$, where $\overline{y}$ is a given cutoff outcome that is necessary and sufficient for re-election after an office holder's first term. Then stationary electoral equilibria are characterized by three conditions. First, the cutoff outcome must satisfy the indifference condition that, conditional on observing $\overline{y}$, voters are indifferent between re-electing the first-term incumbent and electing a challenger. Formally, letting $V^C(\sigma)$ be the continuation value of electing a challenger common to all voters, this condition is

$$\sum_j \mu_T(j|p,\overline{y}) \left[\mathbb{E}[u(y)|\hat{x}_j] + \delta V^C(\sigma) \right] \quad = \quad V^C(\sigma).$$

Simplifying, this means that the expected utility from the incumbent's policy choice in the second term is equal to the (normalized) continuation value of a challenger:

$$\sum_j \mu_T(j|p,\overline{y}) \mathbb{E}[u(y)|\hat{x}_j] \quad = \quad (1-\delta)V^C(\sigma). \tag{29}$$

Second, each politician type, knowing that she is re-elected if and only if $y \geqslant \overline{y}$, mixes over optimal actions in her first term of office, i.e., she solves

$$\max_{x \in X} w_j(x) + \delta \left[(1 - F(\overline{y} - x))[w_j(\hat{x}_j) + \beta + \delta V^C(\sigma)] + F(\overline{y} - x)V^C(\sigma) \right],$$

and the first order condition for this problem is

$$w'_j(x) \quad = \quad -f(\overline{y} - x)\delta[w_j(\hat{x}_j) + \beta - (1-\delta)V^C(\sigma)]. \tag{30}$$

This parallels the politicians' objective function in the two-period model of Subsection 3.4, the difference being that the payoff from re-election is discounted and reflects the continuation of the game with a challenger taking office, and the politician's expected payoff from a challenger is endogenized in a more complex way. Third, as always, updating of voter beliefs follows Bayes rule: conditional on observing outcome y, the posterior probability that the politician is type j is

$$\mu_T(j|p,y) \quad = \quad \frac{p_j \sum_x f(y-x)\pi_j^1(x)}{\sum_k p_k \sum_x f(y-x)\pi_k^1(x)}.$$

As in Subsection 3.4, we can consider the constrained version of the first term politician's optimization problem with objective function

$$U_j(x,r;V) \quad = \quad w_j(x) + r\delta[w_j(\hat{x}_j) - (1-\delta)V],$$

with the difference that we now include a parameter V, which in equilibrium will be the continuation value of a challenger. We use the formulation of (C3) from the preceding subsection extended to an arbitrary number of types,

$$\text{(C3)} \qquad w_1(\hat{x}_1) + \beta \quad > \quad \mathbb{E}[u(y)|\hat{x}_n],$$

which means that a first-term politician prefers to remain in office, even if she can return to the electorate and in all future periods, outcomes are determined by the ideal policy of the highest type. We assume provisionally that $(1-\delta)V \leqslant \mathbb{E}[u(y)|\hat{x}_n]$, and we will see that $\mathbb{E}[u(y)|\hat{x}_n]$ does indeed bound the (normalized) continuation value of a challenger in equilibrium, so the new (C3) means that all politicians are in principle interested in re-election. By (C3), the right-hand side of the first order condition in (30) is positive, and it follows that for arbitrary cutoff and continuation value satisfying the above bound, a politician exerts positive effort in the first term.

As in Subsection 3.4, an office holder will not choose policies below her ideal policy, and the politician will not choose arbitrarily high policies, so each type of office holder has an optimal policy in the first term of office. We impose condition (C4) to obtain the result that given an arbitrary cutoff and a continuation value $V \leqslant \frac{1}{1-\delta}\mathbb{E}[u(y)|\hat{x}_n]$, each type of politician has at most two optimal policies. Again, the objective function $U_j(x, 1-F(\bar{y}-x); V)$ is supermodular in (j,x), with the implication that optimal policies are strictly ordered by type. The arguments for these results proceed as for the two-period model, and we omit their formal statement and proof. Given cutoff $\bar{y}$ and continuation value of a challenger $V \leqslant \frac{1}{1-\delta}\mathbb{E}[u(y)|\hat{x}_n]$, we let $x_j^*(\bar{y},V)$ and $x_{*,j}(\bar{y},V)$ denote the greatest and least optimal policies, respectively, of the type j politician in the first term of office. It follows that optimal policies are strictly ordered by type, i.e.,

$$\text{for all } j < n, \qquad x_j^*(\overline{y},V) < x_{*,j+1}(\overline{y},V),$$

and standard arguments imply that $x_j^*(\cdot)$ and $x_{*,j}(\cdot)$ are upper and lower semi-continuous, respectively.

Now consider mixed policy strategies $\pi_1^1,\ldots,\pi_n^1$ with supports that are strictly ordered according to type, and let V be a continuation value satisfying $\mathbb{E}[u(y)|\hat{x}_1] \leqslant (1-\delta)V \leqslant \mathbb{E}[u(y)|\hat{x}_n]$. The induced cutoff for voters is the unique solution in $\overline{y}$ to the equation $V^I(p,\overline{y}|\sigma) = V$, or more explicitly,

$$\sum_k \mu_T(k|p,\overline{y})[\mathbb{E}[u(y)|\hat{x}_k] + \delta V] \quad = \quad V,$$

reflecting the fact that a re-elected incumbent chooses her ideal policy and is replaced by a challenger. Simplifying, we obtain the indifference condition

$$\sum_k \mu_T(k|p,\overline{y})\mathbb{E}[u(y)|\hat{x}_k] \quad = \quad (1-\delta)V, \tag{31}$$

and we denote the unique solution by $y^*(\pi_1^1,\ldots,\pi_n^1,V)$. Again, $y^*(\pi_1^1,\ldots,\pi_n^1,V)$ is a continuous function of its arguments.

Existence of equilibrium follows from a fixed point argument, as in the two-period model. In contrast to Subsection 3.4, however, we must complete an intermediate step, in which mixed policy strategies $\pi_1^1,\ldots,\pi_n^1$ and a cutoff $\overline{y}$ determine the continuation values from re-electing a first-term incumbent, conditional on an outcome y, and from electing a challenger. Specifically, we define $V^*(\pi_1^1,\ldots,\pi_n^1,\overline{y})$ as the unique solution to the recursion

$$V \quad = \quad \sum_j p_j \sum_x \left[\mathbb{E}[u(y)|x] + \delta[(1-F(\overline{y}-x))(\mathbb{E}[u(y)|\hat{x}_j] + \delta V) + F(\overline{y}-x)V]\right]\pi_j^1(x),$$

or more explicitly, we define

$$V^*(\pi_1^1,\ldots,\pi_n^1,\overline{y}) \quad = \quad \frac{\sum_j p_j \sum_x \left[\mathbb{E}[u(y)|x] + \delta(1-F(\overline{y}-x))\mathbb{E}[u(y)|\hat{x}_j]\right]\pi_j^1(x)}{1-\delta\sum_j p_j \sum_x[(1-F(\overline{y}-x))\delta + F(\overline{y}-x)]\pi_j^1(x)}.$$

This raises the technical challenge that the voters' cutoff is uniquely defined only when the (normalized) continuation value of a challenger is strictly between the expected utility from the ideal policies of the lowest and highest types; yet it is possible, in principle, for $V^*(\pi_1^1,\ldots,\pi_n^1,\overline{y})$ to violate those bounds. Nevertheless, the continuation values generated by optimal policies and induced cutoffs can be appropriately bounded and a fixed point argument applied.

Existence of equilibrium requires a fixed point argument, and under (C4), a proof similar to that for the two-period model could be used. Duggan (2016) establishes existence of stationary electoral equilibria using a different proof approach in a more general model, and that paper provides a characterization familiar from Proposition 3.7. We state it next, specializing it by adding (C4) and obtaining a somewhat sharper result. Note that the equilibrium cutoff is always finite: otherwise, all types of politicians would choose their ideal policy, but then choices are ordered by type, so the cutoff must be finite after all.

Proposition 6.14 *In the infinite-horizon model of adverse selection and one-sided learning with two-period term limit, assume (C1)–(C4). Then there is a stationary electoral equilibrium, and in every electoral equilibrium, there exist mixed policy strategies $\pi_1^*, \dots, \pi_n^*$ and a finite cutoff y^* such that:*

(i) each type j politician mixes over policies in the first term of office using π_j^*, which places positive probability on at most two policies, say x_j^* and $x_{*,j}$, where $\hat{x}_j < x_{*,j} \leqslant x_j^*$,

(ii) the supports of policy strategies are strictly ordered by type, i.e., for all $j < n$, we have $x_j^* < x_{*,j+1}$,

(iii) each type j politician chooses $\hat{x}_j$ in the second term of office, if re-elected,

(iv) voters re-elect an office holder after the first term if and only if $y \geqslant y^*$.

The analysis of the infinite-horizon model with two-period term limit has so far relied on a close parallel to the two-period model. Endogeneity of the continuation value $V^C(\sigma)$ does not affect existence in an essential way, as the nature of this endogeneity is continuous, allowing the existence argument to carry over with few changes. This is not so for the qualitative nature of equilibria. Proposition 3.8 established that when the office benefit β is high in the two-period model, all politician types exert arbitrarily high effort in the first period, providing a strong policy responsiveness result for that model. The limits of responsive democracy were revised in Proposition 6.11 for the infinite-horizon model with no term limit, the upper bound being the expected payoff from the ideal policy of the high type. In the general model with a two-period term limit, we find that equilibrium policy choices are subject to the same upper bound. The reason is that second-term politicians exert zero effort in equilibrium, and so the voters' expected payoff from electing a politician to a second term is bounded above by the payoff generated by the highest ability type choosing zero effort; and if first-term politicians' effort levels are too high, then the voter would rather elect a challenger, but then politicians will exert zero effort in equilibrium. As in the pure moral hazard model, the voter cannot commit to decline the option of an untried challenger, so *electoral incentives are attenuated in the model with term limits.* In fact, the next result, due to Duggan (2016), is stated in somewhat stronger terms: the voters' expected utility from the policy choices of first-term office holders—and therefore the continuation value of a challenger—is less than the expected utility from the ideal policy of the highest type.

Proposition 6.15 *In the infinite-horizon model of adverse selection and one-sided learning with two-period term limit, assume (C1)–(C4). For all levels of office benefit* $\beta \geqslant 0$ *and all discount factors* $\delta \in [0,1)$*, in every stationary electoral equilibrium* σ*, the expected utility to voters from policies chosen by first-term office holders is no more than the expected utility from the ideal policy of the type n politician, i.e.,*

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \quad < \quad \mathbb{E}[u(y)|\hat{x}_n].$$

To prove the proposition, suppose that for some parameterization of the model, we have

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \quad \geqslant \quad \mathbb{E}[u(y)|\hat{x}_n].$$

Recall that the continuation value of a challenger satisfies

$$V^C(\sigma) \quad = \quad \sum_j p_j \sum_x \Big[\mathbb{E}[u(y)|x] + \delta[(1-F(\bar{y}-x))(\mathbb{E}[u(y)|\hat{x}_j] + \delta V^C(\sigma)) + F(\bar{y}-x)V^C(\sigma)] \Big] \pi_j^1(x). \tag{32}$$

Note that

$$\begin{aligned} &\sum_j p_j \sum_x (1-F(y^*-x))(\mathbb{E}[u(y)|\hat{x}_j] + \delta V^C(\sigma))\pi_j^1(x) \\ &= \int_{y^*}^{\infty} \left[\sum_j \left(p_j \sum_x f(\tilde{y}-x)\pi_j^1(x) \right) (\mathbb{E}[u(y)|\hat{x}_j] + \delta V^C(\sigma)) \right] d\tilde{y} \\ &= \int_{y^*}^{\infty} \left(\sum_k p_k \sum_x f(\tilde{y}-x)\pi_k^1(x) \right) \left[\sum_j \mu_T(j|p,\tilde{y})(\mathbb{E}[u(y)|\hat{x}_j] + \delta V^C(\sigma)) \right] d\tilde{y} \end{aligned}$$

$$\begin{aligned} &= \int_{y^*}^{\infty} \left(\sum_k p_k \sum_x f(\tilde{y}-x)\pi_k^1(x) \right) V^I(p,\tilde{y}|\sigma)d\tilde{y} \\ &\geqslant V^C(\sigma) \sum_k p_k \sum_x (1-F(y^*-x))\pi_k^1(x), \end{aligned}$$

where the second equality follows by multiplying and dividing by the denominator in the expression for Baye's rule. Thus, we infer from (32) that the voters' (normalized) continuation value of a challenger is at least equal to the expected utility from the policy choices of first-period office holders, i.e.,

$$(1-\delta)V^C(\sigma) \geqslant \sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x). \tag{33}$$

Combining our observations, we have $(1-\delta)V^C(\sigma) > \mathbb{E}[u(y)|\hat{x}_n]$, but the indifference condition (29) yields

$$V^C(\sigma) = \frac{\sum_j \mu_T(j|p,y^*)\mathbb{E}[u(y)|\hat{x}_j]}{1-\delta} < \frac{\mathbb{E}[u(y)|\hat{x}_n]}{1-\delta},$$

a contradiction. This establishes the result.

For the infinite-horizon model with no term limits, Proposition 6.13 establishes that the effort levels of politicians approach the upper bound from Proposition 6.11 as citizens become patient. This result fixes the office benefit at a given level and uses an optimality principle argument to deduce that the continuation value of a challenger converges to the expected utility from the ideal policy of the high type. Here, the analysis of the model with two-period term limit diverges from the model with no term limit, as the commitment problem of voters imposes further constraints on responsive democracy in the former model. We next show that for a given level of office benefit, the voters' expected utility from the policy choices of first-term office holders—and therefore the continuation value of a challenger—is bounded strictly below the expected utility from the ideal policy $\hat{x}_n$ as we vary the discount factor. In fact, the statement of the result is stronger than this, in the sense that we can allow the office benefit to become large, as long as the discount factor eventually offsets the increase in office motivation. For the result, we add the Inada condition (C5).

Proposition 6.16 *In the infinite-horizon model of adverse selection and one-sided learning with two-period term limit, assume (C1)–(C5) hold. For every constant* $c > 0$*, there is a bound* $\overline{u} < \mathbb{E}[u(y)|\hat{x}_n]$ *such that for all levels of office benefit* $\beta \geqslant 0$ *and all discount factors* $\delta \in [0,1)$ *satisfying* $\beta\delta \leqslant c$*, in every stationary electoral equilibrium* σ *for parameters* (β,δ)*, the expected utility to voters from policies* *chosen by first-term office holders is below this bound, i.e.,*

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \quad \leqslant \quad \overline{u}.$$

To deduce a contradiction, suppose there is a constant $c > 0$ and a sequence of parameters (β, δ) such that $\beta\delta \leqslant c$ and for which the voters' expected utility from the choices of first-term office holders approaches the expected utility from the ideal policy of the type n politician, i.e.,

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \;\rightarrow\; \mathbb{E}[u(y)|\hat{x}_n].$$

Note that the right-hand side of the first order condition in (30) is bounded, and thus, by (C5), we can bound the optimal policies of the politicians along the sequence by some $\overline{x}$. From the argument in the proof of Proposition 6.15, inequality (33) holds, and we infer the voters' continuation value of a challenger converges to the expected utility from the ideal policy of the type n politicians, as

$$\mathbb{E}[u(y)|\hat{x}_n] \;\geqslant\; (1-\delta)V^C(\sigma) \;\geqslant\; \sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \;\rightarrow\; \mathbb{E}[u(y)|\hat{x}_n].$$

The indifference condition (29) then implies that the posterior probability that the incumbent is type n conditional on observing y^* goes to one, i.e.,

$$\mu_T(n|p,y^*) \;=\; \frac{p_n \sum_x f(y^*-x)\pi_n^1(x)}{\sum_k p_k \sum_x f(y^*-x)\pi_k^1(x)} \;\rightarrow\; 1.$$

Because the equilibrium policy choices of the politicians belong to the compact interval $[0,\overline{x}]$, this implies that $y^* \rightarrow \infty$, and thus the probability of re-electing an incumbent goes to zero, i.e., for all politician types j, we have $\sum_x F(y^*-x)\pi_j^1(x) \rightarrow 1$. In particular, we have $f(y^*-x_j^*) \rightarrow 0$ for each type j, and thus the right-hand side of the first order condition converges to zero when evaluated at the greatest optimal policy of the politician. It follows that $x_j^* \rightarrow \hat{x}_j$ for each type j, but then

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \;\rightarrow\; \sum_j p_j \mathbb{E}[u(y)|\hat{x}_j],$$

a contradiction. This establishes the result.

To apply the previous result for a given level of office benefit, say β, we simply set $c = \beta$.

Corollary 6.1 *In the infinite-horizon model of adverse selection and one-sided learning with two-period term limit, assume (C1)–(C5) hold, and fix the office benefit* $\beta \geqslant 0$*. Then there is a bound* $\overline{u} < \mathbb{E}[u(y)|\hat{x}_n]$ *such that for all discount factors* $\delta \in [0,1)$ *and every stationary electoral equilibrium* σ*, the expected utility to voters from policies chosen by first-term office holders is below this bound, i.e.,*

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j^1(x) \quad \leqslant \quad \overline{u}.$$

The preceding corollary is reminiscent of Proposition 5.8, which established a bound on responsive democracy in the infinite-horizon model of pure adverse selection with a two-period term limit. A difference is that the latter bound holds for arbitrary levels of office benefit, while Corollary 6.1 fixes β. This difference owes to the fact that voter utility is unbounded in the model with combined adverse selection and moral hazard, leading to the possibility that as office benefit increases, type n politicians choose arbitrarily high policies with sufficiently low probability so that the (normalized) continuation value of a challenger converges to $\mathbb{E}[u(y)|\hat{x}_n]$. To satisfy our equilibrium conditions, it must then be that the voters' cutoff y^* diverges to infinity, and type $j < n$ politicians are re-elected with probability converging to zero. The next proposition makes these equilibrium conditions more precise and leaves open the possibility that the bound $\overline{u}$ in Proposition 6.16 converges to $\mathbb{E}[u(y)|\hat{x}_n]$ as the product $\beta\delta$ increases, approximating our responsiveness result, Proposition 6.13, for the model with no term limits as politicians become substantially office motivated.

Proposition 6.17 *In the infinite-horizon model of adverse selection and one-sided learning with two-period term limit, assume (C1)–(C4) hold. Let the office benefit* $\beta \geqslant 0$ *and* $\delta \in [0,1)$ *vary arbitrarily subject to* $\lim \beta\delta = \infty$*. Then for every selection of stationary electoral equilibria* σ*, the voters' cutoff diverges to infinity; the type n politicians in their first term mix between policies that are close to their ideal policy and ones that are arbitrarily high, with small, positive probability on the latter; and all other type* $j < n$ *politicians choose policies close to their ideal policies in the first term, i.e.,*

(i) $y^* \to \infty$,

(ii) $x_n^* \to \infty$ *and* $x_{*,n} \to \hat{x}_n$,

(iii) $\pi_n^1(x_n^*) > 0$ *for large enough* β*, and* $\pi_n^1(x_n^*) \to 0$,

(iv) for all $j < n$, $x_j^* \to \hat{x}_j$.

We first prove (i). Indeed, suppose there is a subsequence such that y^* is bounded, so going to a subsequence if needed, we can assume that $y^* \to \bar{y}$. We claim that for each type j politician, the least optimal policy diverges to infinity, i.e., $x_{*,j} \to \infty$. Otherwise, we can go to a subsequence if needed to assume that $x_{*,j} \to \tilde{x}_j < \infty$. By the first order condition in (30), we then have

$$w'_j(\tilde{x}_j) \;=\; \lim w'_j(x_{*,j}) \quad \leqslant \quad -f(\bar{y}-\tilde{x}_j)\lim \beta\delta - C,$$

where C is the limit infimum of $f(\bar{y}-x)\delta[w_j(\hat{x}_j)-(1-\delta)V^C(\sigma)]$, which is finite. Since $\lim \beta\delta = \infty$, the right-hand side of the above inequality has infinite magnitude, a contradiction. Thus, we have $x_{*,j} \to \infty$ for each politician type j, as claimed. But then the voters' expected utility from the policy choices of first-term office holders diverges to infinity, contradicting Proposition 6.15. We conclude that $|y^*| \to \infty$. Now suppose there is a subsequence such that $y^* \to -\infty$. Because policy choices are strictly ordered by type, it follows that for all $x_1,\ldots,x_n$ in the support of the politicians' policy strategies, we have

$$f(y^*-x_1) \;>\; f(y^*-x_2) \;>\; \cdots \;>\; f(y^*-x_n).$$

Therefore, the coefficients on prior beliefs from Bayes rule are ordered by type, i.e.,

$$\frac{\sum_x f(y^*-x)\pi^1_1(x)}{\sum_k p_k \sum_x f(y^*-x)\pi^1_k(x)} \;>\; \cdots \;>\; \frac{\sum_x f(y^*-x)\pi^1_n(x)}{\sum_k p_k \sum_x f(y^*-x)\pi^1_k(x)},$$

and we conclude that voters' prior first order stochastically dominates the posterior distribution $\mu_T(\cdot|p,y^*)$, which implies

$$\sum_j p_j \mathbb{E}[u(y)|\hat{x}_j] \quad > \quad \sum_j \mu_T(j|p,y^*)\mathbb{E}[u(y)|\hat{x}_j].$$

From the argument in the proof of Proposition 6.15, inequality (33) holds, and since $\hat{x}_j < x_{*,j}$ for all j, we have $V^C(\sigma) > \frac{1}{1-\delta}\sum_j p_j \mathbb{E}[u(y)|\hat{x}_j]$. Using (**??**), we then have

$$V^I(p,y^*|\sigma) \;<\; \sum_j p_j \Big[\mathbb{E}[u(y)|\hat{x}_k] + \delta V^C(\sigma)\Big] \;<\; V^C(\sigma),$$

contradicting the voters' indifference condition. Therefore, $y^* \to \infty$, as desired.

Next, we adapt the argument for part (ii) of Proposition 3.8 to show that for all types j, there is no subsequence of greatest optimal policy choices x^*_j that converge to a finite policy greater than the ideal policy; by the same argument, the least optimal policy choices $x_{*,j}$ also cannot converge to a finite policy greater than the ideal policy. Indeed, suppose that there is some type j such that $x_j^* \to \tilde{x}_j$ with $\hat{x}_j < \tilde{x}_j < \infty$. Then for sufficiently large β and some δ (which may depend on β), we have $\hat{x}_j < x_j^*$. For these parameters, the current gain to the type j politician from choosing $\hat{x}_j$ instead of x_j^* is non-positive, and thus we note that

$$\delta(F(y^* - \hat{x}_j) - F(y^* - x_j^*))[w_j(\hat{x}_j) + \beta - (1-\delta)V^C(\sigma)] \quad \geqslant \quad w_j(\hat{x}_j) - w_j(x_j^*).$$

That is, the current gains from choosing the ideal policy are offset by future losses. Since $y^* \to \infty$, the limit of

$$\frac{F(y^* - x_j^*) - F(y^* - \tilde{x}_j - 1)}{F(y^* - \hat{x}_j) - F(y^* - x_j^*)}$$

as β becomes large (with δ possibly depending on β) is indeterminate, and by L'Hôpital's rule, the limit is equal to

$$\lim \frac{f(y^* - x_j^*) - f(y^* - \tilde{x}_j - 1)}{f(y^* - \hat{x}_j) - f(y^* - x_j^*)} \;=\; \lim \frac{f(y^* - \tilde{x}_j - 1)\left(\frac{f(y^* - x_j^*)}{f(y^* - \tilde{x}_j - 1)} - 1\right)}{f(y^* - x_j^*)\left(\frac{f(y^* - \hat{x}_j)}{f(y^* - x_j^*)} - 1\right)} \;=\; \infty,$$

where we use (C1) and (C2). Then, however, the future gain from choosing $\tilde{x}_j + 1$ instead of x_j^* strictly exceeds current losses, i.e.,

$$\begin{aligned} &\delta(F(y^* - x_j^*) - F(y^* - \tilde{x}_j - 1))[w_j(\hat{x}_j) + \beta - (1-\delta)V^C(\sigma)] \\ &\quad > \quad w_j(x_j^*) - w_j(\tilde{x}_j + 1), \end{aligned} \tag{34}$$

for some parameters (β', δ'). To be specific, let

$$\begin{aligned} A &= \delta[w_j(\hat{x}_j) + \beta - (1-\delta)V^C(\sigma)] \\ B &= w_j(\hat{x}_j) - w_j(x_j^*) \\ C &= w_j(x_j^*) - w_j(\tilde{x}_j + 1), \end{aligned}$$

where A is evaluated at sufficiently large β (with δ possibly depending on β). Note that since $\hat{x}_j < \tilde{x}_j < \infty$, we have $\lim B > 0$ and $\lim C < \infty$. We have noted that $(F(y^* - \hat{x}_j) - F(y^* - x_j^*))A \geqslant B$ for sufficiently large β (with δ possibly depending on β), and we have shown that as β becomes large (with δ possibly depending on β), we have

$$\frac{F(y^* - x_j^*) - F(y^* - \tilde{x} - 1)}{F(y^* - \hat{x}_j) - F(y^* - x_j^*)} \quad > \quad \frac{C}{B}.$$

Combining these facts, we have

$$(F(y^* - \hat{x}_j) - F(y^* - x_j^*))A\left(\frac{F(y^* - x_j^*) - F(y^* - \tilde{x} - 1)}{F(y^* - \hat{x}_j) - F(y^* - x_j^*)}\right) \quad > \quad B\left(\frac{C}{B}\right),$$

which yields (34) for some (β', δ'). This gives the type j politician a profitable deviation from x_j^*, a contradiction.

Now, suppose there is a subsequence such that the greatest optimal policy x_n^* of the type n politicians is bounded above by some policy, say $\bar{x}$. It follows that for all politician types j, we have $x_j^* \to \hat{x}_j$, so that the probability of re-electing an incumbent goes to zero, i.e., for all politician types j, we have $\sum_x F(y^* - x)\pi_j^1(x) \to 1$. Re-writing (32), we have

$$\begin{aligned}(1-\delta)V^C(\sigma) \quad &= \quad \left(\frac{1-\delta}{1-\delta\sum_j p_j \sum_x [(1-F(\bar{y}-x))\delta + F(\bar{y}-x)]\pi_j^1(x)}\right) \\ &\quad \cdot \left(\sum_j p_j \sum_x \left[\mathbb{E}[u(y)|x] + \delta(1-F(\bar{y}-x))\mathbb{E}[u(y)|\hat{x}_j]\right]\pi_j^1(x)\right).\end{aligned}$$

Taking limits as $y^* \to \infty$, and using L'Hôpital's rule in case $\delta \to 1$, we see that $(1-\delta)V^C(\sigma) \to \sum_j p_j \mathbb{E}[u(y)|\hat{x}_j]$. But we also have

$$\mu_T(n|p, y^*) \;=\; \frac{p_n \sum_x f(y^* - x)\pi_n^1(x)}{\sum_k p_k \sum_x f(y^* - x)\pi_k^1(x)} \;\leqslant\; \frac{1}{1 + \sum_{k<n} p_k \frac{f(y^* - x_j^*)}{p_n f(y^* - \bar{x})}} \;\to\; 1.$$

By the indifference condition (29), we then also have $V^C(\sigma) \to \mathbb{E}[u(y)|\hat{x}_n]$, a contradiction. We conclude that $x_n^* \to \infty$. By Proposition 6.15, it cannot be that the type n politicians place probability one on x_j^* as β becomes large (with δ possibly depending on β), and it follows that $x_{*,n} \to \hat{x}_n$, which proves (ii). Moreover, since policy choices are ordered by type, this implies that for all $j < n$, we have $x_j^* \to \hat{x}_j$. This proves (iv).

Finally, if there is a subsequence such that $\pi_n^1(x_n^*) = 0$ for arbitrarily large β (with δ possibly depending on β), then (32) again yields the implication $(1-\delta)V^C(\sigma) \to \sum_j p_j \mathbb{E}[u(y)|\hat{x}_j]$, and choosing any $\bar{x} > \hat{x}_n$, we obtain a contradiction as in the previous paragraph. We conclude that the type n politicians place positive probability on the greatest optimal policy for sufficiently large β (with δ possibly depending on β), and Proposition 6.15 implies that $\pi_n^1(x_n^*) \to 0$. This proves (iii) and establishes the result.

6.4 Symmetric learning

A class of models related to the one-sided learning setting of the previous subsections are the symmetric learning models, inspired by Holmstrom's (1999) model of career concerns. Here, a politician may be one of several valence types, but neither the politician nor the other citizens directly observe the politician's ability prior to the election; rather, voters and the politician receive public signals and update their beliefs about the politician's ability in the same way. Rather than being a preference parameter indexing cost of effort, the politicians' types are interpreted as an ability parameter, where outcome distributions for higher types dominate those for lower types. Political agency models using the informational assumption of symmetric learning encompass work of Persson and Tabellini (2000), Ashworth (2005), and Ashworth and Bueno de Mesquita (2008), discussed in Subsection 3.4. In addition, Martinez (2009) analyzes a three-period model in which effort is chosen in the first two periods before an election in period three, and he shows that in equilibrium, effort increases as the election is approached, and he discusses equilibrium dynamics for the finite-horizon model using numerical methods.

An advantage of the symmetric learning model over the pure moral hazard model is that it precludes some arbitrariness of the equilibrium selection, as the trivial "shirking equilibrium" will not generally persist: instead of shirking, an office holder will have an incentive to manipulate the updating of the voter's beliefs to increase her chances of re-election. A technical simplification over the one-sided learning model is that in a pure strategy equilibrium, politicians and voters update their beliefs the same way, precluding complications due to private information; because information is symmetric, all types of politicians face the same optimization problem and make the same policy choice along the equilibrium path of play. Of course, because the symmetric learning model assumes away private information about policy preferences, it may not be suitable for the analysis of elections as mechanisms for solving conflicts of interest between citizens and their elected delegates. Moreover, this class of models encounters the same issues with equilibrium existence as does the model with private information: as in Subsections 3.4 and 6.3, equilibria must solve a non-trivial fixed point problem, where the voters' cutoff rule determines an optimal effort choice for an office holder, and the effort choice of the politicians determines (via Bayes rule) a cutoff for the voters; and again, as in Figures 1 and 5, the optimization problem of an office holder suffers from potential non-convexities.

Modifying the formalism of the dynamic elections framework slightly, the utility of a politician is now $w(x) = v(x) - c(x)$ and is independent of type, and we let $\hat{x}$ denote the unique ideal policy of the politicians. Given policy choice x by a type j politician, the outcome y is realized from the density $f_j(y-x)$. To bring this closer to the framework of this paper, we fix parameters $\tilde{z}_1 < \tilde{z}_2 < \cdots < \tilde{z}_n$ for each politician type, and we simply assume that $f_j(y-x) = f(y-\tilde{z}_j-x)$, effectively incrementing the policy choices of higher types by larger amounts. Then under (C1) and (C2), higher outcomes are evidence that the politician is a higher type. We let

$\tilde{F}(\cdot|x)$ be the ex ante distribution of the policy outcome given policy choice x, so that $\tilde{F}(y|x) = \sum_j p_j F(y - \tilde{z}_j - x)$ is the probability of an outcome realization below y given policy choice x; and we let $\tilde{f}(y|x)$ be the associated ex ante density. We let $\mathbb{E}[u(y)|x, j]$ denote the voters' expected utility when a type j politician chooses x. Note that even if the density $f(\cdot)$ satisfies (C1) and (C2), it does not follow that the ex ante density inherits these properties, so to maintain desirable quasi-concavity properties of politician payoffs, we strengthen these conditions, without re-stating them here, to apply to the ex ante density $\tilde{f}(\cdot)$ as well.

Existence of equilibrium in the infinite-horizon model without term limits is an open question that is fraught with the same technical difficulties encountered in the analysis of the model of adverse selection and moral hazard without term limits. We therefore focus in this subsection on the model with a two-period term limit. Without going into formalities, we modify the concept of stationary electoral equilibrium so that policy choices are independent of the office holder's type (since politicians do not observe their own types), and we let politicians condition their choices on the term of office; of course, in equilibrium all politicians choose the ideal policy $\hat{x}$ in their second term, if re-elected. As always, the strategies of voters are summarized by a cutoff $\overline{y}$ such that a first-term incumbent is re-elected if and only if the realized outcome satisfies $y \geqslant \overline{y}$. Letting $V^C(\sigma)$ be the continuation value of a challenger, the voters' cutoff must satisfy the indifference condition

$$\sum_j \mu_T(j|p,\overline{y})\mathbb{E}[u(y)|\hat{x}, j] \quad = \quad (1-\delta)V^C(\sigma)$$

in equilibrium. In the first term of office, a politician chooses policy to solve

$$\max_{x\in X} w(x) + \delta\left[(1-\tilde{F}(\overline{y}|x))[w(\hat{x}) + \beta + \delta V^C(\sigma)] + \tilde{F}(\overline{y}|x)V^C(\sigma)\right],$$

and of course the voters' posterior beliefs $\mu_T(\cdot|p,y)$ are determined by Bayes rule.

We adapt condition (C3) in the obvious way, to account for symmetric learning, so that politicians prefer to be re-elected, and office holders will not choose policies below the ideal policy:

$$\text{(C3)} \qquad w(\hat{x}) + \beta \quad > \quad \mathbb{E}[u(y)|\hat{x}, n].$$

We re-phrase (C4) in terms of the ex ante density, so that given a cutoff $\overline{y}$ and a continuation value of a challenger $V \leqslant \frac{1}{1-\delta}\mathbb{E}[u(y)|\hat{x}, n]$, the politicians have at most two optimal policies, say $x^*(\overline{y}, V)$ and $x_*(\overline{y}, V)$. Specifically, we assume that for all finite $\overline{y}$ and all $x, \tilde{x}, z$ with $\hat{x} < x < \tilde{x} < z$, we have

$$\text{(C4)} \qquad \text{if } \frac{w''(x)}{w'(x)} \leqslant -\frac{\tilde{f}'(\overline{y}|x)}{\tilde{f}(\overline{y}|x)} \text{ and } \frac{w''(z)}{w'(z)} \leqslant -\frac{\tilde{f}'(\overline{y}|z)}{\tilde{f}(\overline{y}|z)}, \quad \text{then } \frac{w''(\tilde{x})}{w'(\tilde{x})} < -\frac{\tilde{f}'(\overline{y}|\tilde{x})}{\tilde{f}(\overline{y}|\tilde{x})}.$$

Now supermodularity of the objective function $U(x, 1 - \tilde{F}(\overline{y}|x); V)$ plays no role, as policy choices are symmetric with respect to type by assumption.

Since politicians always shirk in the second term of office, it is not necessary for a re-elected incumbent to condition her policy choices on the updated beliefs of the voter. The absence of such conditioning, which is required in the model with no term limit, significantly simplifies the equilibrium analysis, and existence of stationary electoral equilibrium follows from arguments similar to the proof of Proposition 6.14 in the model with one-sided learning.

Proposition 6.18 *In the infinite-horizon model of symmetric learning with two-period term limit, assume (C1)–(C4). Then there is a stationary electoral equilibrium, and in every stationary electoral equilibrium, politicians mix in the first term of office over at most two policies, say x^* and x_*, where $\hat{x} < x_* \leqslant x^*$, politicians choose $\hat{x}$ in the second term of office if re-elected, and voters re-elect an office holder after the first term if and only if $y \geqslant y^*$.*

As in the model of adverse selection and moral hazard with a two-period term limit and one-sided learning, we immediately obtain an upper bound on the voters' expected utility from policy choices of first-term office holders.

Proposition 6.19 *In the infinite-horizon model of symmetric learning with two-period term limit, assume (C1)–(C4). For all levels of office benefit $\beta \geqslant 0$ and all discount factors $\delta \in [0, 1)$, in every stationary electoral equilibrium σ, the expected utility to voters from policies chosen by first-term office holders is less than the discounted expected utility from the choice of the ideal policy by the type n politician, i.e.,*

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x, j]\pi^1(x) \quad < \quad \mathbb{E}[u(y)|\hat{x}, n].$$

The next proposition parallels the characterization in Corollary 6.1 by showing that for a given level of office benefit, the (normalized) continuation value of a challenger is bounded strictly below the voters' expected utility from the choice of the ideal point $\hat{x}$ by the highest type of politician.

Proposition 6.20 *In the infinite-horizon model of symmetric learning with two-period term limit, assume (C1)–(C5) hold, and fix the office benefit $\beta \geqslant 0$. Then there is a bound $\overline{u} < \mathbb{E}[u(y)|\hat{x}, n]$ such that for all discount factors $\delta \in [0, 1)$ and every stationary electoral equilibrium σ, the expected utility to voters from policies chosen by first-term office holders is below this bound, i.e.,*

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x, j]\pi^1(x) \quad \leqslant \quad \overline{u}.$$

Finally, paralleling Proposition 6.17, we note that politicians must place positive probability on arbitrarily high policies and policies close to the ideal, and again the probability of arbitrarily high policies must go to zero. Here, we simplify the statement of the result by fixing the discount factor.

Proposition 6.21 *In the infinite-horizon model of symmetric learning with two-period term limit, assume (C1)–(C4) hold, fix the discount factor* $\delta \in [0,1)$*, and let the office benefit* $\beta \geqslant 0$ *be arbitrarily large. Then for every selection of stationary electoral equilibria* σ*, the voters' cutoff diverges to infinity, and first-term office holders mix between policies that are close to the ideal policy and ones that are arbitrarily high, with small, positive probability on the latter, i.e.,*

(i) $y^* \to \infty$*,*

(ii) $x^* \to \infty$ *and* $x_* \to \hat{x}$*,*

(iii) $\pi^1(x^*) > 0$ *for large enough* β *and* $\pi^1(x^*) \to 0$*.*

7 Applied work

In this section, we touch on the applied literature related to the dynamic elections framework. Topics that have received attention include the possibility of political inefficiency, the effectiveness of elections in attaining accountability, and the effect of term limits, partisanship, and information on accountability. There is, in addition, a large applied literature on the existence of political cycles of different types. Besley and Case (2003) and Ashworth (2012) provide reviews of the accountability literature, and reviews of the political business cycle literature can be found in Persson and Tabellini (1990,2000), Alesina, Roubini, and Cohen (1997), and Drazen (2000).

7.1 Political inefficiency

The literature on political inefficiency has considered a number of different mechanisms that can lead to outcomes that are undesirable for voters. A recent literature has contemplated the possibility that electoral incentives may actually operate in a perverse manner, as the pursuit of re-election induces politicians to take actions that actually reduce the welfare of voters. In recent work on pandering, Acemoglu et al. (2013) consider a two-period model of elections with adverse selection and moral hazard. A representative voter has quadratic utility and ideal policy at zero, and there are two politician types: an honest type with quadratic utility and ideal point $\hat{z}_h = 0$ equal to the median, and a type that accepts bribes with quadratic utility and (effective) ideal point $\hat{z}_b > 0$. The authors assume that conditional on policy choice z, the policy outcome y is distributed normally with mean zero, and they assume the variance is sufficiently high to permit an analysis of pure strategy equilibria. The authors show that although the politicians are conservative, the honest politician type chooses a liberal policy $z < 0$, which is bad for both the politician and the voter. When office benefit is sufficiently high, in fact both politician types choose liberal policies, in order to signal that they are not extreme.

The above model differs from the rent-seeking environment of Subsection 3.4, but the technical differences between the models are inessential, permitting us to apply the analysis of the two-period model of adverse selection and moral hazard to the problem of populism. The key insight that permits this application is the fact that in the analysis of Propositions 3.4–3.8, the full force of monotonicity of voter utility is not needed; because only ideal policies are chosen in the second period, the properties of voter utility at other policies only affect payoffs in the first period, but these are sunk. The equilibrium cutoff and, therefore, policy choices in the first period only depend on voter preferences over ideal policies, and as long as voters strictly prefer higher types to hold office in the second period, i.e.,

$$\mathbb{E}[u(y)|\hat{x}_1] \ < \ \mathbb{E}[u(y)|\hat{x}_2] \ < \ \cdots \ < \ \mathbb{E}[u(y)|\hat{x}_n], \tag{35}$$

the voters' indifference condition remains the same, as do the politicians' incentives in the first period. Therefore, the analysis carries over unchanged as long as voter utility is monotonic over the range of politician ideal policies.

This allows us to map the model of Acemoglu et al. (2013) into the two-type case of the framework of Subsection 3.4. Specifically, we identify the type 1 politician with the dishonest politician and type 2 with the honest politician, and we identify a policy choice $x \geqslant 0$ as the increment subtracted from the ideal policy $\hat{z}_b$ of the dishonest type. That is, the policy choice x corresponds to $\hat{z}_b - x$ in the model of populism, so higher effort in our model corresponds to more liberal policies in their model. Translated to our model, the ideal policy choice of the type 1 politician is $\hat{x}_1 = 0$, and the ideal policy choice of the type 2 politician is $\hat{x}_2 = \hat{z}_b$. This yields a special case of the rent-seeking environment, with the modification that voter utility is quadratic with ideal policy outcome $\hat{y} = \hat{x}_2$, i.e.,

$$u(y) \quad = \quad -(y-\hat{y})^2.$$

The outcome density $f(\cdot)$ is normal with mean zero and variance σ, and mean-variance analysis gives us

$$\mathbb{E}[u(y)|x] \quad = \quad -(x-\hat{y})^2 - \sigma.$$

In particular, $\mathbb{E}[u(y)|\hat{x}_1] < \mathbb{E}[u(y)|x_2]$, so the voters' preferences are increasing in politician ideal policies. As explained above, this leaves the incentives of voters and politicians unchanged from the rent-seeking environment, and the equilibrium analysis of Subsection 3.4 applies.

Of course, the analysis extends to more general settings, as long as (35) holds. Proposition 3.7 extends the existence result of Acemoglu et al. (2013) to the model with an arbitrary number of conservative politician types and more general assumptions on utility and the outcome distribution, permitting a politician to mix over at most two policy choices in the first period to accommodate outcome distributions with low variance. All politicians exert positive effort in equilibrium, so if the ideal policy of the type n politician is equal to zero, then the type n politician chooses a liberal policy in the first period. Proposition 3.8 shows that even with multiple conservative types, high office motivation leads all above average politician types to choose policies to the left of the median voter,[26] and in fact, as office benefit becomes high, these policy choices become arbitrarily extreme. Thus, the responsive democracy result in the rent-seeking environment is a double-edged sword, for it suggests the potential for arbitrarily low expected payoffs for the voter in the first period of the single-peaked model of Acemoglu et al. (2013).

We can apply the same logic to the infinite-horizon model of adverse selection and moral hazard with one-sided learning and a two-period term limit. Once again, if all politician types are conservative, so that voter preferences are increasing on politician ideal policies, then Propositions 6.14–6.17 can be applied. In particular, a stationary electoral equilibrium exists, and all politician types exert positive effort in the first term of office; thus, if the type n politician has ideal policy equal to zero, then she chooses a liberal policy in her first term. In a recent paper, Kartik and Van Weelden (2016) examine an infinite-horizon model with a two-period term limit. They note that the signaling incentives of first-term politicians can have either good or bad reputation effects, as the informational content contained in first-term policy choices can be at odds with the desirability of those choices.

A related literature considers situations in which politicians are better informed than voters about desirable policies for voters, but have an incentive to pander to voters by knowingly choosing policies that are not in the voters' best interest in order to avoid the appearance of being a bad type. Canes-Wrone, Herron, and Shotts (2001) study a two-type model of pandering in which politicians differ in ability. Maskin and Tirole (2004) study pandering in a two-type model in which politicians differ in preferences; they use this model to study when decision making powers should be allocated to elected representatives versus non-accountable

[26]Duggan and Martinelli (2017) show that for a class of preferences generalizing quadratic utility, at least one below average type chooses high policy, so in the two-type model, the result implies that both politician types will engage in populism.

officials like judges or bureaucrats (or directly to voters themselves). Prat (2005), Fox and Shotts (2009), Fox and Van Weelden (2012), and Morelli and Van Weelden (2013) offer more recent papers in this vein. Pandering models illustrate the possibility of a conflict between responsive democracy and monotonic voting: it may be necessary to punish politicians for bad policy outcomes, even if they are payoff-irrelevant in the current election, to achieve responsiveness. Intriguingly, evidence from lab experiments conducted by Woon (2012) indicates that voters pursue retrospective strategies even in an environment, taken from Fox and Shotts (2009), in which prospective and retrospective voting cannot be reconciled in a perfect Bayesian equilibrium.

Besley and Coate (1998) analyze a different mechanism in the context of a two-period electoral model, and they show that coordination and commitment problems can lead to inefficient policy choices, as potential Pareto improving choices by an office holder in the first period may affect choices in the second period or have adverse electoral consequences. Some authors, e.g., Persson and Svensson (1989) and Alesina and Tabellini (1990), have focussed on inefficiencies arising from the incentive to "tie the hands" of the future party in power via the issuance of debt; Aghion and Bolton (1990) consider a related mechanism, in which the issuance of debt can decrease the probability that a liberal party wins the election. In the political cycles literature, which we touch on in the next subsection, Rogoff and Sibert (1988) assume a distorting seignorage tax that competent politicians use to signal their types, and Persson and Tabellini (1990) permit a competent office holder to create unexpected inflation, expanding the economy and signaling her type. Casamatta and De Paoli (2007) show that inefficiency can arise when an office holder uses an inefficient production technology in order to conceal the state of the world.

We have taken as given the nature of the rents extracted by politicians. The "Virginia school" of political economy (e.g., Tullock 1990) has stressed in the past the idea that electoral accountability may induce politicians to disguise their rent-seeking activities in inefficient ways. This point has been formalized by Coate and Morris (1995), who consider a special interest group and the possibility of a transfer from voters to the group. They show that in equilibrium, bad types of politician may confer benefits to the interest group using a risky public project, rather than direct transfers. Electoral incentives may also act in a perverse manner by inducing politicians to divert their effort and resources to campaigning. In Daley and Snowberg (2011), for example, politicians divert effort toward campaigning, which is unproductive but serves to influence voter beliefs. Ash, Morelli, and Van Weelden (2016) consider a multi-task model in which an office holder allocates efforts across two dimensions and may focus on the issue that generates lower utility but greater electoral impact.

7.2 Political cycles

The political cycles literature is extensive and spans work from early models of Nordhaus (1975), Lindbeck (1976), and Hibbs (1977), who assumed myopic voters, to later models of Alesina (1987,1988a), Rogoff and Sibert (1988), Persson and Tabellini (1990), Rogoff (1990), and others, in which voters rationally anticipate the unobserved actions of politicians (or parties) when elected to office. This literature considers the possibility of several kinds of cycles arising from electoral incentives: political business cycles, in which electoral incentives influence real economic variables prior to elections; political budget cycles, in which real economic variables are affected by fiscal decisions, which vary with the party in power, rather than by monetary policy; and partisan cycles, in which economic variables reflect the partisan affiliation of the politician who holds office. Reviews can be found in Persson and Tabellini (2000), Alesina, Roubini, and Cohen (1997), and Drazen (2001).

We discuss the interpretation of the two-period model of adverse selection and moral hazard in the context of political business cycles, specifically, the model of Persson and Tabellini (1990). They consider a two-period model in which the first-period incumbent chooses a level x_1 of inflation and a level z_1 of employment,[27] voters form an expected level x_1^e of inflation in the first period, and their utility is $u(x_1,y_1) = -x_1^2/2 + z_1$; and in the second period, the winner of the election chooses x_2 and z_2, with similar payoffs. One difference between this model and the electoral framework of Subsection 3.4 is that politician types θ_1 and θ_2 are interpreted as competence, i.e., as an exogenous increment to the level of employment generated by any given choice of inflation. Another difference is that politicians have two choice variables, but in fact these are related by the constraint $z = x - x^e + \theta_j$, and thus the reduced form of voters' utility is

$$u(x,\theta_j) \quad = \quad -x^2/2 + x - x^e + \theta_j,$$

$j = 1,2$. Note that this is a quadratic form with ideal policy $\hat{x} = 1$. Furthermore, while the expected inflation term x_t^e in period $t = 1,2$ is important for the welfare analysis, it does not affect equilibrium: because it is pinned down by the politicians' equilibrium strategies, and because it enters voter utility linearly, it does not affect the voters' re-election decision. Politician payoffs are the same, except for the addition of an office benefit term when the politician holds office. When a type θ_j politician holds office in the second period, she simply chooses her ideal policy, which is $\hat{x} = 1$ with the functional form above, and the voters' utility is simply

[27] Note that Persson and Tabellini (1990) use x to denote employment; we adapt notation to suit the present discussion.

$\theta_j - \frac{1}{2}$. The key feature of this formulation is that voters prefer the higher type of politician to hold office in the second period, and this generates an incentive for the first-period incumbent to signal that she is the high type, thereby improving her re-election prospects.

A third, more substantial difference between the models is that in the model of Persson and Tabellini (1990), the first-period incumbent's controls inflation completely, so that her policy decision determines employment precisely without any intervening noise. Employment does depend on the politician's type, so voters cannot perfectly infer the inflation choice x_1. This makes it possible for the low politician type to perfectly mimic the high politician type, raising the possibility of pooling or perfect separation in equilibrium, a phenomenon that cannot arise in the model with adverse selection and moral hazard: there, for example, it is always possible that a high policy choice will lead to a low policy outcome, after which voters update negatively.

We re-interpret the two-period model of adverse selection and moral hazard model of Subsection 3.4 as follows. Let $x_t \geqslant 0$ be the choice of a monetary instrument that influences inflation, and assume that the choice x_t is not observed by voters, but it determines an economic variable (analogous to inflation) $y_t = x_t + \varepsilon_t$ that is observed, where ε_t is a normally distributed noise term with mean zero and variance σ. That is, monetary instruments influence inflation, but they do not pin down inflation with uncertainty. Let y_t^e be the level of y_t expected by voters in period $t = 1, 2$, and assume that employment is given by $z_t = y_t - y_t^e$. Voter utility in period t is quadratic in y and linear in z, plus the constant $\sigma/2$, i.e.,

$$u(y) \quad = \quad -y^2/2 + y - y^e + \sigma/2.$$

Using the fact that ε has mean zero, we have

$$\mathbb{E}[u(y)|x] \quad = \quad -x^2/2 + x - y^e,$$

mirroring the structure of Persson and Tabellini (1990). The utility for the type j politician is

$$w_j(x) \quad = \quad -x^2/2 + x - \frac{1}{\theta_j} c(x),$$

where $c(\cdot)$ is a cost function satisfying the assumptions of the rent-seeking model. This means that politicians incur an extra cost in generating inflation, perhaps due to personal financial interests or commitments to interested parties. Politician ideal points are ranked $\hat{x}_1 < \hat{x}_2 < \cdots < \hat{x}_n$, and voter preferences are increasing in this range, and as discussed in relation to pandering in the previous subsection, the equilibrium analysis of Subsection 3.4 can then be applied.

We conclude that electoral equilibria exist, that each politician type mixes over at most two policy choices, and that all policy choices are strictly positive: incumbents have an incentive to inflate the economy to influence voter perceptions and increase their probability of re-election. An implication is that expected inflation is positive in equilibrium, i.e., $y^e > 0$. It follows that the voters' expected payoff in the first period is

$$\sum_j p_j \sum_x \mathbb{E}[u(y)|x]\pi_j(x) \quad = \quad -\frac{1}{2}\sum_j p_j \sum_x \mathbb{E}[y^2|x]\pi_j(x) \; < \; 0.$$

Moreover, as politicians become highly office motivated, expected inflation becomes arbitrarily high, and voter welfare decreases without bound, illustrating a form of political inefficiency arising from manipulation of the economy prior to an election.

We also mention the work of Ales, Maziero, and Yared (2014), which considers jointly determined economic and political cycles and is close to the accountability setting. They assume an economy in which the politician in office has private information about the tightness of the government budget and rents. In particular, the budget constraint is subject to i.i.d. shocks every period. In the best equilibrium for voters, the politician in office is replaced with positive probability after a sequence of bad shocks to the government budget leading to high observable taxes. The probability of reelection is history dependent, despite the fact that the challenger and the politician in office are identical. Off the equilibrium path, the behavior of the challenger depends on history before she wins the election, for otherwise voters could be tempted to replace the politician in office. In their setting, in constrast, our notion of simple equilibrium would prescribe a stationary re-election rule. Azzimonti (2014) studies the dynamics of public investment in an infinite-horizon model paying attention to asymmetries in re-election probabilities. The model of politics in this line of work differs considerably from the accountability approach; the authors typically assume the existence of two political parties representing well-defined groups that alternate in power; preferences and actions are observed, and there is no opportunity to build a reputation.

Empirically, Alesina et al. (1997) find support for "rational partisan" cycles in the US but not for models of myopic voting or models in which politicians engage in "opportunistic," pre-election manipulation of the economy. Data for OECD countries are roughly consistent with the US, with support for the rational partisan model. Shi and Svensson (2006) report evidence for political budget cycles in a large cross-country data set, with stronger effects in developing countries. Work by Brender and Drazen (2005) suggests that policy manipulation may be more prevalent in new and weaker democracies, where voters are inexperienced or lack information to evaluate policies.

7.3 Accountability

One of the fundamental predictions of the electoral accountability approach is that some, if not all, politicians will compromise their policy choices in order to improve their prospects for being re-elected; or in the rent-seeking contest, some politicians will exert greater effort to improve their chances of re-election. In sum, politicians respond to electoral incentives. As a corollary, politicians will tend to compromise less (or "shirk" more) when electoral constraints are relaxed. Early evidence supporting this comparative static is reported by Kalt and Zupan (1984,1990) and Zupan (1990), the latter paper comparing politician behavior before and after the decision to retire.

Another potentially useful test is to compare the choices of term-limited politicians to their choices earlier in their tenure of office. Using data from 1950–1986, Besley and Case (1995) find that there is a difference between the first and the second term in office for US state governors who are incumbents and face term limits: state taxes and spending are higher in the second term when term limits bind. That is, state governors behave differently when not subject to re-election incentives. Besley and Case (2003) update these results using data from 1950–1997. They still find an effect on state spending; intriguingly, however, the earlier effect on taxes is reversed, so that lame duck governors instead generate lower state taxes. Alt, Bueno de Mesquita, and Rose (2011) account for the fact that term limits vary from one to two terms across states, and they recover the earlier positive effect on state taxes.

Other work examines the connection between responsiveness and electoral incentives (or lack thereof) from term limits in different ways. Ferraz and Finan (2011) use data from a program of anticorruption audits carried out by the Brazilian central government to study term limits and incumbent behavior at the municipality level. They find that, consistent with decreased incentives for re-election, lame duck mayors are more corrupt than other mayors. They also show that the effects of re-election incentives are more pronounced in municipalities in which there are no local public prosecutors, implying weaker judicial constraints on behavior, and in municipalities where the elections are competitive.

At a more fundamental level, the effects of term limits may be limited by the extent to which politicians facing them may care for their successor belonging to the same political party, or the same faction of the party. Politicians facing term limits, for instance, may be motivated by the expectation of continuing a political career in a different position with the support of the party, or may partake in the office rents if the successor belongs to the same faction. An interesting empirical question, to which we pay some attention below, is whether accountability works by punishing the political party in countries like Mexico with tight no-reelection rules.

Frey (2016) analyzes the effect of income on electoral accountability. He proposes a dynamic model of elections in which a component of voter preferences is unobserved by politicians and an office-motivated incumbent chooses public good levels for rich and poor voters, along with a monetary payment to a subset of poor voters. A poor voter who receives the payment from the incumbent may not receive it from the challenger, so payments are a form of vote buying. He deduces that if the income of poor voters increases, then, due to income effects, those voters are less concerned with the risk that a challenger will fail to provide those payments. Frey examines data from the Bolsa Família, a large scale cash transfer program in Brazil that offers poor families an increment to their income, and his results confirm that an increase in the income of poor families leads to a decrease in clientelism and an increase in the level of public good consumed by the poor.

Two recent papers have applied methods of structural estimation to examine empirical support for infinite-horizon models of electoral accountability. Sieg and Yoon (2017) provide a structural analysis of the infinite-horizon model of pure adverse selection using data from US gubernatorial elections from 1950–2012. They replicate and extend the empirical results of Besley and Case (1995), and they estimate the distribution of candidate ideologies for each party, the distribution of voter preferences, and the office benefit of politicians from each party. The authors find that candidates from the two parties are drawn from distinct distributions with non-overlapping support, and that the distribution of voter ideal policies is similar to, and somewhat more polarized than, the distribution of potential challengers. For the estimated parameter values, the authors find that election standards are tighter, i.e., the win set is smaller, in the presence of a two-period term limit, providing support for term limits.

Aruoba et al. (2015) estimate a structural model using voter ratings of job performance from US gubernatorial elections for 1982–2012. They find that politicians respond to electoral incentives by exerting positive effort, and they find that elections have a positive selection effect on the quality of re-elected politicians. Like Sieg and Yoon (2017), the authors examine an infinite-horizon model, but in contrast to the latter paper, they consider the model of adverse selection and moral hazard with a two-period term limit, two type, and two effort levels. They find that a substantial proportion of "bad" governors are disciplined to exert high effort in their first terms, and that most governors who are re-elected are the "good" type.

The electoral accountability approach also emphasizes the importance of the availability of information to judge performance and the existence of political alternatives to allow voters to punish or reward politicians in office. A growing literature uses field and natural experiments to study the effects of information and media coverage on voting behavior. Ferraz and Finan (2008) make use of the data on Brazilian municipalities to examine the effect of the disclosure of information about corruption on electoral outcomes. They find that the likelihood of re-election was substantially reduced for mayors who were found to have committed several violations associated to corruption, and that the effects were magnified in municipalities where local radio was present, presumably enabling the dissemination of information.

In a similar spirit, Chong et al. (2015) conduct a randomized control experiment and randomly assigned voting precincts to a campaign that disseminated information about the performance of mayors just before municipal elections in Mexico. They observe that information on corruption decreases electoral support both for the incumbent party and for the challenger party, so that it mostly reduces turnout. The interpretation is that in an environment such that voters perceive they have little opportunity to alter the electoral result (because of high partisanship, barriers to entry, electoral fraud, etc.), more damaging information may simply reduce overall trust in politicians. Thus, better information may be necessary, but it is not sufficient to improve the performance of representative democracy in settings with scarce actual political competition.

There is also a growing literature on the influence of mass media on accountability via voters' information. Besley and Burgess (2002) use panel data from India to illustrate the impact of mass media circulation on government responsiveness. In the US, Snyder and Stromberg (2010) use the mismatch between media markets and congressional districts to estimate the effect of media coverage on the incumbent on policy choices. They find that a better match between media markets and congressional districts improves the coverage of incumbent politicians, which in turn leads to policy choices that are more congruent with the citizens' preferences. Prat and Stromberg (2013) survey the recent work.

8 Modeling challenges

We conclude with a discussion of modeling challenges that are not addressed in this survey and which we view as important steps in the development and applicability of the electoral accountability approach.

First, with respect to the elections themselves, the framework we have presented assumes that voter preferences are known to politicians. A more realistic framework, which we believe would preserve many of the results covered above (in spirit, if not literally), would assume "probabilistic voting," as in Lindbeck and Weibull (1993); see also Alesina (1988b) for a model of repeated elections with probabilistic voting. The framework also abstracts away from the role of money, through either special interest lobbying (e.g., Snyder and Ting (2008)) or campaign finance; see also Dixit, Grossman, and Helpman (1997), Bergemann and Valimaki (2003) for common agency models of interest group lobbying. Importantly, the framework also abstracts from the role of media, particularly through information about the challenger's intended policies (e.g., Duggan and Martinelli (2011)), and from the role of electoral platforms as conveyors of information from candidates to voters (e.g., Martinelli (2001)). The incorporation of these realistic features of politics would permit the analysis of a number of interesting issues and could inform the current debate about the desirability of limits on campaign contributions or of media regulation.

Second, with respect to politicians, the framework should be extended to incorporate a meaningful model of political careers, including endogenous challenger selection and the possibility that a former office holder re-enters the political scene (rather than the current standard of a random draw without replacement). Such an extension may incorporate aspects of Mattozzi and Merlo (2008), in which the career decision to enter politics is endogenized in an overlapping generations setting. This is an important issue to the extent that term limits may not be so detrimental to accountability when politicians are motivated by a career beyond their current position.

Third, with respect to policy making, the accountability framework typically considers the policy choice problem of a single office holder in isolation, but the paradigm must be extended to capture interaction among multiple political office holders, as in Alesina and Rosenthal (1996); more recently, Cho (2009) analyzes a model of political representation in a single-member district system, and Fox and Van Weelden (2010) and Fox and Stephenson (2011) consider the effect of a veto player in the electoral accountability framework. This is essential to better understand the effects of division of powers on long run policy outcomes, for the comparison of different political systems, and the study of constitutional design issues introduced in formal modeling by Persson, Roland, and Tabellini (1997) and Laffont (2000).

Fourth, and related, the framework largely abstracts away from the role of political parties. One direction of investigation is to expand the analysis to cover party primaries, as in Serra (2011). The party system considered does not need to be restricted to the two-party, majoritarian system reflecting politics in the US, and applicability of the electoral accountability framework would be significantly increased by incorporating structure of multi-party, PR systems. Austen-Smith and Banks (1988) and Baron and Diermeier (2001) consider two-period models of PR systems, while Cho (2014) considers an infinite-horizon model with an endogenous status quo. These models assume complete information, so issues of adverse selection and moral hazard, which are prevalent in the electoral accountability literature, do not arise.

Fifth, and of most importance for applications, the framework must be extended to accommodate a state variable that evolves over time. This is a necessary antecedent, for example, to the detailed study of the political determinants of growth, inequality, and redistribution, continuing the work of Bertola (1993), Perotti (1993), Alesina and Rodrik (1994), Persson and Tabellini (1994), Krusell, Quadrini, and Rios-Rull (1997), Krusell and Rios-Rull (1999), and Benabou (2000), some of whom assume a form of the median voter theorem. More generally, the detailed modeling of politics is relevant for the integration models of elections into dynamic macroeconomic models; see Battaglini and Coate (2007, 2008), Acemoglu, Golosov, and Tsyvinski (2008), and Yared (2010) for recent contributions in this spirit. Camara (2012) includes an extension to growth economies that preserve the stationary structure of his equilibrium. Duggan (2012) contains a general existence result for the complete information model with a general state variable and idiosyncratic preference shocks; moreover, the result allows for a political game played by multiple politicians each period. Battaglini (2014) analyzes a dynamic model of elections in which two parties simultaneously announce fiscal policy platforms each period, voter preferences are subject to idiosyncratic shocks, the variance of the shocks varies stochastically over time, and each party myopically maximizes the number of its elected representatives. He establishes existence and characterization results for Markov perfect equilibria in which players condition on the level of borrowing in the previous period (in addition to real variables), and he gives necessary and sufficient conditions under which political equilibria are efficient. Duggan and Forand (2014) allow for a countable state space and do not require preference shocks. They assume complete information, establish existence of equilibrium, and assuming a small amount of commitment (in place of private information), provide a number of responsiveness results.

Finally, a theoretical issue, but one that poses an obstacle to potentially interesting applications of the accountability framework, is the question of equilibrium existence in the infinite-horizon model of adverse selection and moral hazard with no term limit. We have provided some characterization results in Subsection 6.2, but these assume only two types and are restricted to reputation-monotonic equilibria, and until the existence problem is solved, these results may be vacuous.

References

Acemoglu, D., Egorov, G., Sonin, K., 2013. A political theory of populism. Quarterly Journal of Economics 128, 771–805.

Acemoglu, D., Golosov, M., Tsyvinski, A., 2008. Political economy of mechanisms. Econometrica 76, 619–641.

Acemoglu, D., Johnson, S., Robinson, J., 2005. Institutions as a fundamental cause of long-run growth. Handbook of Economic Growth 1A, 385–472.

Aghion, P., Bolton, P., 1990. Government domestic debt and the risk of default: A political-economic model of the strategic role of debt, in: Dornbusch, R., Draghi, M. (Eds.), Public Debt Management: Theory and History. Cambridge: Cambridge University Press, pp. 315–344.

Ales, L., Maziero, P., Yared, P., 2014. A theory of political and economic cycles. Journal of Economic Theory 153, 224–251.

Alesina, A., 1987. Macroeconomic policy in a two-party system as a repeated game. Quarterly Journal of Economics 102, 651–678.

Alesina, A., 1988a. Credibility and policy convergence in a two-party system with rational voters. American Economic Review 78, 796–806.

Alesina, A., 1988b. Macroeconomics and politics. NBER Macroeconomics Annual 3, 13–52.

Alesina, A., Rodrik, D., 1994. Distributive politics and economic growth. Quarterly Journal of Economics 109, 465–490.

Alesina, A., Rosenthal, H., 1996. A theory of divided government. Econometrica 64, 1311–1341.

Alesina, A., Roubini, N., Cohen, G., 1997. Political Cycles and the Macroeconomy. Berkeley: University of California Press.

Alesina, A., Tabellini, G., 1990. A positive theory of budget deficits and government debt. Review of Economic Studies 57, 403–414.

Alt, J., Bueno de Mesquita, E., Rose, S., 2011. Disentangling accountability and competence in elections: Evidence from us term limits. Journal of Politics 73, 171–186.

Aragones, E., Palfrey, T., Postlewaite, A., 2007. Political reputations and campaign promises. Journal of the European Economic Association 5, 846–884.

Aruoba, B., Drazen, A., Vlaicu, R., 2015. Structural estimation of political accountability. Unpublished paper .

Ash, E., Morelli, M., Van Weelden, R., 2016. Elections and divisiveness: Theory and evidence. Journal of Politics forthcoming.

Ashworth, S., 2005. Reputational dynamics and political careers. Journal of Law, Economics, and Organization 21, 441–466.

Ashworth, S., 2012. Electoral accountability: Recent theoretical and empirical work. Annual Review of Political Science 15, 183–201.

Ashworth, S., Bueno de Mesquita, E., 2006. Delivering the goods: Legislative particularism in different electoral and institutional settings. Journal of Politics 68, 169–179.

Ashworth, S., Bueno de Mesquita, E., 2008. Electoral selection, strategic challenger entry, and the incumbency advantage. Journal of Politics 70, 1006–1025.

Ashworth, S., Bueno de Mesquita, E., 2014. Is voter competence good for voters? Information, rationality, and democratic performance. American Political Science Review forthcoming.

Austen-Smith, D., Banks, J., 1988. Elections, coalitions, and legislative outcomes. American Political Science Review 82, 405–422.

Austen-Smith, D., Banks, J., 1989. Electoral accountability and incumbency, in: Peter Ordeshook (Ed.), Models of Strategic Choice in Politics. University of Michigan Press, pp. 121–150.

Banks, J., Duggan, J., 2000. A bargaining model of collective choice. American Political Science Review 94, 73–88.

Banks, J., Duggan, J., 2008. A dynamic model of democratic elections in multidimensional policy spaces. Quarterly Journal of Political Science 3, 269–299.

Banks, J., Sundaram, R., 1993. Adverse selection and moral hazard in a repeated election model, in: Barnett, W., Hinich, M., Schofield, N. (Eds.), Political Economy: Institutions, Information, Competition, and Representation. Cambridge University Press, pp. 295–311.

Banks, J., Sundaram, R., 1998. Optimal retention in agency problems. Journal of Economic Theory 82, 293–323.

Barganza, J.C., 2000. Two roles for elections: Disciplining the incumbent and selecting a competent candidate. Public Choice 105, 165–193.

Baron, D., Diermeier, D., 2001. Elections, governments, and parliaments in proportional representation systems. Quarterly Journal of Economics 116, 933–967.

Baron, D., Ferejohn, J., 1989. Bargaining in legislatures. American Political Science Review 83, 1181–1206.

Barro, R., 1973. The control of politicians: An economic model. Public Choice 14, 19–42.

Bassetto, M., Benhabib, J., 2006. Redistribution, taxes, and the median voter. Review of Economic Dynamics 9, 211–223.

Battaglini, M., 2014. A dynamic theory of electoral competition. Theoretical Economics 9, 515–554.

Battaglini, M., Coate, S., 2007. Inefficiency in legislative policy-making: A dynamic analysis. American Economic Review 97, 118–149.

Battaglini, M., Coate, S., 2008. A dynamic theory of public spending, taxation, and debt. American Economic Review 98, 201–236.

Benabou, R., 2000. Unequal societies: Income distribution and the social contract. American Economic Review 90, 96–129.

Bergemann, D., Valimaki, J., 2003. Dynamic common agency. Journal of Economic Theory 111, 23–48.

Bernhardt, D., Camara, O., Squintani, F., 2011. Competence and ideology. Review of Economic Studies 78, 487–522.

Bernhardt, D., Campuzano, L., Squintani, F., Camara, O., 2009. On the benefits of party competition. Games and Economic Behavior 66, 685–707.

Bernhardt, D., Dubey, S., Hughson, E., 2004. Term limits and pork barrel politics. Journal of Public Economics 88, 2383–2422.

Bertola, G., 1993. Factor shares and savings in endogenous growth. American Economic Review 83, 1184–1210.

Besley, T., 2006. Principled Agents? The Political Economy of Good Government. Oxford University Press.

Besley, T., Burgess, R., 2002. The political economy of government responsiveness: Theory and evidence from India. Quarterly Journal of Economics 17, 1415–1451.

Besley, T., Coate, S., 1997. An economic model of representative democracy. Quarterly Journal of Economics 112, 85–114.

Besley, T., Coate, S., 1998. Sources of inefficiency in a representative democracy: A dynamic analysis. American Economic Review 88, 139–156.

Border, K., 1985. Fixed Point Theorems with Applications to Economics and Game Theory. Cambridge University Press.

Brender, A., Drazen, A., 2005. Political budget cycles in new versus established democracies. Journal of Monetary Economics 52, 1271–1295.

Calvert, R., 1985. Robustness of the multi-dimensional voting model: Candidate motivations, uncertainty and convergence. American Journal of Political Science 29, 69–95.

Camara, O., 2012. Economic policies of heterogeneous politicians. Unpublished paper .

Canes-Wrone, B., Herron, M., Shotts, K., 2001. Leadership and pandering: A theory of executive policymaking. American Journal of Political Science 45, 532–550.

Casamatta, G., Paoli, C.D., 2007. Inefficient public provision in a repeated elections model. Journal of Public Economic Theory 9, 1103–1126.

Cho, S.j., 2009. Retrospective voting and political representation. American Journal of Political Science 53, 276–291.

Cho, S.j., 2014. Voting equilibria under proportional representation. American Political Science Review 108, 281–296.

Chong, A., De La O, A., Karlan, D., Wantchenkon, L., 2015. Does corruption information inspire the fight or quash the hope? A field experiment in Mexico on voter turnout, choice, and party identification. Journal of Politics 77, 55–71.

Coughlin, P., Nitzan, S., 1981. Electoral outcomes with probabilistic voting and Nash social welfare maxima. Journal of Public Economics 15, 113–121.

Dixit, A., Grossman, G., Helpman, E., 1997. Common agency and coordination: General theory and application to government policy making. Journal of Political Economy 105, 752–769.

Downs, A., 1957. An Economic Theory of Democracy. New York: Harper and Row.

Drazen, A., 2001. The political business cycle after 25 years, in: Bernanke, B., Rogoff, K. (Eds.), NBER Macroeconomics Annual 2000, Volume 15. MIT Press, pp. 75–138.

Duggan, J., 2000. Repeated elections with asymmetric information. Economics and Politics 12, 109–136.

Duggan, J., 2012. Noisy stochastic games. Econometrica 80, 2017–2045.

Duggan, J., 2014a. A folk theorem for repeated elections with adverse selection. Political Science Research and Methods 2, 213–242.

Duggan, J., 2014b. Majority voting over lotteries: Conditions for existence of a decisive voter. Economics Bulletin 34, 263–270.

Duggan, J., 2016. Term limits and bounds on policy responsiveness in dynamic elections. Journal of Economic Theory forthcoming.

Duggan, J., Fey, M., 2006. Repeated downsian electoral competition. International Journal of Game Theory 35, 39–69.

Duggan, J., Forand, J.G., 2014. Markovian elections. Unpublished paper .

Duggan, J., Martinelli, C., 2011. A spatial theory of media slant and voter choice. Review of Economic Studies 78, 640–666.

Duggan, J., Martinelli, C., 2017. Electoral accountability and responsive democracy. Unpublished paper .

Fearon, J., 1999. Electoral accountability and the control of politicians: Selecting good types versus sanctioning poor performance, in: Przeworski, A., Stokes, S., Manin, B. (Eds.), Democracy, Accountability, and Representation. Cambridge University Press, pp. 55–97.

Ferejohn, J., 1986. Incumbent performance and electoral control. Public Choice 50, 5–25.

Ferraz, C., Finan, F., 2008. Exposing corrupt politicians: the effect of Brazil's publicly released audits on electoral outcomes. Quarterly Journal of Economics 123, 703–745.

Ferraz, C., Finan, F., 2011. Electoral accountability and corruption: Evidence from the audits of local governments. American Economic Review 101, 1274–1311.

Forand, J.G., 2014. Two-party competition with persistent policies. Journal of Economic Theory 152, 64–91.

Fox, J., Stephenson, M., 2011. Judicial review as a response to political posturing. American Political Science Review 105, 397–414.

Fox, J., Weelden, R.V., 2010. Partisanship and the effectiveness of oversight. Journal of Public Economics 94, 684–697.

Fox, J., Weelden, R.V., 2012. Costly transparency. Journal of Public Economics 96, 142–150.

Frey, A., 2016. Cash transfers, clientelism and political enfranchisement: Evidence from brazil. Unpublished paper .

Herrera, H., Levine, D.K., Martinelli, C., 2008. Policy platforms, campaign spending and voter participation. Journal of Public Economics 92, 501–513.

Herrera, H., Martinelli, C., 2006. Group formation and voter participation. Theoretical Economics 1, 461–487.

Hibbs, D., 1977. Political parties and macroeconomic policy. American Political Science Review 71, 1467–1487.

Hinich, M., 1977. Equilibrium in spatial voting: The median voter result is an artifact. Journal of Economic Theory 16, 208–219.

Holmstrom, B., 1999. Managerial incentive problems: A dynamic perspective. Review of Economic Studies 66, 169–182.

Hotelling, H., 1929. Stability in competition. Economic Journal 39, 41–57.

Kalandrakis, T., 2009. A reputational theory of two-party competition. Quarterly Journal of Political Science 4, 343–378.

Kalt, J., Zupan, M., 1984. Capture and ideology in the economic theory of politics. American Economic Review 74, 279–300.

Kalt, J., Zupan, M., 1990. The apparent ideological behavior of legislators: Testing for principal-agent slack in political institutions. Journal of Law, Economics, and Organization 33, 103–131.

Kang, I., 2005. Strategic challengers and the incumbency advantage. Unpublished paper .

Kartik, N., Van Weelden, R., 2016. Reputation, term limit, and incumbency (dis)advantage. Unpublished paper .

Kramer, G., 1977. A dynamical model of political equilibrium. Journal of Economic Theory 15, 310–334.

Krusell, P., Quadrini, V., Rios-Rull, J.V., 1997. Politico-economic equilibrium and economic growth. Journal of Economic Dynamics and Control 21, 243–272.

Krusell, P., Rios-Rull, J.V., 1999. On the size of the US government: Political economy in the neoclassical growth model. American Economic Review 89, 1156–1181.

Laffont, J.J., 2000. Incentives and Political Economy. Clarendon Lectures in Economics, Oxford University Press.

Lindbeck, A., 1976. Stabilization policy in open economies with endogenous politicians. American Economic Review Papers and Proceedings 66, 1–19.

Lindbeck, A., Weibull, J., 1993. A model of political equilibrium in a representative democracy. Journal of Public Economics 51, 195–209.

Martinelli, C., 2001. Elections with privately informed parties and voters. Public Choice 108, 147–167.

Martinez, L., 2009. A theory of political cycles. Journal of Economic Theory 144, 1166–1186.

Maskin, E., Tirole, J., 2004. The politician and the judge: Accountability in government. American Economic Review 94, 1034–1054.

Mattozzi, A., Merlo, A., 2008. Political careers or career politicians? Journal of Public Economics 92, 597–608.

Meirowitz, A., 2007. Probabilistic voting and accountability in elections with uncertain policy constraints. Journal of Public Economic Theory 9, 41–68.

Meltzer, A., Richard, S., 1981. A rational theory of the size of government. Journal of Political Economy 89, 914–927.

Morelli, M., Van Weelden, R., 2013. Ideology and information in policymaking. Journal of Theoretical Politics 25, 412–436.

Nordhaus, W., 1975. The political business cycle. Review of Economic Studies 42, 169–190.

Osborne, M., Slivinski, A., 1996. A model of competition with citizen candidates. Quarterly Journal of Economics 111, 65–96.

Perotti, R., 1993. Political equilibrium, income distribution, and growth. Review of Economic Studies 60, 755–776.

Persson, T., Roland, G., Tabellini, G., 1997. Separation of powers and political accountability. Quarterly Journal of Economics 112, 1163–1202.

Persson, T., Svensson, L., 1989. Why a stubborn conservative would run a deficit: Policy with time-inconsistent preferences. Quarterly Journal of Economics 104, 325–345.

Persson, T., Tabellini, G., 1994. Is inequality harmful for growth? American Economic Review 84, 600–621.

Persson, T., Tabellini, G., 2000. Political Economics: Explaining Economic Policy. MIT Press.

Prat, A., 2005. The wrong kind of transparency. American Economic Review 95, 862–877.

Prat, A., Stromberg, D., 2013. The political economy of mass media, in: Acemoglu, D., Arellano, M., Dekel, E. (Eds.), Advances in Economics and Econometrics, Tenth World Congress, Volume 2, Applied Economics. Cambridge University Press, pp. 135–187.

Reed, R., 1994. A retrospective voting model with heterogeneous politicians. Economics and Politics 6, 39–58.

Reny, P., 2011. The existence of monotonic pure strategy equilibria in bayesian games. Econometrica 79, 499–553.

Roberts, K., 1977. Voting over income tax schedules. Journal of Public Economics 8, 329–340.

Roemer, J., 1994. A theory of policy differentiation in single issue electoral politics. Social Choice and Welfare 11, 355–380.

Roemer, J., 1997. Political-economic equilibrium when parties represent constituents: The unidimensional case. Social Choice and Welfare 14, 479–502.

Rogoff, K., 1990. Equilibrium political budget cycles. American Economic Review 80, 21–36.

Rogoff, K., Sibert, A., 1988. Elections and macroeconomic policy cycles. Review of Economic Studies 55, 1–16.

Romer, T., 1975. Individual welfare, majority voting, and the properties of a linear income tax. Journal of Public Economics 4, 163–185.

Schwabe, R., 2011. Reputation and accountability in repeated elections. Unpublished paper .

Serra, G., 2011. Why primaries? The party's tradeoff between policy and valence. Journal of Theoretical Politics 23.

Shi, M., Svensson, J., 2006. Political budget cycles: Do they differ across countries and why? Journal of Public Economics 90, 1367–1389.

Sieg, H., Yoon, C., 2017. Estimating a dynamic game of gubernatorial elections to evaluate the impact of term limits. American Economic Review forthcoming.

Snyder, J., Ting, M., 2008. Interest groups and the electoral control of politicians. Journal of Public Economics 92, 482–500.

Tullock, G., 1990. The costs of special privilege, in: Alt, J., Shepsle, K. (Eds.), Perspectives on Positive Political Economy. Cambridge University Press, pp. 195–211.

Van Weelden, R., 2013. Candidates, credibility, and re-election incentives. Review of Economic Studies 80, 1622–1651.

Van Weelden, R., 2015. The welfare implications of electoral polarization. Social Choice and Welfare 45, 653–686.

Wittman, D., 1977. Candidates with policy preferences: A dynamic model. Journal of Economic Theory 14, 180–189.

Wittman, D., 1983. Candidate motivations: A synthesis of alternatives. American Political Science Review 77, 142–157.

Woon, J., 2012. Democratic accountability and retrospective voting: A laboratory experiment. American Journal of Political Science 56, 913–930.

Yared, P., 2010. Politicians, taxes and debt. Review of Economic Studies 77, 806–840.

Zupan, M., 1990. The last period problem in politics: Do congressional representatives not subject to a reelection constraint alter their voting behavior? Public Choice 65, 167–180.