The Political Economy of Dynamic Elections: Accountability, Commitment, and Responsiveness

(Forthcoming at Journal of Economic Literature)

John Duggan* César Martinelli†

May 23, 2017

Abstract

We survey the literature on dynamic elections in the traditional settings of spatial preferences and rent-seeking, under perfect and imperfect monitoring of politicians. We define stationary electoral equilibrium, which encompasses notions used by Barro (1973), Ferejohn (1986), Banks and Sundaram (1998), and others. We show that repeated elections mitigate the commitment problems of politicians and voters, so that a responsive democracy result holds under general conditions. Term limits, however, attenuate the responsiveness result. 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 economics.

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

*Duggan: University of Rochester, dugg@ur.rochester.edu
†Martinelli: George Mason University, cmarti33@gmu.edu

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 actively developing, and 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 essay, we survey and synthesize the literature on political 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 politicians are sufficiently office motivated or politicians and other citizens place sufficient weight on the future, *responsive democracy* is possible under certain conditions, in the sense that elected politicians choose policies that approach the majority winning policy. For example, when policies form a one-dimensional ideological spectrum and politician types are unobserved by voters, if elections are infinitely repeated with no term limits, then equilibrium policies converge to the median ideal policy if either office motivation or patience is sufficiently high.

Although superficially similar to median voter results in the traditional Hotelling-Downs competition framework, the mechanism underlying responsive democracy is different: in the electoral accountability model, candidates cannot make binding campaign promises, and they do not compete for votes in the spirit of Hotelling-Downs; 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 to shirk 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 we frame our discussion in terms of representative democracies, and consequently focus on elections as the means to discipline politicians, these incentives are to some extent at work in nondemocratic polities via accountability in the form of protests, coups, and revolutions.

[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.

Convergence to the median policy in dynamic elections arises from politicians' 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, dynamic elections engender the possibility of responsive democracy, despite the paucity of instruments available to voters 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 explore each informational assumption under general specifications of preferences, but we survey the most relevant specifications from the point of view of existing work on the topic.

Accordingly, throughout this essay, 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 by long-run incentives in a dynamic framework. 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 dynamic 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 among 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 expense 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 captured elected politicians from serving the interests of an elite, 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 to democratic responsiveness when politicians in office face the realistic 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.

In the presence of term limits, electoral incentives still induce some responsiveness, in the sense that some politician types move from their ideal policies in the first term of office, but the strongest form of the responsiveness result, in which equilibrium policy choices converge to the median voter's ideal policy, fails in this setting. When a two-period 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, as it is still conceivable that the responsiveness result could hold for policy choices of first-term office holders; this could arise if the equilibrium threshold used by voters and the incentive of re-election induced politicians to compromise in their first terms, even though they choose their ideal policies in the second term. Now, however, 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 challengers are also expected to shirk if elected, 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, when politicians are highly office motivated. 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.

Formal proofs and technical discussion are sketched or omitted.[2] The remainder of this essay is organized as follows. Section 2 provides an overview of the classical static framework of electoral competition and sets 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. Section 5 deals with the case of perfect monitoring. 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 critical areas for future research.

2 Classical electoral competition

In this section, we review classical results in the theory of elections, and we 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 that 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;

[2]Duggan and Martinelli (2017b) is an expanded version of the current paper with many of the proofs omitted from the present discussion. The interested reader may consult it for details.

and parties seek to maximize their chances of winning the election. We denote the policy space by X, and for simplicity we assume throughout this essay that $X \subseteq \mathbb{R}$. A continuum of voters, N, is partitioned into a finite set $T = \{1, \ldots, n\}$ of types,[3] 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., $\hat{x}_1 < \hat{x}_2 < \cdots < \hat{x}_n$.

(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 $u_j(x) = -(x - \hat{x}_j)^2 + K_j$, where K_j 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 $u_j(x) = -e^x + xe^{\hat{x}_j} + K_j$, which determines ideal policy $\hat{x}_j$ and also satisfies (A1) and (A2).

The distribution of types in the electorate is given by $(q_1, \ldots, q_n)$, 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,[4] in the sense that a 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

[3] Finiteness of the set of types is assumed only for simplicity. Results stated in this section carry over, with some technical modifications, to the framework with a continuum of types.

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

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:[5]

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

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 median 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 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.

[5] If indifferent, voters of the median type split their votes to create a tie, which is decided by the toss of a fair coin.

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, i.e., "ego rents," 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 extends 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 now strengthen (A3) to

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

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 mixing 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 $\hat{x}_a = -\hat{x}_b < 0$; 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 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 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 the first order condition*

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

Intuitively, the marginal disutility from locating closer to the ex ante ideal policy of the median voter in case of winning, $u_a'(x)$, multiplied by the equilibrium probability of winning, $1/2$, is just offset by the increase in the probability of winning the election, $f(0)$, times the net gain from winning. 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 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:

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

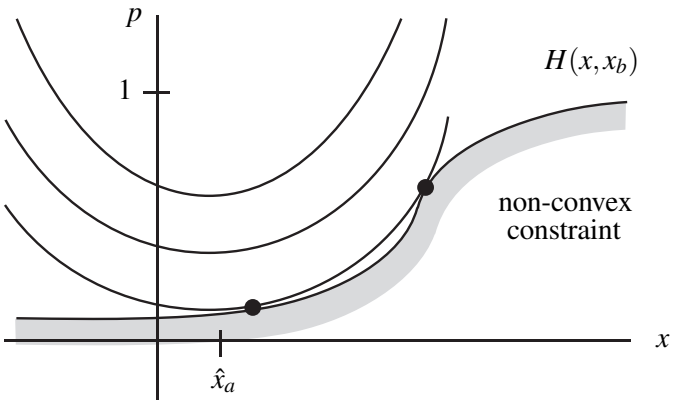

Figure 1: Multiple best responses

The objective function of the transformed problem continues to be concave, and and we only require that $\ln(H(x,x_b))$ be concave for the constraint set to be 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 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 result 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 $v: X \to \mathbb{R}$

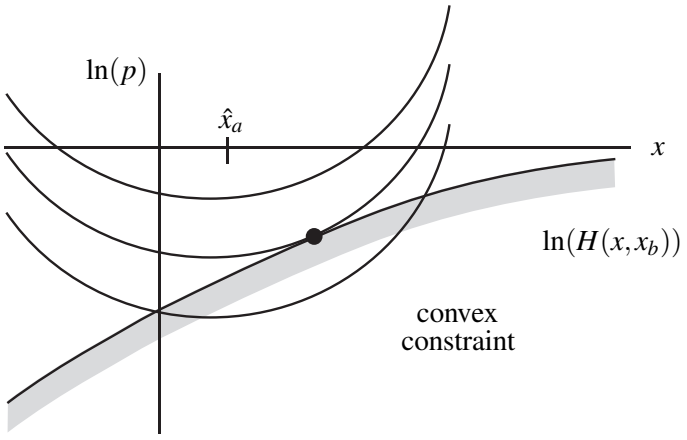

Figure 2: Log concavity

and $c:X \to \mathbb{R}$ such that for all $x \in X$,

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

where $0 < \theta_1 < \theta_2 < \cdots < \theta_n$.

While (A2) implies that preferences over policies are ordered restricted, (A5) implies that the difference in expected utility between any two lotteries over policy is monotonic in the type j, so that preferences over lotteries on policies are order restricted.[6] Thus, the median voter type acts as a representative voter in pairwise votes over lotteries.

As a special case of (A5), we obtain *quadratic utility*

$$u_j(x) \quad = \quad -(x-\hat{x}_j)^2 + K_j \tag{1}$$

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*

$$u_j(x) \quad = \quad xe^{\hat{x}_j} - e^x + K_j \tag{2}$$

by setting $v(x) = x$, $c(x) = e^x$, $\theta_j = e^{\hat{x}_j}$, and $\kappa_j = K_j$. 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. For instance, in Basseto and Benhabib's (2006) economy, all

[6]Under (A5), the difference in expected utility from any two lotteries, say λ and λ', is $\theta_j(\mathbb{E}_\lambda[v(x)] - \mathbb{E}_{\lambda'}[v(x)])$ plus a constant. This difference is monotonic in type, and thus the set of types with a given strict preference between λ and λ' is connected. See Duggan (2014b) for further details.

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.

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; then the discounted utility (normalized by $(1-\delta)$) from a sequence of policies is mathematically equivalent to the expected utility from a particular lottery. To see this more clearly, consider a sequence that alternates between two policies, x and y. The (normalized) discounted utility from this sequence for a type j voter is,

$$(1-\delta)\left[\frac{u_j(x)}{1-\delta^2} + \frac{\delta u_j(y)}{1-\delta^2}\right],$$

which is the expectation of u_j with respect to the lottery that places probability $1/(1+\delta)$ on x and probability $\delta/(1+\delta)$ on y. This observation generalizes to arbitrary sequences, 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 equilibrium path of play, and in equilibrium both parties locate at the median; but

when players place more weight on future periods than the current one, a folk theorem holds: arbitrary paths of policies can be supported in equilibrium. Alesina (1987) 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 a temporal dimension into the electoral framework, however, dynamic incentives rise to the fore. These incentives can, in turn, interact with information considerations—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—and play an important role in escaping the shirking equilibrium. The literature on electoral accountability, which is the subject of the remainder of this review, addresses these issues: elections are modeled as a dynamic 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

In this section, we introduce dynamics in the form of the two-period model of electoral accountability. We define the concept of electoral equilibrium, and we characterize equilibria in the pure adverse selection model and in the model with imperfect monitoring.

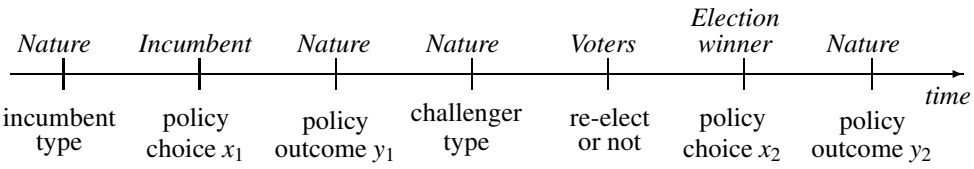

Figure 3: Timeline in two-period model

3.1 Timing and preferences

This section 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 $T = \{1,\dots,n\}$, with $n \geqslant 2$ and $q_j > 0$ denoting the fraction of type $j \in T$ in the population. We assume a generic distribution of types among voters, so there is a unique median type m. There are now two periods, $t = 1,2$. In period 1, a politician is randomly drawn from the population of citizens to hold office, with each type j having probability $p_j > 0$, and the politician chooses a policy $x_1 \in X$, where X is a nonempty, 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 convex (possibly unbounded) outcome space $Y \subseteq \mathbb{R}$. 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 $f(y|x)$.

Neither the incumbent nor the challenger can make binding promises before an election. Voters cannot commit their vote either, 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, which is private information. Once in office, the incumbent chooses the first-period policy action x_1. Next, a publicly observed outcome y_1 is realized from $F(\cdot|x_1)$. Nature draws the challenger type, and voters participate in a majority-rule election and decide to

re-elect the incumbent or to replace the incumbent with the challenger. Finally, the winner of the election chooses the second-period policy x_2, and the policy outcome y_2 is realized from $F(\cdot|x_2)$.

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 assume no discounting, as this provides the sharpest electoral incentives and allows us to compare policy responsiveness in the two-period model with the positive results (Propositions 5.3 and 6.9) for the infinite-horizon model as citizens place greater weight on the future.

We consider two *environments* (specifications of payoffs and information structure in the model) that are prominent in the literature:

Spatial preferences In this environment, 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 $Y = [\underline{x}, \overline{x}]$ and conditional on policy choice x, the outcome is $y = x$ with probability one. 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, and $\theta_n > \theta_{n-1} > \cdots > \theta_1 > 0$ and $\kappa_1, \ldots, \kappa_n$ are type-dependent parameters. To ensure interior ideal policy choices, we assume $\theta_1 v'(\underline{x}) > c'(\underline{x})$ and $\theta_n v'(\overline{x}) < c'(\overline{x})$. Without loss of generality, we also 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 explained in the previous section, two useful examples are quadratic utility and exponential utility defined by (1) and (2), respectively, with constant K_j appropriately chosen. Since (A1) and (A5) are satisfied, voter preferences over policy lotteries are order restricted, which ensures the median type is pivotal in pairwise voting. 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 $\theta_n > \theta_{n-1} > \cdots > \theta_1 > 0$ and $\kappa_1, \ldots, \kappa_n$ 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:

$$w_1(\hat{x}_1) \leqslant w_2(\hat{x}_2) \leqslant \cdots \leqslant w_n(\hat{x}_n).$$

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. Again, we admit the quadratic and exponential functional forms, after a suitable normalization: these functional forms are

$$w_j(x) = 2x - \frac{1}{\hat{x}_j}x^2 + K_j \quad \text{and} \quad w_j(x) = -e^{x-\hat{x}_j} + x + K_j,$$

respectively. Note that we can assume politicians share the voters' preferences over policy by setting the term $v(x)$ equal to the expected utility from policy outcomes generated by the choice x, i.e., $v(x) = \mathbb{E}[u(y)|x]$, 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 a distribution over policy choices in period 1 and a distribution over policy choices in period 2 for each possible previous policy choice and observed outcome.[7] For tractability, we impose the restriction that the distribution π_j^1 has

[7]Measurability of strategies and beliefs will be assumed implicitly, as needed, without further mention.

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

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

where $\rho_j(y)$ is the probability of voting 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 signal.

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 would emerge from 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. 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.[8] 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

[8] In the model with perfect monitoring, voters' beliefs place probability one on the policy chosen. Equivalently, policy outcomes are observable and chosen directly by the office holder.

$V_m^I(y_1) > V_m^C$ 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 $V_m^I(y_1) \geqslant V_m^C$.

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:

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

We say an equilibrium is *monotonic* if the acceptance set is closed, and if for every policy outcome belonging to the acceptance set, moving in a way that increases the median voter's utility maintains inclusion in the acceptance set. Formally, we break this definition into two versions, depending on the environment. In the spatial preferences environment, if $y \in A$ and z is between y and $\hat{x}_m$, then $z \in A$; and in the rent-seeking environment, if $y \in A$ and $z \geqslant y$, then $z \in A$. In particular, A is convex, and in the spatial preferences model, 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 $A = \varnothing$ 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.

3.3 Adverse selection

In this subsection, we focus on perfect monitoring in the spatial preferences environment, so that the realized policy outcome in the first period 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 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. Otherwise, a type j office holder's maximum payoff from choosing a policy in the acceptance set, and her maximum payoff from shirking, are

$$\max_{x \in A} u_j(x) + u_j(\hat{x}_j) + 2\beta \quad \text{and} \quad u_j(\hat{x}_j) + V_j^C + \beta,$$

respectively. Thus, for a politician whose ideal policy is not acceptable, the compromise decision hinges on the comparison of $\max_{x \in A} u_j(x) + \beta$ vs. V_j^C. These observations lead to a partition of types as follows: 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 politicians in the set W as "winners," in the set C as "compromisers," and in the set L as "losers."

In this section, we consider simple conditions such that equilibrium incentives lead all politician types to 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

$$\theta_j(v(\hat{x}_m) - \textstyle\sum_k p_k v(\hat{x}_k)) + \beta \quad \geqslant \quad c(\hat{x}_m) - \sum_k p_k c(\hat{x}_k) \tag{3}$$

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, due to risk aversion of politicians. In general, 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 by the incumbent 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, assume the extreme types are willing to compromise, i.e., condition (B1) holds. Then there is an electoral equilibrium such that every politician type chooses the median policy in the first period.*

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 the median voter type strictly prefers a type m politician to an unknown challenger, which may be of a different type, 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 $A = [\underline{x}(\beta), \overline{x}(\beta)]$ defined by two endpoints, with a focus on the lower endpoint. First, note that the median type is obviously willing to compromise at the median, so (3) holds for $j = m$.

As a consequence, if the inequality fails for one extreme type, say $j = 1$, then it holds for the other, say $j = n$. The case in which both extreme types are willing to compromise to the median is covered by (B1), so we assume:

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

which implies that the term $v(\hat{x}_m) - \sum_k p_k v(\hat{x}_k)$ in (3) is strictly positive. 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.,

$$u_\ell(\hat{x}_m) + \beta \quad > \quad V_\ell^C, \tag{4}$$

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:[9] it is the greatest solution to

$$u_\ell(x) + \beta \quad = \quad V_\ell^C. \tag{5}$$

Note that $\underline{x}(\beta)$ increases with β until $\underline{x}(\beta) = \hat{x}_m$, at which point the first case obtains. Moreover, it can be shown that because the type j politician is indifferent between shirking and compromising at $\underline{x}(\beta)$, types $1, \ldots, \ell - 1$ 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 voter's expected payoff of re-electing the incumbent is at least as great as that from drawing the challenger at large:

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

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. The usefulness of the condition is that it ensures that if strategies are such that all types equal to or greater than the median pool at $\hat{x}_m$, then voters are willing to re-elect the incumbent after observing a choice of the median policy in the first period.

[9] 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.

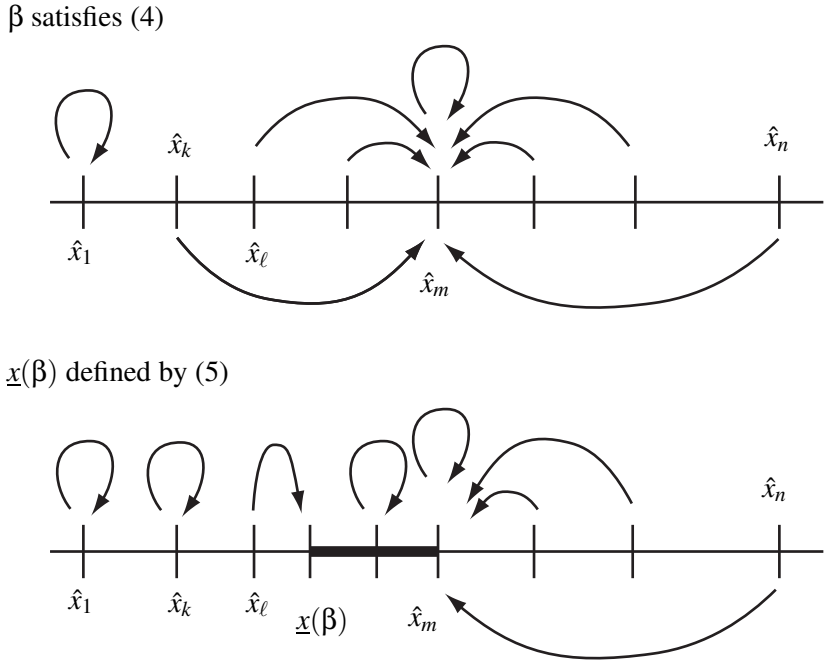

Figure 4: Partial compromise equilibria

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.

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

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

(ii) otherwise, $\underline{x}(\beta)$ *is the greatest solution to (5), and there exists* $k \in \{\ell+1,\ldots,m\}$ *such that*

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

To construct an equilibrium for the case in which β satisfies (4), 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. 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.

For the case in which $\underline{x}(\beta)$ is defined by (5), 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:

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

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.* Behind the result is the fact that when office benefit is large, no politician type will choose a policy that leads to below average expectations on the part of voters; but if the equilibrium acceptance set is a non-degenerate interval, then at least one endpoint of the acceptance set will lead voters to update negatively. Thus, the acceptance set must be a singleton and, by monotonicity of equilibrium, it consists of just the median policy.

Proposition 3.3 *In the two-period model of adverse selection, assume (B1) and (B4) hold. Then in every electoral equilibrium, every politician type chooses the median policy in the first period.*

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 (3) 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. 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.

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 Adverse selection and moral hazard

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.[10] 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

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

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.

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.[11] 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 includes Austen-Smith and Banks (1989), who investigate the voters' ability to discipline politicians when all politicians have the same preferences, so that the model is one of pure moral hazard, and Barganza (2000), who studies a two-type model in which politicians differ in ability.

The development in this subsection is based on Duggan and Martinelli (2017a), 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 of the density on outcomes, so, abusing notation slightly, the density can be written $f(y|x) = f(y-x)$ for strictly positive density $f(\cdot)$ with zero mode, 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

[11] 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.

assume $Y = \mathbb{R}$ and

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

As an example, $f(\cdot)$ may be a normal density with mean zero and arbitrary 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., $A = [\overline{y}, \infty)$, or they always re-elect the incumbent and $\overline{y} = -\infty$. 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

$$w_j(x) + (1 - F(\overline{y} - x))[w_j(\hat{x}_j) + \beta] + F(\overline{y} - x)V^C. \tag{6}$$

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

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

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.,

(C3) $$w_1(\hat{x}_1) + \beta > V^C.$$

In words, type 1 politicians have the greatest cost of providing the policy action, but even they would rather provide (at a cost) their ideal policy than have a stochastically better challenger replace them. 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 of (6) is

$$w_j'(x_j^*) = -f(\overline{y} - x_j^*)[w_j(\hat{x}_j) + \beta - V^C]. \tag{7}$$

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 politician's increased chance of re-election. By (C3), the right-hand side of (7) is negative, and we see that for an arbitrary cutoff, the politician optimally exerts a positive amount of effort, i.e., chooses a policy strictly than her ideal policy, 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

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

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

That is, the set of policies above the ideal policy of a politician such that the coefficient of absolute risk aversion $w_j''(x)/w_j'(x)$ is smaller than the likelihood ratio $-f'(\overline{y}-x)/f(\overline{y}-x)$ is convex. To see the permissiveness of this condition, note that by log concavity of $f(\cdot)$, the likelihood ratio is strictly decreasing in x, and thus (C4) is satisfied if the coefficient of absolute risk aversion does not decrease too

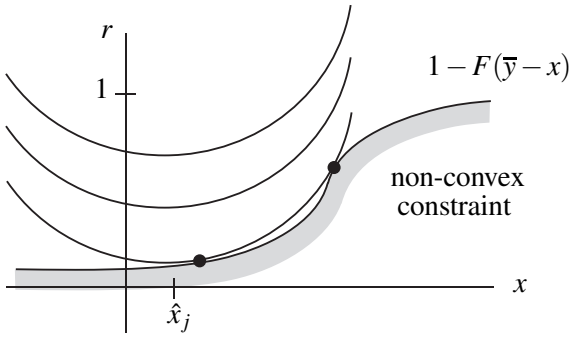

Figure 5: Politician's optimization problem

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 $(x-\hat{x}_j)^{-1}$, and when the density f is normal, the likelihood ratio is linear in 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 $(1-\exp(\hat{x}_j-x))^{-1}$, 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 in the next result, which implies that for arbitrary cutoffs, each type of office holder has at most two optimal policies. Intuitively, assuming the type j politician has two optimal policy choices, say $x_{*,j} < x_j^*$, then the politician is indifferent between the policies: the higher policy 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 is "taking it easy."

Proposition 3.4 *In the two-period model of adverse selection and moral hazard, 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 (6).*

We can illustrate Proposition 3.4 assuming the normal density with standard deviation σ and a variant of exponential utility with decreasing preferences over policy, i.e., $w_j(x) = -e^x + 1$. Then the first order condition is

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

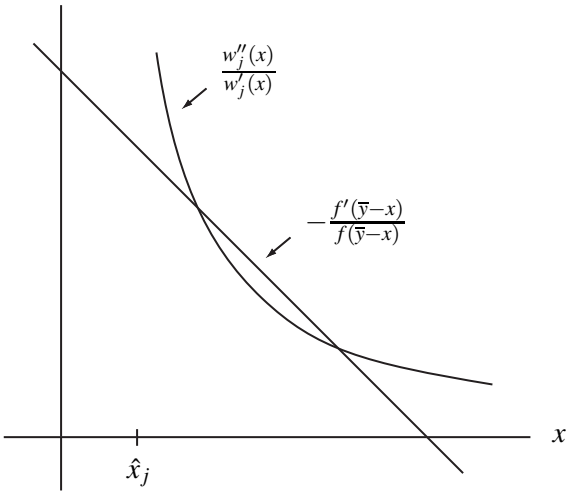

Figure 6: Quadratic-normal special case

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

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

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.[12] Alternatively, the solutions are real if the variance of the observed outcome is sufficiently small. Note that even though the politician's utility is decreasing in effort, the greatest optimal solution increases without bound as the office benefit becomes large; this suggests the possibility of responsive democracy when electoral incentives are large, a topic to which we return below.

To establish existence of equilibrium, we rely on two intermediate observations. First, for each politician type $j \in T$ and all $x, z \in X$ with $x > z$, the utility difference

$$U_j(x, 1 - F(\overline{y} - x)) - U_j(z, 1 - F(\overline{y} - z))$$

is strictly increasing in j, i.e., the objective function of politicians is supermodular. This implies an ordering property over policy choices; given an arbitrary value $\overline{y}$ of

[12]This is an artifact of the functional form modification, which drops the additive x term from exponential utility.

the cutoff, the optimal policy choices of the types are strictly ordered by type, i.e., for all $j < n$,

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

Second, given arbitrary policy choices $x_1 < x_2 < \cdots < x_n$ of the politician types in the first period, there is a unique outcome $y^*(x_1, \ldots, x_n)$ such that conditional on realizing this value, voters are indifferent between re-electing the incumbent and electing a challenger. Moreover, this extends to the case of mixed policy strategies $\pi_1, \ldots, \pi_n$. As long as the optimal policy choices of the types are ordered by type, there is a unique solution $y^*(\pi_1, \ldots, \pi_n)$ to the equation

$$V^I(y) \quad = \quad V^C. \tag{8}$$

Using the two properties described above, Duggan and Martinelli (2017a) establish existence of electoral equilibrium, along with a characterization of equilibria. Note that even in the two-period model, where second-period policies are pinned down by end-game effects, the equilibrium existence argument must address 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.

Proposition 3.5 *In the two-period model of adverse selection and moral hazard, assume (C1)–(C4). Then there is an electoral equilibrium, and in every electoral 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 policy choices, i.e., $x_{*,1} \leqslant y^* \leqslant x_n^*$.

By the MLRP, policy outcomes below the cutoff y^* influence the updating of voters' beliefs by increasing weight on worse types, while policy outcomes above the cutoff do the opposite. In equilibrium, these forces are balanced equally at the cutoff, fulfilling the indifference condition in (8). For the special case of two types using pure policy strategies, the voters' cutoff y^* is simply the solution to $\mu_T(2|y^*) = p_2$, i.e., conditional on the cutoff, the probability the incumbent is the

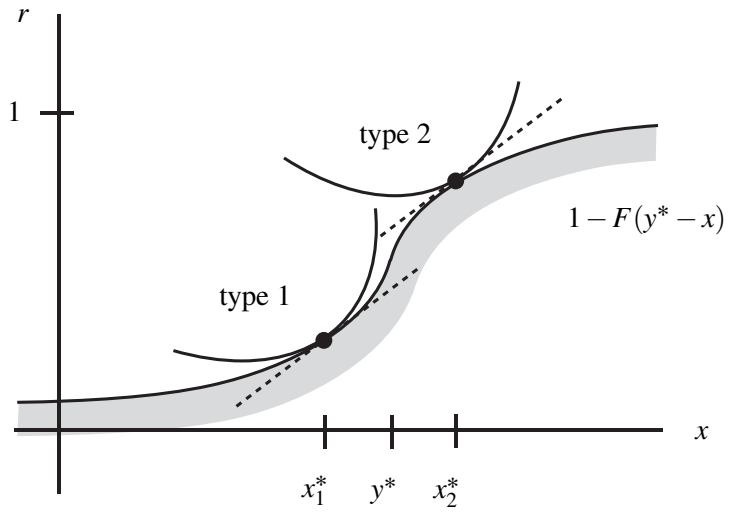

Figure 7: Electoral equilibrium

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^* solves

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

and using $p_1+p_2=1$, this implies $f(y^*-x_1)=f(y^*-x_2)$. If the density $f(\cdot)$ is symmetric around zero, then the cutoff is simply the midpoint of the politicians' choices, i.e.,

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

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, the indifference curves of the politician types become flatter, so, ceteris paribus, the optimal policy choices of the politicians move to the right. In equilibrium, of course, 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.

We have not yet touched on the possibility for responsive democracy in the two-period model with imperfect monitoring, 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 (2017a), 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, and that all above average types exert unbounded effort. An immediate implication, since there is at least one above average type, 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.

Proposition 3.6 *In the two-period model of adverse selection and moral hazard, assume (C1)–(C5), and let the office benefit become arbitrarily large. Then for every selection of electoral equilibria, the voters' cutoff diverges to infinity; 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* $x_{*,j} \to \infty$.

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 (2017a) 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.

The main starting point of our analysis has been the statement of equilibrium existence in Proposition 3.5, 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 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 observable 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 τ, then we refer to policy outcomes $\mathbf{y}^{\tau-1} = (y_t, y_{t+1}, \ldots, y_{t+\tau-1})$ 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 signal 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. Note that voter payoffs are derived from observable outcomes, but the observable outcome of policy choice x_t

is just $y_t = x_t$ in the model with perfect monitoring; and in the model with imperfect monitoring, recall that we can set $v(x_t) = \mathbb{E}[u(y)|x_t]$ to bring politician preferences in line with voters', if desired.

We continue to focus on the *spatial preferences* environment in the context of 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 can 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 $(\pi_j^\tau)_{\tau=1}^\infty$, where π_j^τ specifies the politician's mixture over policy choices in term τ of office as a function of personal history; we write $\pi_j^\tau(\mathbf{y}^{\tau-1}) \in \Delta(X)$ for the mixture with finite support over policy choices in term τ given personal history $\mathbf{y}^{\tau-1}$ over the first $\tau-1$ terms of office. We write $\pi_j^\tau(Z|\mathbf{y}^{\tau-1})$ for the probability that a type j politician chooses a policy in the set Z. An alternative interpretation is that $\pi_j^\tau(Z|\mathbf{y}^{\tau-1})$ 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 $(\rho_j^\tau)_{\tau=1}^\infty$, where $\rho^\tau(\mathbf{y}^{\tau-1}, y_{t+\tau}) \in [0,1]$ determines the probability that the voter votes to re-elect the incumbent as a function of the personal history $\mathbf{y}^{\tau-1}$ of the incumbent in prior terms of office and the current policy outcome $y_{t+\tau}$. And a *belief system* is a sequence $(\mu^\tau)_{\tau=1}^\infty$, where $\mu^\tau(\cdot|\mathbf{y}^{\tau-1}, y_{t+\tau})$ is a probability distribution on $T \times X$ as a function of the personal history of the incumbent and current policy outcome. 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.

A strategy profile $\sigma = ((\pi_j^\tau)_{\tau=1}^\infty, (\rho_j)_{\tau=1}^\infty)_{j\in T}$ is *sequentially rational* given belief system μ if for every term of office τ and every personal history $\mathbf{y}^{\tau-1}$, no politi-

cian can gain by deviating to a different policy choice, and for all policy outcomes $y_{t+\tau}$, 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 τ and every personal history $\mathbf{y}^{\tau-1}$ and outcome $y_{t+\tau}$ on the path of play, $\mu^{\tau}(j,x|\mathbf{y}^{\tau-1},y_{t+\tau})$ is derived using Bayes rule. A *perfect Bayesian equilibrium* is a pair (σ,μ) 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 two type j politicians with different personal histories, 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 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 (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 $V_j^I(b,y|\psi)\geqslant V_j^C(\psi)$. 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.,

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

Our main equilibrium concept for the infinite-horizon model is defined as follows: we say ψ is a *stationary electoral equilibrium* if it is a perfect Bayesian equilibrium that is stationary, deferential, and monotonic. 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 satisfy (A5). Thus, the median voter is a representative voter over lotteries on policies; in the dynamic electoral framework, payoffs from the incumbent and challenger can each be reduced to the expected utility from a lottery over policy, by arguments of Section 2.4. Given a stationary electoral equilibrium, it follows that the median voter type is pivotal in elections, in the sense that the incumbent 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, given any stationary electoral equilibrium, the median voter type is a representative voter, i.e., the incumbent wins if and only if the median voter weakly prefers the incumbent to the challenger.*

Proposition 4.1 holds because the continuation values of electing the incumbent and the challenger can be rewritten as expected utilities with respect to two lotteries over policies; using (A5), it follows that the median type m is representative. Thus, the *acceptance set* of policy outcomes that lead to re-election given updated prior b is

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

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

We 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. Proposition 4.1 shows that 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.

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. 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 a hypothetical optimal retention problem. That is, *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, 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, thereby removing reputational considerations from the optimization problem of politicians, 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. The result is that equilibria in this class become degenerate: office holders always choose their ideal policies.

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.*

To see why the result holds, note that since the voters' beliefs about the incumbent's type do not depend on her policy choices in this formulation of the model, stationarity implies that the expected discounted utility of voters from a type j incumbent is a constant independent of the history of play. By the assumption of deferential strategies, each voter's behavior will be independent of the policy choice of the politician, but this means that electoral incentives have been removed: the only reason a politician would choose a policy other than her ideal policy is to increase her chances of re-election, but absent reputational concerns, a politician's electoral prospects are not affected by her policy choice, so politicians simply choose their ideal policies. Thus, elections 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 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 reputational incentives, we allow multiple politician types, and we assume these are private information—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 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 maximal acceptance set that is consistent with this criterion is

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

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 ψ 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, there exists a simple electoral equilibrium.*

The proof 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. Uniqueness of equilibrium is proved in symmetric models with continuous types by Duggan (2000) and Bernhardt et al. (2009a).

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

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

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

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

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 $\mathscr{r} = \max W$ 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 adverse selection, let $\psi = (\sigma, \mu)$ be a simple electoral equilibrium. Then there exist integers k' and k'' such that $k' \leqslant \ell \leqslant m \leqslant \mathscr{r} \leqslant k''$ and*

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

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. To see this, define $u(x|\theta) = \theta v(x) - c(x)$ and $\hat{x}(\theta) = \arg\max_x u(x|\theta)$, and let $\mathbb{E}[v(x)]$ and $\mathbb{E}[c(x)]$ be the expected discounted value of $v(x)$ and $c(x)$ on the equilibrium path; then for given x' the indifference condition

$$\theta(v(x') - \delta\mathbb{E}[v(x)]) + \beta \quad = \quad (1-\delta)u(\hat{x}(\theta)|\theta) + c(x') - \delta\mathbb{E}[c(x)]$$

can be shown to be satisfied by at most one θ on the same side of θ' such that $x(\theta') = x'$. 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 8, where types $k' = 2$ and $k'' = 5$ mix in equilibrium.

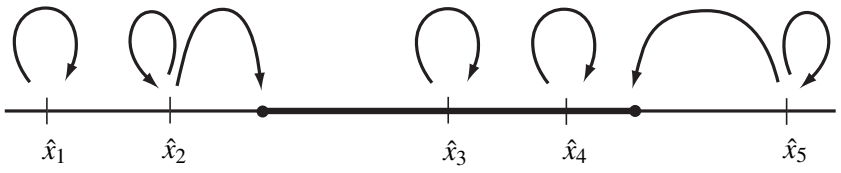

Figure 8: Partitional structure

The preceding observation implies that a simple electoral equilibrium can be characterized as a solution to a system of equations, depending on the cutoff types $k', \ell, \hat{\imath}, k''$. This characterization implies that the model is computationally tractable.

5.3 Responsive democracy

Next, we investigate the implications for equilibrium outcomes when politicians are highly office motivated, and in particular we focus on the possibility of obtaining a dynamic version of the Downsian median voter 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

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

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 $j \in T$,

$$\frac{\beta\delta}{1-\delta} \quad \geqslant \quad u_j(\hat{x}_j). \tag{12}$$

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. Using this logic, Banks and Duggan (2008) provide the following strong responsive democracy result for simple electoral equilibria. Assuming $\beta\delta > 0$, the assumptions of the result are satisfied if politicians are sufficiently office motivated (β large) or citizens are sufficiently patient (δ close to one).

Proposition 5.3 *In the infinite-horizon model of adverse selection, assume inequality (12) holds for all types. In every simple electoral equilibrium, each type j politician chooses the median policy in the first term in office and in every future term along the path of play.*

Although Proposition 5.3 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. Driving this is the idea that although the acceptance set may be a non-degenerate interval, the width of that interval relies on a wedge created by the voters' cost of trying a challenger; as voters become patient, this cost goes to zero, and the acceptance set must collapse to a singleton consisting of the median policy.

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-period 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 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.

Behind the need for mixed voting is the fact that given the opportunity to secure re-election, highly office-motivated politicians will do so, but this means that voters have an incentive to replace a first-term incumbent with a fresh challenger. Specifically, 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's chance of re-election after choosing the median policy 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.

Electoral incentives still induce some compromise, in the sense that some politician types move from their ideal policy in the first term of office, but the full responsiveness result cannot be obtained in the term limit model. To see this, suppose there is an equilibrium σ 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 expected payoff to the median voter from equilibrium policy choices of first-term office holders cannot approach the voter's ideal payoff. 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 et al. (2009a) 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 stronger incentives to moderate policy choices in the partisan model, as 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 $\beta/(1-\delta)$ of potential office benefits can be used to induce an office holder to choose the worst policy in the interval X. Duggan (2014a) shows that in the model of pure adverse selection, equilibrium behavior can 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

We now extend the pure adverse selection model of Section 5 to the case of imperfect monitoring. As in Subsection 3.4, a policy choice x stochastically determines an outcome y, which is realized from the density $f(y-x)$. We retain assumptions (C1) and (C2); that is, f satisfies the MLRP, and extreme outcomes are arbitrarily informative. 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, to an extent depending on their types,

prefer lower exertion of effort. By the MLRP, higher outcomes are evidence of higher effort choices, consistent with the monotonicity property that equilibria are characterized by a cutoff outcome $\overline{y}$ that is necessary and sufficient for re-election of the incumbent. The question of responsive democracy then 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 a large office benefit.

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 voters to set an arbitrary cutoff in the space of outcomes. Since voters are in fact indifferent between the two candidates, every cutoff is time-consistent. Following the literature, we analyze the optimal cutoff for voters, an approach that amounts to a selection of equilibria. Of course, it may be unrealistic to assume that the electorate coordinates on the equilibrium that is optimal for voters. Nevertheless, the selection criterion establishes a (tight) upper bound to voters' welfare in equilibrium.

We drop Ferejohn's productivity shock and examine the voters' optimization problem in the simple rent-seeking environment, effectively allowing voters to commit to a cutoff before the beginning of the game; again, since voters are always indifferent between the candidates, this commitment assumption does not rely on an outside enforcement mechanism. We find that voters can induce positive effort, and that *greater office-motivation leads to arbitrarily high effort levels, delivering responsive democracy* in the optimal equilibrium for voters. The logic is simple: if voters use 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.

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 $y \geqslant \overline{y}$, where y is the realized outcome and $\overline{y}$ 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 $\overline{y}$. Then the continuation value of a challenger is simply the expected utility generated by mixtures over policy, i.e., $V^C(\psi) = \frac{1}{1-\delta}\sum_x \mathbb{E}[u(y)|x]\tilde{\pi}(x)$, and the value of the office holder's optimization problem net of current office benefit, denoted $W(\tilde{\pi},\overline{y})$, satisfies

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

The first order condition for this problem is of course

$$w'(x) + f(\overline{y}-x)\delta\left[W(\tilde{\pi},\overline{y}) + \beta - \frac{\sum_x \mathbb{E}[u(y)|x]\tilde{\pi}(x)}{1-\delta}\right] \quad = \quad 0. \tag{13}$$

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

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

now parameterizing the objective function by the mixed policy strategy, $\tilde{\pi}$. A difficulty in reformulating the condition that politicians desire re-election, (C3), arises in the present setting, because the mixture $\tilde{\pi}$, affecting future payoffs in case of removal, is endogenous. We cannot rephrase the condition so that the payoff from holding office is positive for all possible values that $\tilde{\pi}$ might take, because policy is unbounded, and so the voter's expected utility is, in principle, unbounded. It suffices for the equilibrium analysis, however, to specify the condition for the case in which politicians choose their ideal policy with probability one,

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

and this 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 $\beta > c(\hat{x})$.

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

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

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, 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, then the objective function of the politician has at most two local maximizers and, therefore, at most two maximizers; a result paralleling Proposition 3.4. Of course, if there are two optimal policies, then every mixture over these 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.

Our analysis is facilitated by the following result, which shows that if the policy mixture used by politicians is degenerate on a policy $\tilde{x}$, then the value of office is decreasing in $\tilde{x}$. Abusing notation, 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 one on that choice.

Proposition 6.1 *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 $\tilde{x}$*.*

In words, if all other politicians choose policy $\tilde{x}$, then an office holder's optimal policy choices decrease with $\tilde{x}$, as the prospect of being out of office becomes more attractive. If there are multiple optimal policy choices given $\tilde{x}$, then an increase in $\tilde{x}$ may lead to discontinuities 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 $\tilde{x}$; see Figure 9.

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 (13), politicians optimally choose policies larger than $\hat{x}$. Increasing the policy used by politicians to $\tilde{x} > \hat{x}$, we see that the "best response" policy either decreases continuously until

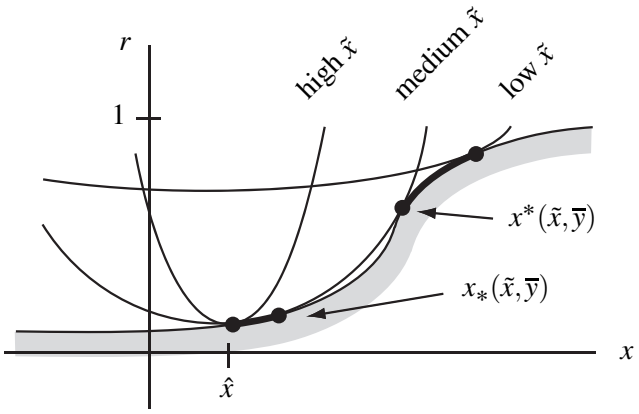

Figure 9: Multiple optimal policies

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 10. 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 mentioned above.

Proposition 6.2 *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, and $\pi^*(\overline{y})$ places positive probability on at most two policies. Moreover, $\pi^*(\overline{y})$ is a continuous function of the cutoff.*

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. 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

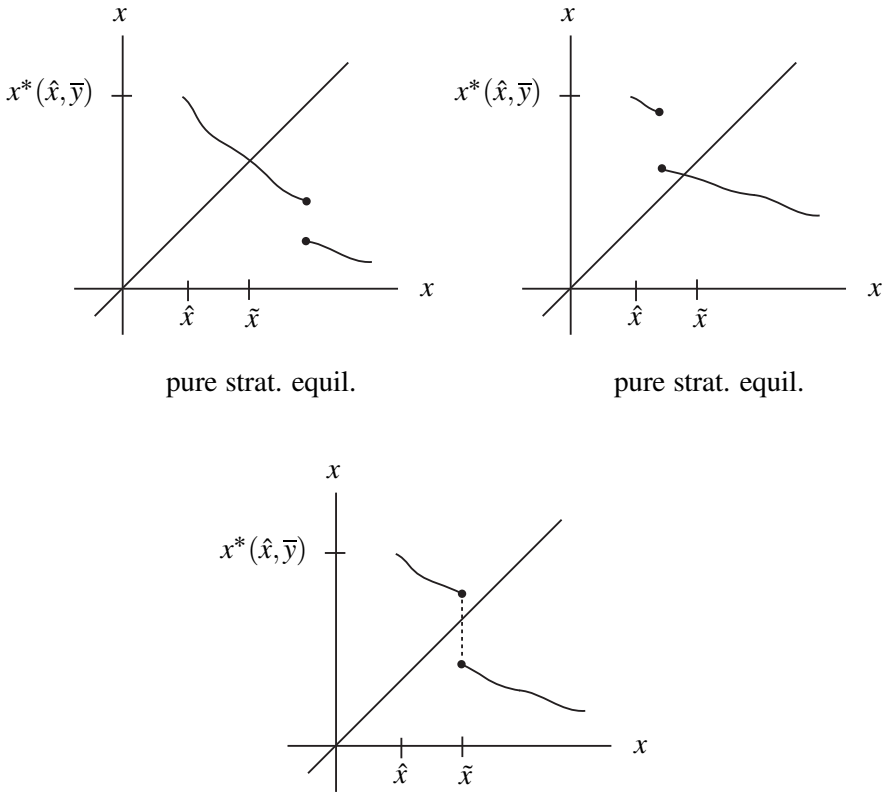

Figure 10: Existence and uniqueness

politicians' constraint to the right; see Figure 11. Second, there is an indirect effect on the net value of office, which can be positive or negative, and this is reflected in the marginal rate of substitution of the politician's objective function. These effects in turn lead to a change in the equilibrium policies. Voters do indeed have an optimal cutoff.[13] 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 11. Here, we depict the case in which politicians choose a single policy x^* with probability one given the optimal cutoff.

We have not yet considered the possibility of responsive democracy in the infinite-horizon model of pure moral hazard. The next proposition establishes a

[13] See Duggan and Martinelli (2017b).

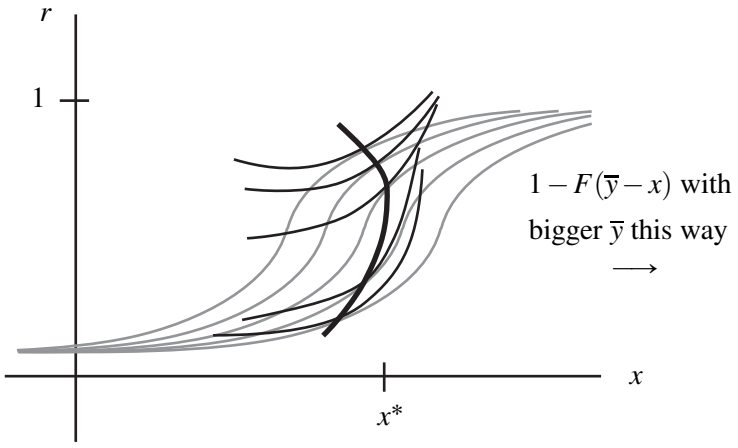

Figure 11: Locus of policy choices

strong responsiveness result when politicians are substantially office-motivated: when the office benefit β is large, the equilibrium policy choices of the politicians increase without bound for an arbitrarily fixed cutoff. Obviously, the result is reinforced if the cutoff is set optimally.

Proposition 6.3 *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 become arbitrarily large. Then the politicians' optimal policy choice increases without bound, i.e.,*

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

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, paralleling the difficulties encountered in Subsection 5.4. In particular, assume that politicians can hold office for at most K terms, and allow voters to use cutoffs $\overline{y}_1, \ldots, \overline{y}_{K-1}$ that depend on the incumbent's term of office. Assuming these cutoffs are time-consistent, in the sense that they never dictate that the electorate votes against their preferences, politicians always choose $\hat{x}$ in equilibrium. Indeed, suppose otherwise. If elected for K terms, then the politician of course shirks in the final term of office; but then there is a last term during which she exerts positive effort, and after that term, voters have an incentive to replace the incumbent with a fresh challenger. Thus, in the presence of a term limit, electoral incentives unravel, leaving only the trivial shirking equilibrium.

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.

In contrast to the positive result of Proposition 3.6 for the two-period model of moral hazard and adverse selection, we provide a bound on responsive democracy in the infinite-horizon model with no term limit: for a simplified model with two types, Proposition 6.7 establishes that the voters' continuation value of a challenger cannot, in equilibrium, exceed the discounted expected payoff from the ideal policy of the high type. Nevertheless, Proposition 6.9 delivers a qualified responsiveness result for a simplified model with two types: 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 3.6 for the two-period model, where the strategic structure of the game implies that office-motivated politicians' effort levels increase without bound; this issue does not arise in Proposition 5.3 for the model of pure adverse selection, where the ideal policy of the "best type" of politician is the median ideal policy.

The focus on monotonic, deferential strategies 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 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' updated beliefs about an office holder's type conditional on policy outcome y and given beliefs b at the beginning of the period. 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. Thus, the indifference condition determining the voters' cutoff is

$$V^I(\mu_T(b,y)|\sigma) \quad = \quad V^C(\sigma),$$

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. 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 and strictly prefer to reject the incumbent following low outcomes.

Proposition 6.4 *In the infinite-horizon model of one-sided learning, assume (C1) and (C2). In every stationary electoral equilibrium σ, we have $-\infty < y^*(p) < \infty$.*

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. 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.* To avoid perverse incentives of politicians 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 type 2 politician. As we argue below, $\mathbb{E}[u(y)|\hat{x}_2]$ 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. This means that an office holder is willing to choose a policy distinct from her ideal policy in order to gain re-election, but in contrast to the two-period model, we cannot take for granted that equilibrium policy choices will exceed the politicians' ideal policies.

Existence of stationary electoral equilibrium is a thorny issue; leaving it aside, it is instructive to consider the incentives of politicians to exert effort as office benefit becomes large in the two-type model, *assuming equilibria exist.* 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

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

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.

The following result, which can be obtained by inspection of the politicians' first order condition, establishes that in a reputation-monotonic equilibrium of the two-type model, regardless of the voters' beliefs about the incumbent's type, the politician exerts a positive amount of effort.

Proposition 6.5 *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, we have $\pi_j((\hat{x}_j,\infty)) = 1$.*

As a 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 6.6 *In the infinite-horizon model of one-sided learning with two types, assume (C1)–(C3). 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 $x < z$.*

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 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. An implication is that the net value of office is indeed positive in equilibrium.

Proposition 6.7 *In the infinite-horizon model of one-sided learning with two types, assume (C1)–(C3) hold. For all levels of office benefit and all discount factors, in every reputation-monotonic stationary electoral equilibrium, the voters' continuation value of a challenger is no more than the expected discounted utility from the ideal policy of the type* 2 *politician:*

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

The proof of the above 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.

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. By Proposition 6.5, the incumbent's policy choice always exceeds her ideal policy, and the principle of optimality for dynamic elections, Proposition 4.2, implies that the voters' equilibrium expected payoff cannot fall below this bound. This logic delivers the following result.

Proposition 6.8 *In the infinite-horizon model of 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 6.7, 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 3.6 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.

Proposition 6.9 *In the infinite-horizon model of 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:*

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

The proof again 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. By the optimality principle for dynamic elections, the equilibrium strategy used by the representative voter can do no worse than the non-stationary strategy described above, and the result follows.

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' updated beliefs. In particular, if the realized policy outcome falls below a given trigger 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 single bad outcome will be realized and cause the voter to 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 two-type 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 6.7 for the two-type model of moral hazard and adverse selection with no term limit, but the qualified responsiveness demonstrated in Proposition 6.9 for the model without term limits fails in the model with term limits. Thus, the commitment problem of voters has a significant impact on the effectiveness of electoral incentives, and this impact is magnified in the presence of term limits.

In the model with a two-period term limit, we extend our definition of stationary strategy profile to allow for politicians to condition their choices on the term of office, as 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. 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. Third, as always, updating of voter beliefs follows Bayes rule.

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 stated for an arbitrary number of politician types,

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

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 will see that $\mathbb{E}[u(y)|\hat{x}_n]$ bounds 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), for arbitrary cutoff and continuation value, a politician exerts positive effort in the first term.

We impose condition (C4) to obtain the result that given a such a cutoff and a challenger 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. We let $x_j^*(\overline{y},V)$ and $x_{*,j}(\overline{y},V)$ denote the greatest and least optimal policies, respectively, of the type j politician in the first term of office. Again, the utility difference $U_j(x,1-F(\overline{y}-x);V)-U_j(z,1-F(\overline{y}-z);V)$ is strictly increasing in j for $x>z$. It follows that optimal policies are strictly ordered by type, i.e.,

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

Now consider mixed policy strategies $\pi_1^1,\ldots,\pi_n^1$ with supports that are strictly ordered according to type. The induced cutoff for voters is the unique solution in $y^*(\pi_1^1,\ldots,\pi_n^1,V)$ to the equation

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

reflecting the fact that a re-elected incumbent chooses her ideal policy and is replaced by a challenger.

Existence of equilibrium requires a fixed point argument, but an additional complication is that we must now complete an intermediate step, in which mixed policy strategies $\pi_1^1,\ldots,\pi_n^1$ and a cutoff $\overline{y}$ determine the continuation value of a challenger, $V=V^*(\pi_1^1,\ldots,\pi_n^1,\overline{y})$. Duggan (2016) establishes existence of stationary electoral equilibria in a more general model, and that paper provides a characterization familiar from Proposition 3.5. We state it next, specializing it by adding (C4) to obtain 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. Furthermore, first-term politicians always exert a positive amount of effort, as $\hat{x}_j < x_{*,j}$, so that re-election incentives do induce some level of responsiveness.

Proposition 6.10 *In the infinite-horizon model of one-sided learning with two-period term limit, assume (C1)–(C4). Then there is a stationary electoral equilibrium, and in every stationary electoral 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 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* $y \geqslant y^*$.

Recall that in the two-period model of adverse selection and moral hazard, Proposition 3.6 establishes a responsive democracy result as politicians become highly office motivated. Specifically, as office benefit grows, voters become arbitrarily demanding, and the policy choice of every above average type of politician increases without bound. In the infinite-horizon model with two period term limit, the strength of electoral incentives is controlled by the product $\beta\delta$ of office benefit and the discount factor, and the next result, due to Duggan (2016), examines the effect of increased office motivation on electoral equilibria. Again, voters become arbitrarily demanding as politicians become more office motivated, and in the first term of office, the highest type of politician mixes over arbitrarily high policy choices and policies near her ideal policy, but the probability of high policy choices goes to zero. All other politician types choose policies close to their ideal policies. We will see that this reduced responsiveness has a substantial effect on voter welfare in the infinite-horizon model.

Proposition 6.11 *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* $j < n$, $x_j^* \to \hat{x}_j$.

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, but the qualitative nature of equilibria are quite different in the two models. As in the adverse selection model

with a two-period term limit, in Section 5, the *commitment problem of voters* imposes an inherent bound on policy responsiveness in the model: if newly elected challengers exerted high effort in their first term of office, then because a second-term politician would choose policy no higher than $\hat{x}_n$, voters would always have an incentive to replace an incumbent after her first term. In fact, the next result, due to Duggan (2016), establishes that the voters' expected utility from the policy choices of first-term office holders—and therefore the continuation value of a challenger—is strictly less than the expected utility from the ideal policy of the highest type.

Proposition 6.12 *In the infinite-horizon model of adverse selection and one-sided learning with two-period term limit, assume (C1)–(C4). Fix a level office benefit* $\beta \geqslant 0$*. Then there exists* $\overline{u} < \mathbb{E}[u(y)|\hat{x}_n]$ *such that for 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* $\overline{u}$*, i.e.,*

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

A similar bound was deduced in Proposition 6.7 for the two-type model of adverse selection and moral hazard with no term limit. There, however, Proposition 6.9 provided a qualified responsiveness result: the upper bound on voter utility, i.e., the expected utility from the ideal policy of the type 2 politician, is achieved as citizens become patient. In the model with two-period term limit, however, the voters' expected utility from a first-term politician is bounded uniformly below the expected utility from the ideal policy of the type n politician, regardless of the discount factor. Thus, electoral incentives are attenuated in the model with term limits, and the prospects for responsive democracy are more problematic.

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 $f_j(y-x)$. 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. 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

$$\max_{x\in X} w(x) + \delta\Big[(1-\tilde{F}(\overline{y}|x))[w(\hat{x})+\beta+\delta V^C(\sigma)] + \tilde{F}(\overline{y}|x)V^C(\sigma)\Big],$$

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:

(C3) $$w(\hat{x})+\beta \quad > \quad \mathbb{E}[u(y)|\hat{x},n].$$

Without re-stating it explicitly, we formulate (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 $x_*(\overline{y},V)$.

Techniques used in the proof for the model of adverse selection and moral hazard with one-sided learning and two-period term limit can be used to establish existence of electoral equilibrium in the symmetric learning model.

Proposition 6.13 *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 $y \geqslant y^*$.*

Paralleling Proposition 6.11, we note that as politicians become highly office motivated, politicians must place positive probability on arbitrarily high policies

and policies close to the ideal $\hat{x}$, 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 6.14 *In the infinite-horizon model of symmetric learning with two-period term limit, assume (C1)–(C4) hold, fix the discount factor δ, and let the office benefit β be arbitrarily large. Then for every selection of stationary electoral equilibria, 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.*

Finally, arguments similar to those for Proposition 6.12 imply a bound on policy responsiveness in the symmetric learning model. Again, voters cannot commit to electing an untried challenger, and this commitment problem of voters implies that the voters' expected utility from the policy choice of a first-term politician is less than the expected utility from the choice of the ideal policy by the highest politician type.

Proposition 6.15 *In the infinite-horizon model of symmetric learning with two-period term limit, assume (C1)–(C5) hold. 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 expected utility from the choice of the ideal policy by the type n politician.*

7 Applied work

The possibility of electoral incentives leading to undesirable outcomes for voters has been the object of a recent literature, mostly employing versions of the two-period model with adverse selection and moral hazard. The existence of political cycles, and the effects of these cycles on government policy and economic performance, has been the subject of an important literature in macro political economy, and these topics continue to attract attention. In addition, there is also a large and growing literature exploring the empirical evidence on the effect of term limits, information, and media on accountability. In this section, we touch on some research along these different lines, and we illustrate how the dynamic elections framework developed in this essay can be adapted to address these topics—in some cases, extending and generalizing existing applied work.

7.1 Political inefficiency

Throughout this essay, we have emphasized the disciplining effect of elections on the behavior of politicians in office and the implications for the possibility of responsive democracy. 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 3.4–3.6, 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.,

$$\mathbb{E}[u(y)|\hat{x}_1] \;<\; \mathbb{E}[u(y)|\hat{x}_2] \;<\; \cdots \;<\; \mathbb{E}[u(y)|\hat{x}_n], \tag{14}$$

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 poli-

cies 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 (14) holds. Proposition 3.5 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 3.6 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,[14] 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 6.10–6.12 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

[14] Duggan and Martinelli (2017a) 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.

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 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 inflate the economy, thereby 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 (1987,1988), 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,[15] 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

[15]Note that Persson and Tabellini (1990) use x to denote employment; we adapt notation to suit the present discussion.

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 $\theta_j - \frac{1}{2}$. 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. 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 a constant term $\sigma/2$, 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 Evidence

The empirical literature on accountability is reviewed by Ashworth (2012) and Pande (2011), the latter with a focus on natural and field experiments in developing countries; here, we summarize some of this research and provide pointers to more recent work. 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 in-

centives (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 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) or Banks and Duggan (2005); see also Alesina (1988) 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) and 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., Golosov, M., Tsyvinski, A., 2011. Political economy of Ramsey taxation. Journal of Public Economics 95, 467–475.

Acemoglu, D., Johnson, S., Robinson, J., 2005. Institutions as a fundamental cause of long-run growth. Handbook of Economic Growth 1A, 385–472.

Acemoglu, D., Robinson, J., 2012. Why Nations Fail: The Origins of Power, Prosperity, and Poverty. Crown Publishers, New York.

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., 1988. Credibility and policy convergence in a two-party system with rational voters. American Economic Review 78, 796–806.

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., 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.

Azzimonti, M., 2014. The dynamics of public investment under persistent electoral advantage. Review of Economic Dynamics forthcoming.

Banks, J., Duggan, J., 2000. A bargaining model of collective choice. American Political Science Review 94, 73–88.

Banks, J., Duggan, J., 2005. Probabilistic voting in the spatial model of elections: The theory of office-motivated candidates, in: Social Choice and Strategic Decisions: Essays in Honor of Jeffrey S. Banks. Springer.

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., 2009a. 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.

Bernhardt, D., Duggan, J., Squintani, F., 2009b. The case for responsible parties. American Political Science Review 103, 570–587.

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., Case, A., 1995. Does electoral accountability affect economic policy choices? Evidence from gubernatorial term limits. Quarterly Journal of Economics 110, 769–798.

Besley, T., Case, A., 2003. Political institutions and policy choices: Evidence from the United States. Journal of Economic Literature 41, 7–73.

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.

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.

Coate, S., Morris, S., 1995. On the form of transfers to special interests. Journal of Political Economy 103, 1210–1235.

Coughlin, P., Nitzan, S., 1981. Electoral outcomes with probabilistic voting and Nash social welfare maxima. Journal of Public Economics 15, 113–121.

Daley, B., Snowberg, E., 2011. Even if it is not bribery: the case for campaign finance reform. Journal of Law, Economics, and Organization 27, 301–323.

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., 2017a. Electoral accountability and responsive democracy. Unpublished paper .

Duggan, J., Martinelli, C., 2017b. The political economy of dynamic elections: A survey and some new results. 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., Shotts, K., 2009. Delegates or trustees? A theory of political accountability. Journal of Politics 71, 1225–1237.

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 .

Gans, J., Smart, M., 1996. Majority voting with single-crossing preferences. Journal of Public Economics 59, 219–237.

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., 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.

Osborne, M., Slivinski, A., 1996. A model of competition with citizen candidates. Quarterly Journal of Economics 111, 65–96.

Pande, R., 2011. Can informed voters enforce better governance? Experiments in low income democracies. Annual Review of Economics 3, 215–237.

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.

Roberts, K., 1977. Voting over income tax schedules. Journal of Public Economics 8, 329–340.

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.

Rothstein, P., 1990. Order restricted preferences and majority rule. Social Choice and Welfare 7, 331–342.

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., Stromberg, D., 2010. Press coverage and political accountability. Journal of Political Economy 118, 355–408.

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.