

Making Elections Work: Accountability with Selection *and* Control*

Vincent Anesi[†] Peter Buisseret[‡]

January 30, 2019

Abstract

We study dynamic electoral accountability in a long-run interaction between voters and politicians when voters are uncertain about politicians' characteristics (*adverse selection*) and their actions in office (*moral hazard*); when elections unfold over an infinite time horizon, and when voters cannot pre-commit to their electoral strategies. While existing work highlights a trade-off between elections as mechanisms for sanctioning poor performance versus mechanisms for selecting good politicians, we show that voters can obtain arbitrarily close to the first-best payoff, i.e., that they would obtain absent uncertainty about either politicians' types or actions. We characterize a decision rule for voters that achieves this payoff, imposing no bounds on history dependence and allowing no electoral pre-commitment. In particular, there need be no trade-off between selection and sanctioning, despite the coarseness of the electoral mechanism.

*We are grateful to ...

[†]Department of Economics, University of Nottingham, *Email:* Vincent.Anesi@nottingham.ac.uk

[‡]Harris School of Public Policy, University of Chicago, *Email:* pbuisseret@uchicago.edu

1. Introduction

In a representative democracy, elections are critical for enforcing the accountability of politicians to citizens. However, elections suffer from obvious deficiencies. First, they offer voters only a single instrument with which to achieve two distinct objectives: selecting good politicians and controlling their behavior in office. Second, they pit uninformed voters against politicians that have superior knowledge about both their policy choices and their own characteristics. Third, elections are coarse instruments: they allow voters to choose only whether to replace or retain an incumbent. Finally, voters cannot pre-commit to their decisions: their electoral choices must be optimal at the time they are made.

What chance, then, do voters have of simultaneously selecting competent politicians, and inducing them to serve the public interest? Existing theoretical work identifies a trade-off between *selection* and *control* (e.g., [Fearon, 1999](#), [Besley, 2006](#), [Ashworth, Bueno de Mesquita and Friedenberg, 2017](#)). The argument is straightforward: when incumbents and challengers are homogeneous (e.g., [Ferejohn, 1986](#)), voters are indifferent between retention and replacement. They can therefore commit to a reelection rule that sanctions poor incumbent performance, and thus *control* politicians' incentives. If, however, politicians are heterogenous in a way that matters for future payoffs—for example, if they are distinguished by expected ability—then voters are not indifferent and prefer to use their decision to maximize the quality of *selection*. But this removes their flexibility to control politicians by making commitments to sanction and reward past behavior.

The trade-off appears starkly in two-period frameworks—the focus of the vast majority of theoretical work on political accountability. In these contexts, there is only a *single* election, and thus a single opportunity for voters to hold politicians accountable. Yet in real-world democratic contexts elections are both *recurrent* and *frequent*. To what extent can voters do better when they have the opportunity to hold policymakers accountable not merely once, but over the course of multiple elections? Does the trade-off between selection and control remain an insurmountable obstacle?

Existing theoretical work in dynamic electoral accountability—surveyed recently in [Duggan and Martinelli \(2017\)](#)—partially exploits the potential of repeated interactions between voters and politicians. Nonetheless, it maintains a trade-off between selection and control. On the one hand, voters in a dynamic context can punish poorly performing candidates with prolonged or indefinite non-election ([Banks and Sundaram, 1993](#)) when they could not in a static election setting. The reason is that—as in any dynamic interaction—choices that are not optimal in the short-run may be

sustained by the prospect of larger future payoffs. On the other hand, given the inherent uncertainty in policymaking, even a very good incumbent realistically faces a positive prospect of failing to meet the voter's expectations when holding office. Punishing this candidate with indefinite exclusion therefore weakens future selection.

Nonetheless, the true potential of dynamic elections to achieve both selection and control remains elusive. The reason is that existing theoretical work under-exploits the potential of repeated interactions between voters and politicians. It does this by limiting the extent to which voters can condition their strategies on the full history of their interaction with politicians—for example, imposing by stationarity, as in [Duggan \(2017\)](#) and [Banks and Sundaram \(1998\)](#). These strategies reduce the complexity of the analysis, but also reduce the benefit the voter can obtain from the repeated interaction.

The questions therefore remain: do there exist equilibria of the dynamic interaction between voters and politicians in which voters are able to use elections to optimally sanction *and* select politicians, in order to achieve at or near the first best?

In this paper, we obtain a positive and remarkable result: despite their shortcomings, if elections are sufficiently frequent, there need not be any trade-off between selection and control. We show that voters can closely approximate the first-best, i.e., the payoff that they would obtain were they perfectly informed about *both* politicians' actions *and* their characteristics: coarseness of the electoral mechanism does *not* impede accountability.

Our framework studies a long-run (infinite-horizon) relationship between a set of political candidates (or parties) and a representative voter. In each period, an election is held, and the voter derives a stage payoff from a policy outcome that depends on the winning politician's type (e.g., her ability), the action that she takes in office (e.g., effort), and unforeseen contingencies (a stochastic shock). The voter observes neither politicians' types nor their actions: her only sources of information are cheap-talk messages that candidates may send during each election campaign, and the policy outcome generated by the incumbent at each date. Politicians' types may vary over time: the best candidate, today, may not be the best candidate, tomorrow.

We start with a hypothetical static benchmark that equips the voter with the ability to commit to (i) a probability of election, conditioned on candidates' campaign messages, and (ii) monetary transfers. This allows us to ask how well the voter could do if she were endowed with a richer set of contracts than simply retaining or replacing the incumbent. A preliminary result—Observation

1—establishes that, indeed, the voter can implement the first-best outcome in this idealized environment. We then transition to an infinite horizon setting in which the voter’s only instrument for selection and control is her decision to reelect or replace the incumbent. Naturally, we do not allow the voter to make any binding commitments.

We adopt the classical mechanism-design approach to dynamic games with imperfect information by using a candidate’s future payoffs—her *continuation values*—to simulate the transfers that she would receive in our hypothetical static benchmark. The idea underlying this approach is straightforward: if the variation in a candidate’s future payoffs from today’s choices replicates how her transfers vary with her choices in the static interaction, her incentives to report her type and choose maximal effort in the static context with transfers can also be replicated in the dynamic environment *without* transfers. Our main result—Proposition 1—shows that this can be achieved.

Our result does not obtain from arguments related to the folk theorem on repeated games.¹ Also in contrast with that the linear programming approaches used in that literature, our approach explicitly describe the voter’s decision rule. This rule has a number of substantively interesting properties, some of which play an important role in other work on elections and accountability.

First, the rule is “retrospective” in a very strong sense: a voter’s decision depends not only on the current incumbent’s performance, but also on the past performances and campaigns of *all* candidates. Thus, the candidates’ previous choices affect their electoral fortunes even if they were not in office when these choices were made. We therefore join [Gordon, Huber and Landa \(2007\)](#) in emphasizing the importance of electoral challengers for outcomes even when they do not win.

Second, the voter chooses a lowest quality candidate at each election with an arbitrarily small probability, which increases as the past performance of the parties deteriorates, or when the voter believes that candidates’ announcements during previous campaigns were lies. Thus, it is part of an optimal strategy to elect on rare occasions a bad candidate when the voter feels that she has been repeatedly betrayed by the remaining candidates (e.g., [Di Tella and Rotemberg, 2018](#)).

Third, when the voter decides that a candidate’s campaign messages are no longer credible, she disregards that candidate in future elections. In essence, she “tunes out” and ignores that candidate in the future.

The paper is as organized as follows. [Section 2](#) presents our theoretical framework, [Section 3](#)

¹The *Road Map* in our Supplemental Appendix further elaborates on the non-application of existing techniques from repeated games to the present framework.

states our main result, and [Section 4](#) develops intuition for the voter’s decision rule, and why it works. [Section 5](#) discusses our contribution to the most closely related work. Formal statements and proofs of all of our results are in a Supplemental Appendix.

2. The Framework

Agents. We consider an interaction between a representative voter and a finite set $C \equiv \{1, \dots, n\}$ of candidates, and a single low-quality candidate, denoted by 0. The interaction unfolds over an infinite number of discrete periods $t = 1, 2, \dots$.

Adverse Selection. In each period $t \in \mathbb{N}$, each candidate i has a type θ_i^t . This type represents the candidate’s quality, or competence. The set of possible types of candidate $i \in C$, Θ_i , is a finite subset of \mathbb{R} . We allow each candidate’s type to vary across periods and to be serially correlated: the candidates’ types are distributed according to independent, irreducible Markov chains on $\Theta = \Theta_1 \times \dots \times \Theta_n$. Thus, in each period t , the probability distribution over candidate i ’s type may depend on her type in period $t - 1$.

Moral Hazard. In each period t , the elected candidate, say i , chooses an unobservable action a_i^t from a finite set $A \subset \mathbb{R}_+$ that includes 0. This action can be interpreted as the level of effort exerted by the policymaker in office. The office holder’s action yields an observed policy outcome $y^t = \theta_i^t + a_i^t + \varepsilon^t$, where ε^t is a stochastic shock that is distributed according to a well-behaved density f with $\mathbb{E}(\varepsilon) = 0$. The voter receives no information about the office-holder’s action beyond the realized policy outcome. We let a^* denote the maximal effort level, i.e., $a^* \equiv \max A$.

Timing. Events unfold as follows in each period t :

1. All the candidates privately learn their types.
2. The candidates simultaneously make public announcements from a set of messages M_i .
3. The voter elects a candidate from $\{0, 1, \dots, n\}$.
4. The elected candidate i selects an action a_i^t from A .
5. All agents observe the policy outcome $y^t = \theta_i^t + a_i^t + \varepsilon^t$.

We denote the voter’s (possibly random) decision by $\pi = (\pi_1, \dots, \pi_n) \in [0, 1]^n : \sum_{i=1}^n \pi_i \leq 1$, where π_i is candidate $i \in C$ ’s probability of winning the election. Thus, the low-quality candidate’s probability of winning is $1 - \sum_{i=1}^n \pi_i$.

Payoffs. The voter’s stage-game payoff is given by $v(\theta_i^t + a^t + \varepsilon^t)$ if she elects candidate $i \in \{0, 1, \dots, n\}$, where $v : \mathbb{R} \rightarrow \mathbb{R}$ is a continuous, strictly increasing function; the elected candidate’s stage-game payoff is given by $B - c(a^t)$, where $B > 0$ is an office-rent, and $c : \mathbb{R}_+ \rightarrow \mathbb{R}$ is a strictly convex cost function with $c(0) = 0 < c'(0)$; and all the candidates who lose the election at date t receive a payoff of zero at that date. Time is discounted at a common rate $\delta \in [0, 1)$. All agents seek to maximize their average discounted sums of stage utilities.

Equilibrium. We study perfect Bayesian equilibria, allowing agents to play mixed strategies.

Additional Assumptions. We impose some additional assumptions that ease analysis but which are not needed for our result. First, we assume that candidate 0 has a single type θ_0 , and that $\theta_0 < \min_i(\min \Theta_i)$. In other words, the quality of the candidates in C is always higher than that of the “bad” candidate 0. This assumption could be replaced by a weaker restriction that the probability with which this candidate is the very best candidate at any date is not too large. Second, we assume that the set of messages M_i is sufficiently rich, in the sense that $\Theta_i \times [0, 1]$ is embedded in M_i . This allows us to generate a public randomization device that makes deviations by the voter detectable. Third—and to avoid trivialities—we assume that the benefit of holding office is large enough to outweigh a candidate’s costs from choosing maximal effort, i.e., $B > c(a^*)$, and we assume that for every candidate i , there is a positive probability at each date of a profile θ such that candidate i is the most able candidate, i.e., $\theta_i = \max\{\theta_j : j = 1, \dots, n\}$.

A natural interpretation of our candidates is that they are political parties. Indeed, all of our references to ‘candidates’, below, could be replaced with ‘parties’, further justifying the idea that the candidates are long-lived.

3. Main Result

We begin by characterizing the voter’s optimal selection rule, and her first-best payoff. The optimal rule at each date is:

$$\pi_i^*(\theta_1, \dots, \theta_n) = \begin{cases} 1 & \text{if } \theta_i = \max\{\theta_j : j = 1, \dots, n\}, \\ 0 & \text{otherwise.} \end{cases}$$

for each candidate $i \in C$. Further, the voter prefers that the elected candidate exert maximal effort $a^* = \max A$. The electoral mechanism should therefore aim to give the representative voter a payoff

that is as close as possible to

$$u^* \equiv (1 - \delta)\mathbb{E}_0 \left[\sum_{t=1}^{\infty} \sum_{i=1}^n \delta^{t-1} \pi_i^*(\theta^t) \int_{-\varepsilon}^{\varepsilon} v(\theta_i^t + a^* + \varepsilon) f(\varepsilon) d\varepsilon \right]. \quad (1)$$

We state our main result.

Proposition 1. *For every $\varepsilon > 0$, there exists $\bar{\delta} < 1$ such that the following holds: for any $\delta \in (\bar{\delta}, 1)$, there exists a perfect Bayesian equilibrium in which the voter's payoff is within ε of u^* .*

Alternatively stated: if elections are sufficiently frequent, or, alternatively, if players are sufficiently patient, the voter can achieve arbitrarily close to full electoral accountability.² For the remainder of the paper, we provide intuition for the result, and highlight the features of substantive interest—the full proof is provided in a Supplemental Appendix.

Step 1: Benchmark with Transfers. As a contracting mechanism, the voter's electoral instrument is crude—she can only choose either to retain the incumbent, or replace her with a challenger. This raises the question: could the voter achieve electoral accountability even under a more favorable contracting environment?

The first step in our argument considers such a hypothetical setting, by allowing the voter to supplement her selection decision with a monetary transfer to each candidate. The following observation is formally developed and stated in Lemma 1 of the Supplemental Appendix.

Observation 1. *If the voter had access to transfers, she could design a reward scheme that, when coupled with the election rule π^* , implements the first-best outcome in dominant strategies.*

The precise formulation of Observation 1 is as follows: confronted with the election rule π^* and an appropriately constructed transfer functions $r_i(\hat{\theta}, y)$, truthfully reporting their types and (if elected) choosing maximal effort is a dominant strategy for all candidates. In particular, there exists $\beta > 0$ such that for every candidate i , for any type of candidate i , $\theta_i \in \Theta_i$ and report she might make $\hat{\theta}_i \in \Theta_i$, any action she might take $a \in A$, and set of reports by the remaining candidates $\hat{\theta}_{-i} \in \Theta_{-i}$, the following is true: if $\max\{\pi_i^*(\theta_i, \hat{\theta}_{-i}), \pi_i^*(\hat{\theta}_i, \hat{\theta}_{-i})\} = 1$,³ then for any $\hat{\theta}_i \neq \theta_i$ or $a \neq a^*$ (or both):

$$\pi_i^*(\theta_i, \hat{\theta}_{-i})[B - c(a^*)] + \int_{-\varepsilon}^{\varepsilon} r_i(\theta_i, \hat{\theta}_{-i}, \theta_i + a^* + \varepsilon) f(\varepsilon) d\varepsilon$$

²The interpretation of more frequent elections obtains by parameterizing $\delta \equiv \exp(-\rho\Delta)$, where $\Delta > 0$ is the length of each electoral cycle, and $\rho \in (0, 1)$, so that $\delta \rightarrow 1$ as $\Delta \rightarrow 0$.

³If a candidate is not elected at either message, her choice of message does not affect the voter's payoff.

$$> \pi_i^*(\hat{\theta}_i, \hat{\theta}_{-i})[B - c(a)] + \int_{-\varepsilon}^{\varepsilon} r_i(\hat{\theta}_i, \hat{\theta}_{-i}, \theta_i + a + \varepsilon) f(\varepsilon) d\varepsilon + \beta. \quad (2)$$

Why is it important that these incentives are strict? It implies that the voter can preserve candidates' incentives not only with the reward scheme $r_i(\hat{\theta}, y^t)$, but also by way of other reward schemes that closely—but imperfectly—approximate this scheme. This is crucial, because in real-world elections the voter does not have access to transfers—she can provide incentives only through variation in a candidate's prospect of election. So, the main step is to show that the voter can use carefully chosen perturbations of the optimal selection rule π^* to change candidates' future payoffs in ways that approximate the incentives arising from transfers.

Step 2: Voter's Election Rule. Under the optimal election rule π^* , some candidates are sometimes elected with probability zero, or with probability one. Our first step is to construct a benchmark perturbation of this rule under which all candidates face a prospect of election that lies strictly between zero and one. Let $\hat{\theta}$ denote a profile of messages sent by the candidates; for $\eta \in (0, 1/(2n-1))$, we construct this benchmark perturbation as follows:

$$\pi_i^\eta(\hat{\theta}) = \begin{cases} \pi_i^*(\hat{\theta}) - (2n-1)\eta & \text{if } \pi_i^*(\hat{\theta}) = 1, \\ \pi_i^*(\hat{\theta}) + \eta & \text{otherwise.} \end{cases} \quad (3)$$

Starting from the election rule $\pi_i^\eta(\hat{\theta})$, the voter can further adjust *each* of these n candidates' prospects of election—either as a reward or punishment—by an additional increment $\phi_i \in (-\eta, \eta)$.⁴ The key is that variation in these increments from different announcements and policy outcomes should approximate variation in the transfers.

To that end, we divide time into blocks of T elections, each. At each election in block b , the voter selects each candidate using the perturbed election rule (3) plus an increment $\phi_i^b \in (-\eta, \eta)$ that is constant within the block b , and determined at the end of the previous block, i.e., after the policy outcome at date $t = (b-1)T$. In order to determine ϕ_i , the voter applies two distinct tests: a *credibility* test, and an *eligibility* test. We first explain how these tests are implemented, informally, and subsequently describe their significance.

Credibility test. At each election *within* a block, the voter observes a profile of reports from every candidate, $\hat{\theta}^t$, and generates a profile of candidates' types $\vartheta^t = (\vartheta_1^t, \dots, \vartheta_n^t)$, that constitute her 'best

⁴Notice that under this benchmark decision rule, the low quality candidate is therefore elected with probability $1 - \sum_{i=1}^n \pi_i^\eta(\hat{\theta}) = n\eta$.

guess' about each candidate's type at that date.

To formulate this guess, the voter engages in the following introspection: for each candidate i , is the message that she sent in that election—together with her announcements in previous elections within that block—the messages that I would expect to receive if the candidate were truthful? Specifically: is the empirical transition of each candidate's reported types close to the underlying stochastic process by which her types are generated?

If the answer to this question is *yes*, the voter treats the candidate's announcement as credible, i.e., sets $\vartheta_i^t = \hat{\theta}_i^t$. If, however, the candidate fails this test, the voter ignores that candidate's message and all of her subsequent announcements in that block; instead, she replaces her subsequent messages in that block with simulated messages according to what the voter would have expected under truth-telling. In essence, the voter 'tunes out' of candidates whose messages are no longer credible.

Eligibility test. At the end of the b^{th} block, the voter uses both the profile ϑ^t and the policy outcome y^t at each date $t = (b - 1)T + 1, \dots, bT$ to determine a number:

$$\phi_i^b = \frac{1}{T} \sum_{t=(b-1)T+1}^{bT} r_i(\vartheta^t, y^t), \quad (4)$$

where we recall that $r_i(\vartheta^t, y^t)$ corresponds to the transfer a candidate received in the static benchmark. The voter then assigns each candidate i an "eligibility" status $e_i^{b+1} \in \{0, 1, \dots, \lambda T\}$, where $\lambda \in \mathbb{N}$. The eligibility status is determined by the following rule:

1. If a candidate's eligibility status at the start of the block b was $e_i^b = 0$, the candidate is assigned $e_i^{b+1} = 0$ if the following test is satisfied:

$$\left| \frac{\phi_i^b}{B - c(a^*)} \right| < \eta, \quad (5)$$

otherwise the candidate is assigned $e_i^{b+1} = 1$.

2. If $e_i^b \in \{1, \dots, \lambda T - 1\}$, then $e_i^{b+1} = e_i^b + 1$.
3. If $e_i^b = \lambda T$, then $e_i^{b+1} = 0$.

The voter then elects each candidate i at each period t in subsequent block $b + 1$ with probability

$$\pi_i^b(\vartheta^t) = \begin{cases} \pi_i^\eta(\vartheta^t) + \frac{\phi_i^b}{B-c(a^*)} & \text{if } e_i^b = 0 \\ 0 & \text{otherwise.} \end{cases} \quad (6)$$

A candidate's eligibility status at block b , e_i^b , is interpreted as the number of previous blocks in which the voter has given her a 'time out' by electing her with probability zero at every election within that block. This period of no-election will serve as a punishment that the voter imposes on any previously eligible candidate that fails the criterion (5).

In particular, the integer λT is the number of periods (and thus λ the number of blocks) for which a candidate must forego any prospect of election. If and only if a candidate has not yet been consigned to the punishment phase ($e_i^{b-1} = 0$) or if she has ended the punishment phase ($e_i^{b-1} = \lambda T$), she qualifies for election under the perturbed rule (6) if and only if (5) holds.

The condition defining the eligibility test represents the voter's "budget constraint" in how much variation she can assign to any candidate's prospect of reelection. Recall that, under the perturbed rule (3), the voter can further modify each candidate's prospect of election up or down by η . The eligibility test (5) constrains any further variation in the benchmark rule to respect this constraint.

4. Why the Rule Works

We now provide a broad intuition for why the voter's decision rule induces an equilibrium in which her payoff is close to u^* .

Candidate Incentives. We want to focus on a candidate's incentives in block b , setting aside the strategic behavior of the remaining candidates as well as the voter's credibility and eligibility tests in that and all subsequent blocks.

To that end, we consider the following thought experiment. At the start of block b , suppose that candidate i expects every other candidate to truthfully report her type and select maximal effort a^* at every election, *and* that the candidate i is herself pre-committed to truthful reporting and maximal effort in every election starting in the subsequent block $b + 1$. Suppose also that the candidate expects to be elected under the first-best rule π^* during elections in block b and at each election from block $b + 2$ onward, but anticipates being elected under the perturbed rule $\pi_i^b(\hat{\theta}^t)$ defined in

(6) at each date $t = bT + 1, \dots, (b + 1)T$ in the next block $b + 1$. Suppose, finally, that neither the credibility nor eligibility test will be used: the voter believes each candidate's announcement and every candidate remains eligible at the start of the subsequent block.

In this context, the candidate i 's problem is to choose a strategy for each election in block b . Observe that the candidate's continuation value at the start of block $b + 1$ —under the presumption that she remains eligible—is:

$$\begin{aligned}
& (1 - \delta) \sum_{t=bT+1}^{(b+1)T} \delta^{t-1} \pi_i^b(\hat{\theta}_i^t, \theta_{-i}^t) [B - c(a^*)] + \delta^T \times \text{constant} \\
&= (1 - \delta) \sum_{t=bT+1}^{(b+1)T} \delta^{t-1} \underbrace{[\pi_i^\eta(\hat{\theta}_i^t, \theta_{-i}^t) + \phi_i^b [B - c(a^*)]^{-1}]}_{\pi_i^b(\hat{\theta}_i^t, \theta_{-i}^t)} [B - c(a^*)] + \delta^T \times \text{constant} \\
&= (1 - \delta) \frac{\sum_{t=(b-1)T+1}^{bT} r_i(\hat{\theta}_i^t, \theta_{-i}^t, y^t)}{T[B - c(a^*)]} \sum_{t=bT+1}^{(b+1)T} \delta^{t-1} [B - c(a^*)] + \text{constant} \\
&= (1 - \delta) \frac{\sum_{t=(b-1)T+1}^{bT} r_i(\hat{\theta}_i^t, \theta_{-i}^t, y^t)}{T} \frac{1 - \delta^T}{1 - \delta} + \text{constant}. \tag{7}
\end{aligned}$$

Since $\lim_{\delta \rightarrow 1} (1 - \delta^T)/(1 - \delta) = T$, it follows that for $\delta \approx 1$, the candidate's objective function at the start of block b is approximately

$$\begin{aligned}
& (1 - \delta) \sum_{t=(b-1)T+1}^{bT} \delta^{t-1} \pi_i^*(\hat{\theta}_i^t, \theta_{-i}^t) [B - c(a_i^t)] + (1 - \delta) \sum_{t=(b-1)T+1}^{bT} \int_{-\varepsilon}^{\varepsilon} r_i(\hat{\theta}_i^t, \theta_{-i}^t, y^t) f(\varepsilon) d\varepsilon \\
&\approx (1 - \delta) \sum_{t=(b-1)T+1}^{bT} \left[\pi_i^*(\hat{\theta}_i^t, \theta_{-i}^t) [B - c(a_i^t)] + \int_{-\varepsilon}^{\varepsilon} r_i(\hat{\theta}_i^t, \theta_{-i}^t, y^t) f(\varepsilon) d\varepsilon \right]. \tag{8}
\end{aligned}$$

Comparing this objective function with the inequality (2), we observe that if the candidate's objective were *exactly* given by (8), we could invoke Observation 1 that truthfully reporting her type and choosing the optimal action at every date t is optimal for candidate i .

Recall, however, that the continuation values only *approximate* the incentives under transfers, not perfectly recreate them. Thus, we do not explicitly characterize candidate i 's strategy. Rather, our proof shows that an arbitrarily good approximation in the dynamic game of candidate i 's incentives in the context of transfers leads her to choose strategies whose payoff consequences to the voter are arbitrarily close to the first-best, u^* . Moreover, making η arbitrarily small implies that the

voter's required distortions of the optimal rule π^* are also arbitrarily small.

Our thought experiment presumed that candidate i will indeed remain eligible in the subsequent block regardless of her opponents' choices. But eligibility in subsequent block $b + 1$ depends on the increment ϕ_i^{b+1} , which in turn depends on the announcements and actions of *all* the candidates in the previous block b . Since candidate i cannot perfectly anticipate what her opponents will report at each election within the block, she may fear falling afoul of the eligibility test simply because of her opponent's choices. But if she is indeed correct, her subsequent ineligibility means that she foregoes the continuation values that she is supposed to realize in the next block, and our representation of her objective in expression (8) would not apply. We therefore require two critical steps.

First: if a candidate indeed reports truthfully and chooses maximal effort a^* whenever she is elected—or does not deviate 'too frequently' from this strategy—she will remain eligible with arbitrarily high probability, regardless of her opponents' choices.

Second: ineligibility is sufficiently punitive to deter a candidate from any strategy that renders this prospect non-trivial.

The first step is obtained by way of the 'credibility test' to which the voter subjects the candidate at each date. On the one hand, if the voter is listening to a candidate j 's announcements, the empirical distribution of her reported types must be close to the underlying distribution. On the other hand, if the voter has stopped listening to the candidate's announcements, then the voter likewise simulates her messages from the underlying distribution. In either case, we have that $\vartheta_j^t = \hat{\theta}_j^t$ with arbitrarily high probability.

The credibility test effectively makes candidate i 's announcement strategy a decision problem, i.e., a choice that does not depend on her conjectures about the actions of her opponents. The first step then concludes by observing that—under the test—a candidate's prospect of falling afoul of the condition (5) and thus becoming ineligible if she tells the truth and chooses maximal effort at every election in the block approaches zero as the number of periods T in the block grows large.

In short: the credibility test ensures that regardless of her opponent's choices, a candidate's conviction that she will remain eligible and therefore realize the continuation payoffs in the next block if she abides by truth-telling and maximal effort in the present block is indeed correct.

The second step—i.e., that if a candidate were to choose some other strategy that led to a non-trivial prospect of ineligibility, the punishment phase of λT periods of certain non-election is sufficiently punitive—is straightforward: we simply choose the number of blocks λ to be arbitrarily

large.

Voter Incentives. If a candidate becomes ineligible, the voter’s decision rule calls on her to subject the candidate to λT periods of certain non-election. This is a severe punishment to the candidate. But, it may also severely harm the voter! To see why, recall that at some period of non-election, a candidate i ’s type may be higher than every other candidate’s, i.e., $\theta_i = \max\{\theta_j : j = 1, \dots, n\}$. The first-best rule π^* calls on the voter to select this candidate with probability one. However, the candidate’s ineligibility calls on the voter *not* to elect this candidate. This highlights the classic trade-off between selection and control: in a single-period election, the voter would not wish to implement the punishment.

Our dynamic elections framework therefore raises two questions. First, how can the voter threaten a candidate with non-eligibility if, at one of the subsequent λT elections, the candidate turns out to be the highest quality? Second, to what extent do the voter’s costs from implementing the punishment strategy impede her ability to obtain the first-best payoff u^* ?

To address the first concern, we specify that the following ‘babbling equilibrium’ is played at every subsequent election whenever the voter deviates from her election rule: all candidates send uninformative messages to the voter, who elects the best candidate under her prior belief; and, whichever candidate is elected chooses the lowest level of effort $a = 0$. Since the voter’s payoff from this equilibrium is strictly less than the first-best u^* , the threat of a transition to this equilibrium encourages her to apply the decision rule at each election. This makes her implicit threat to punish a candidate with ineligibility credible.

To address the second concern, we show that under the credibility test, so long as a candidate indeed reports her type truthfully and chooses maximal effort—or does not deviate too far from these choices—the prospect that she becomes ineligible is arbitrarily small. This, in turn, implies that the possibility that the voter is forced to punish a candidate is also arbitrarily small. This makes the payoff consequences of the trade-off between selection and control that arises whenever the punishment is triggered arbitrarily small.

5. Related Work.

Our infinite horizon framework includes uncertainty about politicians’ actions (*moral hazard*) and their types (*adverse selection*), a long-run relationship between voters and politicians, and equilibria with no bounds on the history dependence of agents’ strategies. These distinguish us from

existing work that focuses on complete information (e.g., [Van Weelden, 2013](#)), or in which policy choices are observed (e.g., [Aragonès, Postlewaite and Palfrey, 2007](#), [Meirowitz, 2007](#)), or in which there is no uncertainty about politicians' types at the time policies are chosen (e.g., [Austen-Smith and Banks, 1989](#), [Aragonès, Postlewaite and Palfrey, 2007](#), [Duggan and Forand, 2014](#)). The former also includes [Kalandrakis \(2009\)](#), who considers a setting in which party preferences are private information and exhibit serial correlation.

Other frameworks restrict the history dependence of agents' strategies explicitly, or implicitly via short-lived voters (e.g., [Kartik and Van Weelden, 2015](#)), two-period term limits ([Banks and Sundaram, 1998](#), [Duggan, 2017](#)), or a two-period framework ([Duggan and Martinelli, 2015](#)). As Duggan and Martinelli note, in the context of models feature adverse selection and moral hazard: “[d]ue to difficult theoretical issues, the literature going beyond two periods is small” (p. 26). A notable exception is [Banks and Sundaram \(1993\)](#). However, the authors assume that differences between candidates relate to their marginal costs of effort—types themselves have no direct payoff consequence to voters. This ensures that a voter is indifferent between the set of candidates who choose zero effort at any election, which is crucial for their trigger equilibrium. Our production technology, by contrast, is the same as [Holmström \(1999\)](#), studied in a three-period political economy setting by [Ashworth \(2005\)](#). In those frameworks, however, learning is symmetric across politicians and voters; our framework assumes that a politician's type is private information.

References

- Aragonès, Enriqueta, Andrew Postlewaite and Thomas Palfrey. 2007. “Political reputations and campaign promises.” *Journal of the European Economic Association* 5(4):846–884.
- Ashworth, Scott. 2005. “Reputational dynamics and political careers.” *Journal of Law, Economics, and Organization* 21(2):441–466.
- Ashworth, Scott, Ethan Bueno de Mesquita and Amanda Friedenberg. 2017. “Accountability and information in elections.” *American Economic Journal: Microeconomics* 9(2):95–138.
- Austen-Smith, David and Jeff Banks. 1989. Electoral accountability and incumbency. In *Models of Strategic Choice in Politics*. University of Michigan Press.
- Banks, Jeffrey S and Rangarajan K Sundaram. 1993. Adverse selection and moral hazard in

- a repeated elections model. In *Political Economy: Institutions, Information, Competition and Representation.*, ed. William Barnett et al. Vol. 12 Cambridge University Press pp. 295–311.
- Banks, Jeffrey S and Rangarajan K Sundaram. 1998. "Optimal retention in agency problems." *Journal of Economic Theory* 82(2):293–323.
- Besley, Timothy. 2006. *Principled agents?: The political economy of good government*. Oxford University Press on Demand.
- Di Tella, Rafael and Julio J Rotemberg. 2018. "Populism and the return of the "Paranoid Style": Some evidence and a simple model of demand for incompetence as insurance against elite betrayal." *Journal of Comparative Economics* 46(4):988–1005.
- Duggan, John. 2017. "Term limits and bounds on policy responsiveness in dynamic elections." *Journal of Economic Theory* 170:426–463.
- Duggan, John and César Martinelli. 2015. "Electoral accountability and responsive democracy.".
- Duggan, John and César Martinelli. 2017. "The political economy of dynamic elections: Accountability, commitment, and responsiveness." *Journal of Economic Literature* 55(3):916–84.
- Duggan, John and Jean Guillaume Forand. 2014. "Markovian Elections.".
- Fearon, James D. 1999. "Electoral accountability and the control of politicians: selecting good types versus sanctioning poor performance." *Democracy, accountability, and representation* 55:61.
- Ferejohn, John. 1986. "Incumbent performance and electoral control." *Public choice* 50(1-3):5–25.
- Gordon, Sanford C, Gregory A Huber and Dimitri Landa. 2007. "Challenger entry and voter learning." *American Political Science Review* 101(2):303–320.
- Holmström, Bengt. 1999. "Managerial incentive problems: A dynamic perspective." *The review of Economic studies* 66(1):169–182.
- Kalandrakis, Tasos. 2009. "A reputational theory of two-party competition." *Quarterly Journal of Political Science* 4(4):343–378.
- Kartik, Navin and Richard Van Weelden. 2015. "Reputation, Term Limits, and Incumbency (Dis) Advantage.".

- Meirowitz, Adam. 2007. "Probabilistic voting and accountability in elections with uncertain policy constraints." *Journal of Public Economic Theory* 9(1):41–68.
- Van Weelden, Richard. 2013. "Candidates, credibility, and re-election incentives." *Review of Economic Studies* 80(4):1622–1651.

Supplementary Appendix to *Making Elections Work*

Road map. To prove Proposition 1, it suffices to establish that, for sufficiently high δ , there exists an equilibrium in which the following happens arbitrarily often: (i) the voter uses a decision rule arbitrarily close to $\pi^*(\cdot)$; (ii) the candidates truthfully report their types; and (iii) the winning candidate belongs to $C \equiv \{1, \dots, n\}$ and chooses the action a^* .

We begin with two preliminary observations. Firstly, we will focus on equilibria in which candidate 0's probability of being elected only depends on the other candidates' past reports and outputs. It follows that it is optimal for this candidate to choose the lowest action $a = 0$ whenever she is elected. (As this will only occur with an arbitrarily small probability, it will only have a negligible impact on the voter's expected payoff.)

Secondly, it is readily checked that repeating the following one-period "babbling equilibrium" in every period gives a PBE for the infinite-horizon model: all candidates send uninformative messages to the voter, who always elects the best candidate according to her prior beliefs; and the elected candidate always chooses the costless action $a = 0$. The voter's expected payoff in this equilibrium is strictly lower than u^* . Therefore, for any candidate equilibrium, attaching the punishment of perpetual play of the babbling equilibrium to out-of-equilibrium voting behavior suffices to ensure (for sufficiently large δ) that such behavior is suboptimal. Although the voter may use mixed strategies, the candidates' messages can be used to generate a public correlating device (Aumann, Maschler and Stearns, 1995) allowing them to detect voter deviations.

Put together, the above two observations obviate the need to consider deviations by candidate 0 and the voter. To obtain an equilibrium giving an expected payoff of u^* to the voter, it suffices to construct a dynamic decision rule (or mechanism) that creates the appropriate incentives for the candidates in C . Under this dynamic rule, equilibrium choices in future periods serve a transfer role, rewarding the candidates for their past actions and revelations of their types. As explained in the main text, the first step of the proof establishes existence of a transfer function which, coupled with the decision rule π^* , would induce candidates to truthfully reveal their types and choose action a^* (when elected) in a one-period model. In the second step, we construct a dynamic decision rule to simulate this transfer rule in our dynamic model, using the candidates' continuation values. Finally, step 3 verifies that the game induced by the dynamic decision rule possesses a PBE in which the voter's payoff is arbitrarily close to u^* .

Before proceeding, it is worth noting that the mechanism-design approach to dynamic games

with asymmetric information in which continuation payoffs play the role of transfers is quite standard and has been used extensively (e.g., Mailath and Samuelson 2006, Chapters 8 and 11). As we will see below, however, the way transfers are simulated here is different. Among the distinctive features that prevent us from directly applying existing results to our model are the following: (i) it exhibits both adverse selection and moral hazard; (ii) it has common values, since the voter's utility depends on the winning candidate's type; (iii) types are serially correlated; and (iv) in each period, the public signal (output) is not only generated by the officeholder's hidden action but also by her type. An added attraction of our approach is that it permits us to characterize not only the agents' payoffs but also their behavior in equilibrium.

Step 1: Implementing π^* with a transfer rule. Applying techniques from the principal-agent literature (e.g., Laffont and Tirole 1993), we can construct a transfer function — i.e. a mapping from the observed candidates' reports and outputs to a vector of monetary transfers — which implements the voter's optimal decision rule π^* . To do so, we first define the function $a: \mathbb{R} \rightarrow \mathbb{R}_+$ as any C^1 function that satisfies the following conditions: $a(x) = a^*$ and $a'(x) > 0$, for all $x \in \cup_{i \in C} \Theta_i$. Then, for each $i \in C$, let $\rho_i: \Theta \times \mathbb{R} \rightarrow \mathbb{R}$ be defined by

$$\rho_i(\hat{\theta}, y) \equiv \pi_i^*(\hat{\theta}) \left[c'(a(\hat{\theta}_i)) [y - \hat{\theta}_i - a(\hat{\theta}_i)] + c(a(\hat{\theta}_i)) + \int_{m_i(\hat{\theta})}^{\hat{\theta}_i} c'(a(z)) dz - B \right]$$

where $m_i(\hat{\theta}) \equiv \max_{j \neq i} \hat{\theta}_j$; and let $r_i: \Theta \times \mathbb{R} \rightarrow \mathbb{R}$ be defined by

$$r_i(\hat{\theta}, y) \equiv \rho_i(\hat{\theta}, y) - \mathbb{E}_{\theta, \varepsilon} [\rho_i(\theta, \theta + a(\theta) + \varepsilon)].$$

Here and in what follows, the expectations over type profiles θ are always taken using the invariant distribution p of the Markov chain governing the evolution of types.

Our first result establishes that the voter's optimal decision rule, coupled with the transfer function $r = (r_1, \dots, r_n)$ is strategy-proof, in the sense that under these rules, it is optimal for all candidates to truthfully reveal their types when campaigning and choose action a^* when in office.

Lemma 1. *There exists $\beta > 0$ such that the following holds for all $i \in C$, $\theta_i, \hat{\theta}_i \in \Theta_i$ and $\hat{\theta}_{-i} \in \Theta_{-i}$: $\max \{ \pi_i^*(\theta_i, \hat{\theta}_{-i}), \pi_i^*(\hat{\theta}_i, \hat{\theta}_{-i}) \} = 1$ implies that if either $\hat{\theta}_i \neq \theta_i$ or $a \neq a^*$ (or both), then*

$$\pi_i^*(\theta_i, \hat{\theta}_{-i}) [B - c(a^*)] + R_i^*(\theta_i, \hat{\theta}_{-i} | \theta_i) > \pi_i^*(\hat{\theta}_i, \hat{\theta}_{-i}) [B - c(a)] + R_i(\hat{\theta}_i, \hat{\theta}_{-i}, a | \theta_i) + \beta,$$

where $R_i(\hat{\theta}, a \mid \theta_i) \equiv \mathbb{E}_\varepsilon[r_i(\hat{\theta}, \theta_i + a + \varepsilon)]$ and $R_i^*(\hat{\theta} \mid \theta_i) \equiv R_i(\hat{\theta}, a(\hat{\theta}_i) \mid \theta_i)$.

Before proving Lemma 1, two observations are in order. First it follows from the lemma that, given the transfer rule r , every candidate's cost from choosing an action $a \neq a^*$ or misreporting her type (whenever this affects the outcome of the election) is bounded away from zero. Therefore, the candidates' incentive constraints still hold strictly for sufficiently small perturbations of the rules π^* and r . Second, the lemma does not rule out deviations from truth-telling in cases where a candidate loses the election under π^* both before and after the deviation. However, such deviations do not affect electoral outcomes under π^* and, even if one slightly perturbs π^* , they can only have an arbitrarily small impact on the voter's payoff.

Proof of Lemma 1. Fix $i \in C$ and $\hat{\theta}_{-i} \in \Theta_{-i}$. To prove the lemma, we must show that, given the decision rule π^* and the transfer rule r_i , candidate i is always strictly better off truthfully reporting her true type θ_i and choosing the optimal action a^* than deviating to some untruthful report (that would change the outcome of the election) and/or to some action $a \neq a^*$. As the candidates' type and action sets are finite, this ensures that there exists a positive number $\beta_i > 0$ such that any such deviation would reduce i 's payoff by at least β_i . Setting $\beta \equiv \max_{i \in C} \beta_i$, we thus obtain the lemma.

We begin by showing that whenever it wins the election after making some report $\hat{\theta}_i \in \Theta_i$, candidate i is strictly better off choosing action $a(\hat{\theta}_i) = a^*$. Given the transfer rule $r_i(\cdot)$, an optimal action for i must be a solution to

$$\max_{a \in A} R_i(\hat{\theta}_i, \hat{\theta}_{-i}, a \mid \theta_i) - c(a),$$

or, equivalently,

$$\max_{a \in A} c'(a(\hat{\theta}_i))(\theta_i + a) - c(a).$$

Assuming that the feasible set is \mathbb{R}_+ (instead of A), it follows from the first-order condition (and the strict concavity of the objective function) of this problem that its unique solution is $a(\hat{\theta}_i)$. As $a(\hat{\theta}_i) = a^* \in A$, we conclude that it is also the unique solution to the above maximization problem.

We now turn to candidate i 's optimal report. We must consider two cases: first, where candidate i wins the election by truthfully reporting her type θ_i , and second, where she loses the election by truthfully reporting her type. We begin with the first case, i.e., where $\pi_i^*(\theta_i, \hat{\theta}_{-i}) = 1$. Anticipating that she will choose the optimal action if elected, candidate i 's expected payoff from making any report $\hat{\theta}_i > m_i(\hat{\theta})$ is given by $B - c(a(\hat{\theta}_i)) + R_i(\hat{\theta}_i, \hat{\theta}_{-i} \mid \theta_i)$. Maximizing this function with respect to

$\hat{\theta}_i$ in \mathbb{R} boils down to solving

$$\max_{\hat{\theta}_i \in \mathbb{R}} c'(\mathbf{a}(\hat{\theta}_i))(\theta_i - \hat{\theta}_i) + \int_{m_i(\hat{\theta})}^{\hat{\theta}_i} c'(\mathbf{a}(z))dz.$$

The only solution to the first and second-order conditions in Θ_i is $\hat{\theta}_i = \theta_i$. Moreover, making a report $\hat{\theta}_i < m_i(\hat{\theta})$ (so that $\pi_i^*(\hat{\theta}_i, \hat{\theta}_{-i}) = 0$) gives a payoff of $-\mathbb{E}_{\theta, \varepsilon}[\rho_i(\theta, \theta + a^* + \varepsilon)]$, whereas reporting her true type θ_i yields

$$\begin{aligned} [B - c(\mathbf{a}(\theta_i))] + R_i(\theta_i, \hat{\theta}_{-i} \mid \theta_i) &= \int_{m_i(\theta_i, \hat{\theta}_{-i})}^{\theta_i} c'(\mathbf{a}(z))dz - \mathbb{E}_{\theta, \varepsilon}[\rho_i(\theta, \theta + a^* + \varepsilon)] \\ &> -\mathbb{E}_{\theta, \varepsilon}[\rho_i(\theta, \theta + a^* + \varepsilon)] \end{aligned}$$

(recall that $c' > 0$). Hence, candidate i is strictly better off truthfully reporting her type.

We now turn to the second case where $\pi_i^*(\theta_i, \hat{\theta}_{-i}) = 0$, so that $\theta_i < m_i(\theta_i, \hat{\theta}_{-i})$. Consider a deviation from θ_i to any report $\hat{\theta}_i$ such that $\pi_i^*(\hat{\theta}_i, \hat{\theta}_{-i}) = 1$ (so that $\hat{\theta}_i > m_i(\hat{\theta}) > \theta_i$). The benefit of such a deviation is equal to

$$\begin{aligned} c'(\mathbf{a}(\hat{\theta}_i))(\theta_i - \hat{\theta}_i) + \int_{m_i(\hat{\theta})}^{\hat{\theta}_i} c'(\mathbf{a}(z))dz &\leq c'(\mathbf{a}(\hat{\theta}_i))(\theta_i - \hat{\theta}_i) + [\hat{\theta}_i - m_i(\hat{\theta})]c'(\mathbf{a}(\hat{\theta}_i)) \\ &= -c'(\mathbf{a}(\hat{\theta}_i))[m_i(\hat{\theta}) - \theta_i] < 0, \end{aligned}$$

where the inequality follows from the fact that $c'(\mathbf{a}(\cdot))$ is a strictly increasing function. This completes the proof of Lemma 1. \square

Step 2: The voter's dynamic decision rule. It follows from the previous step that the voter's ideal outcome can be supported as an equilibrium outcome if one can appropriately specify continuation play in the dynamic model, so that a deviation by a candidate adversely changes her continuation value. Our next goal is to show that this can be (approximately) achieved if the voter uses an appropriate history-dependent decision rule.

Fix $\eta \in (0, 1/(2n - 1))$, and let the decision rule π^η be defined by

$$\pi_i^\eta(\hat{\theta}) \equiv \begin{cases} \pi_i^*(\hat{\theta}) - (2n - 1)\eta & \text{if } \pi_i^*(\hat{\theta}) = 1, \\ \pi_i^*(\hat{\theta}) + \eta & \text{otherwise,} \end{cases}$$

for each $i \in C$. Observe that, under this decision rule, candidate 0 is always elected with probability $n\eta$. This allows the voter to reward/punish each candidate $i \in C$ in any period by increasing/decreasing her probability of being elected by some $\varphi_i \in (-\eta, \eta)$.

We define the decision rule as follows. It divides time into blocks of T periods. Each block $b = 1, 2, \dots$ begins with an “eligibility vector” $e^b = (e_1^b, \dots, e_n^b) \in \{0, 1, \dots, \lambda T\}^n$, $\lambda \in \mathbb{N}$, and a vector of promised “rewards” $\varphi^b = (\varphi_1^b, \dots, \varphi_n^b) \in [-\eta, \eta]^n$ inherited from play in the previous block; we set $e^1 \equiv (0, \dots, 0)$ and $\varphi^1 \equiv (0, \dots, 0)$. Then, in each period $t = (b-1)T + 1, \dots, bT$, events unfold as follows:

- (i) The candidates simultaneously report their types.
- (ii) Each candidate i 's report, $\hat{\theta}_i^t$, is submitted to Escobar and Toikka's (2013) statistical test to assess its credibility, and a profile of types $(\vartheta_1^t, \dots, \vartheta_n^t)$ is generated as follows: if all of i 's previous reports (in block b) passed the test, then $\vartheta_i^t = \hat{\theta}_i^t$; otherwise, ϑ_i^t is obtained from the theoretical distribution of i 's types.
- (iii) The voter selects each candidate $i \in C$ with probability $\pi_i^b(\vartheta^t) \equiv \pi_i^\eta(\vartheta^t) + \varphi_i^b$ if $e_i^b = 0$ (i.e., if candidate i is “eligible”), and with probability 0 otherwise.
- (iv) The elected candidate, say j , chooses an action a^t .
- (v) The outcome y^t is observed by all agents.

At the end of the b th block, each candidate i 's eligibility and reward for the next block are determined as follows. First, we calculate

$$\phi_i^b \equiv \frac{1}{T} \sum_{t=(b-1)T+1}^{bT} r_i(\vartheta^t, y^t);$$

then

$$e_i^{b+1} \equiv \begin{cases} 0 & \text{if either } e_i^b = \lambda T, \text{ or } e_i^b = 0 \text{ and } \left| [B - c(a^*)]^{-1} \phi_i^b \right| < \eta, \\ e_i^b + 1 & \text{otherwise;} \end{cases}$$

and then

$$\varphi_i^{b+1} \equiv \begin{cases} [B - c(a^*)]^{-1} \phi_i^b & \text{if } e_i^b = e_i^{b+1} = 0, \\ 0 & \text{otherwise.} \end{cases}$$

Step 3: The equilibrium. The dynamic decision rule defined above induces a dynamic game between the n candidates in C , which we denote by $\Gamma(\eta, \lambda, T, \delta)$. To complete the proof of the theorem, it remains to establish that for appropriately chosen parameters η, λ, T and δ , this game has a PBE in which the voter's expected payoff is approximately u^* .

For each $i \in C$, let σ_i^* be the strategy that prescribes candidate i to truthfully reveal her type when campaigning and choose action a^* when in office, regardless of the previous history of play. It is readily checked that, for every strategy profile σ (and every T), the average discounted payoff to i in each block $b \in \mathbb{N}$ can be approximated by

$$\mathbb{E}_\sigma \left[\mathbf{1}_{\{e_i^b=0\}} \frac{1}{T} \sum_{t=1}^T [\pi_i^\eta(\vartheta^{(b-1)T+t}) + \varphi_i^b] [B - c(a_i^{(b-1)T+t})] \right]$$

by letting $\delta \rightarrow 1$. Now suppose that candidate i plays σ_i^* and, therefore, chooses action a^* , whenever she is elected in block $b-1$. It follows from Lemma 5.1 in Escobar and Toikka (2013) that the empirical distribution of the sequence of report profiles generated by the test in each block converges to the invariant distribution p as $T \rightarrow \infty$, irrespective of the candidates' actual reports. By definition of r_i , the expected values of ϕ_i^{b-1} and φ_i^b (conditional on $e_i^{b-1} = 0$) approach zero if we let $T \rightarrow \infty$. This has the following two implications for i 's payoff. First, i 's expected block- b payoff, conditional on being eligible in block b , is approximately $\mathbb{E}_\theta [\pi_i^\eta(\theta) [B - c(a^*)]]$, regardless of the other candidates' strategies. Second, i 's probability of remaining eligible approaches one as we let $T \rightarrow \infty$; so that for every λ , we can choose T large enough that i is eligible arbitrarily frequently. It follows that candidate i can guarantee herself a payoff arbitrarily close $\mathbb{E}_\theta [\pi_i^\eta(\theta) [B - c(a^*)]]$ by playing σ_i^* ; and if each candidate $i \in C$ plays σ_i^* , then the voter receives an expected payoff arbitrarily close to $\mathbb{E}_\theta [\sum_{i=1}^n \pi_i^\eta(\theta) u(\theta_i + a^*)]$. This is formally stated in the next lemma.

Lemma 2. *For all $\eta > 0$ and $\lambda \in \mathbb{N}$, the following holds for game $\Gamma(\eta, \lambda, T, \delta)$ if we let $T \rightarrow \infty$ and then $\delta \rightarrow 1$:*

(i) *for each $i \in C$, the payoff to candidate i from playing σ_i^* approaches $\mathbb{E}_\theta [\pi_i^\eta(\theta) [B - c(a^*)]]$, irrespective of the other candidates' strategies; and*

(ii) *if each candidate $i \in C$ plays σ_i^* , then the voter's expected payoff approaches*

$$\mathbb{E}_\theta \left[\sum_{i=1}^n \pi_i^\eta(\theta) u(\theta_i + a^*) \right].$$

Proof of Lemma 2. Fix $\eta > 0$ and $\lambda \in \mathbb{N}$. For any strategy profile σ , the payoff to candidate $i \in C$ can be written as

$$U_i(\sigma) = (1 - \delta^T) \mathbb{E}_\sigma \left[\sum_{b=1}^{\infty} \delta^{(b-1)T} \mathbf{u}_i^b(\sigma) \right],$$

where

$$\mathbf{u}_i^b(\sigma) \equiv \frac{1 - \delta}{1 - \delta^T} \mathbf{1}_{\{e_i^b = 0\}} \sum_{t=1}^T \delta^{t-1} [\pi_i^\eta (\vartheta^{(b-1)T+t}) + \varphi_i^b] [B - c(a_i^{(b-1)T+t})].$$

Let \tilde{p}^T denote the empirical distribution of types profile in a block of T periods. For any $\gamma > 0$, the following holds for sufficiently large $T \in \mathbb{N}$:

$$\begin{aligned} \Pr_{\sigma_i^*, \sigma_{-i}} [e_i^b = 0 \mid e_i^{b-1} = 0] &= \Pr_{\sigma_i^*, \sigma_{-i}} \left\{ \left| [B - c(a^*)]^{-1} \phi_i^{b-1} \right| < \eta \right\} \\ &= \Pr_{\sigma_i^*, \sigma_{-i}} \left\{ \left| \frac{\mathbb{E}_{\tilde{p}^T} [r_i(\theta, \theta_i + a^* + \varepsilon)] - \mathbb{E}_\theta [r_i(\theta, \theta_i + a^* + \varepsilon)]}{[B - c(a^*)]} \right| < \eta \right\} \\ &> 1 - \gamma, \end{aligned}$$

where the second equality follows from the observation that $\mathbb{E}_\theta [r_i(\theta, \theta_i + a^* + \varepsilon)] = 0$, and the inequality from Escobar and Toikka's (2013) Lemma 5.1. By the same logic, for any $\gamma > 0$ (independent of b), we have $|\mathbb{E}_{\sigma_i^*, \sigma_{-i}} [\varphi_i^b \mid e_i^{b-1} = 0]| < \gamma$ if T is sufficiently large. Finally, it is well-known that for any $\gamma > 0$, we have

$$\sup \left\{ \left| \frac{1}{T} \sum_{t=1}^T v^t - \frac{1 - \delta}{1 - \delta^T} \sum_{t=1}^T \delta^{t-1} v^t \right| : (v^1, \dots, v^T) \in [0, B]^T \right\} < \gamma$$

if we let $T \rightarrow \infty$ and $\delta \rightarrow 1$.

Now for each $i \in C$, let $v_i^\eta \equiv \mathbb{E}_\theta [\pi_i^\eta(\theta) [B - c(a^*)]]$. Together with Escobar and Toikka's (2013) Lemma 5.1, the above inequalities imply that for all $\gamma \in (0, v_i^\eta)$, taking T sufficiently large and then δ sufficiently close to one, the following holds for every σ_{-i} :

$$\begin{aligned} \mathbb{E}[\mathbf{u}_i^b(\sigma_i^*, \sigma_{-i}) \mid e_i^b = 0] &> \mathbb{E}_\theta [\pi_i^\eta(\theta) [B - c(a^*)]] - \gamma = v_i^\eta - \gamma \\ &> 0 = \mathbb{E}[\mathbf{u}_i^b(\sigma_i^*, \sigma_{-i}) \mid e_i^b > 0]. \end{aligned}$$

Consequently, for all σ_{-i} , $U_i(\sigma_i^*, \sigma_{-i})$ is bounded below by the (fictitious) payoff, denoted \underline{V}_i , which candidate i would obtain if it received $v_i^\eta - \gamma$ in every block where she is eligible, 0 in the other

blocks, and her probability of becoming ineligible was held constant at γ across blocks. As i (like all the candidates) is eligible at the start of the initial block, we have

$$\underline{V}_i = (1 - \delta^T)(v_i^\eta - \gamma) + \delta^T(1 - \gamma + \gamma\delta^{\lambda T})\underline{V}_i$$

or, equivalently,

$$\underline{V}_i = \frac{1 - \delta^T}{1 - \delta^T(1 - \gamma + \gamma\delta^{\lambda T})}(v_i^\eta - \gamma).$$

An application of l'Hôpital's rule shows that

$$\lim_{\delta \rightarrow 1} \underline{V}_i = \frac{1}{1 + \gamma\lambda}(v_i^\eta - \gamma). \quad (\text{A1})$$

By the same logic as above, for every $\gamma > 0$, if we let $T \rightarrow \infty$ and then $\delta \rightarrow 1$, then $\mathbb{E}[\mathbf{u}_i^b(\sigma_i^*, \sigma_{-i}) | e_i^b = 0] < v_i^\eta + \gamma$, for all σ_{-i} . As $v_i^\eta + \gamma > 0$, $U_i(\sigma_i^*, \sigma_{-i})$ is bounded above by the payoff candidate i would obtain if she got $v_i^\eta + \gamma$ in every block. Coupled with (A1), this implies that irrespective of σ_{-i} , $U_i(\sigma_i^*, \sigma_{-i})$ approaches v_i^η as $\gamma \rightarrow 0$ and, therefore, as $T \rightarrow \infty$ and $\delta \rightarrow 1$. This proves the first part of Lemma 2. The proof of part (ii) is analogous. \square

Let $\eta \approx 0$. It follows from the second part of Lemma 2 that, for any $\varepsilon > 0$ and sufficiently large T and δ , the voter's expected payoff is within ε of u^* if every candidate i plays σ_i^* . Therefore, our final task is to show that $\Gamma(\eta, \lambda, T, \delta)$ has a PBE in which every candidate i 's strategy prescribes the same choices as σ_i^* arbitrarily frequently in every block. Our next result draws on Lemma 2(i) to show that this is the case for sufficiently large λ, T and δ .

Lemma 3. *For any $\eta \in (0, 1)$, if we let $\eta \rightarrow 0$, then $\lambda \rightarrow \infty$, then $T \rightarrow \infty$, and finally $\delta \rightarrow 1$, there exists a PBE for $\Gamma(\eta, \lambda, T, \delta)$ in which the expected empirical frequency of deviations from behavior prescribed by $(\sigma_1^*, \dots, \sigma_n^*)$ is less than η .*

We provide a complete proof of Lemma 3 below. Before proceeding, we outline the main steps of the argument. First, one can show that if we let $\eta \rightarrow 0$, then $\lambda \rightarrow \infty$, then $T \rightarrow \infty$, and finally $\delta \rightarrow 1$, the strategy profile $(\sigma_1^*, \dots, \sigma_n^*)$ generates a payoff vector that is arbitrarily close to the utilitarian optimum, and therefore to the Pareto frontier of the set of feasible payoffs in $\Gamma(\eta, \lambda, T, \delta)$. Coupled with Lemma 2(i), this implies that each candidate i 's equilibrium payoff from any block where she is eligible must be arbitrarily close to $\mathbb{E}_\theta[\pi_i^\eta(\theta)[B - c(a^*)]]$. We can then use Lemma 1

to show that if candidate i deviated too often from σ_i^* in an equilibrium, then her expected payoff would be bounded away from $\mathbb{E}_\theta \left[\pi_i^\eta(\theta) [B - c(a^*)] \right]$. It follows that in any PBE of $\Gamma(\eta, \lambda, T, \delta)$, each candidate i must play in accordance with σ_i^* in an arbitrarily large proportion of blocks with an arbitrarily high probability. Finally, Fudenberg and Levine's (1983) existence theorem guarantees that such an equilibrium exists, thus completing the proof of Lemma 3.

Proof of Lemma 3. As explained above, we prove Lemma 3 in two steps. Step 1 shows that the payoff vector $(U_1(\sigma^*), \dots, U_n(\sigma^*))$ is arbitrarily close to the Pareto frontier of the set of payoff vectors of $\Gamma(\eta, \lambda, T, \delta)$, so that the payoff to each candidate $i \in C$ must be arbitrarily close to $v_i^\eta \equiv \mathbb{E}_\theta \left[\pi_i^\eta(\theta) [B - c(a^*)] \right]$ in any PBE. Step 2 then uses Step 1 to establish that each candidate i chooses the messages and actions prescribed by σ_i^* arbitrarily often, so that the voter's expected payoff approximates u^* in any PBE (Lemma 2(ii)). Finally, as action sets are finite in $\Gamma(\eta, \lambda, T, \delta)$, it follows from Fudenberg and Levine's (1983) existence theorem that such a PBE exists.

Step 1. Recall from the proof of Lemma 2 that, given any strategy profile σ , the payoff to candidate $i \in C$ in $\Gamma(\eta, \lambda, T, \delta)$ can be written as

$$U_i(\sigma) = (1 - \delta^T) \mathbb{E}_\sigma \left[\sum_{b=1}^{\infty} \delta^{(b-1)T} \mathbf{u}_i^b(\sigma) \right],$$

where $\mathbf{u}_i^b(\sigma)$ represents i 's discounted payoff in block b .

Consider the choice of a strategy profile $\sigma = (\sigma_1, \dots, \sigma_n)$ be a utilitarian social planner who seeks to maximize $W \equiv \sum_{i \in C} U_i$. For every (arbitrarily small) $\gamma > 0$, let σ^γ be a strategy profile that satisfies $W(\sigma^\gamma) \geq \sup_\sigma W(\sigma) - \gamma$. Now, for any block $b \in \mathbb{N}$, consider a deviation at the start of b (by the social planner) from σ^γ to the strategy profile σ^b , which coincides with $\sigma^* = (\sigma_1^*, \dots, \sigma_n^*)$ in block b , and with σ^γ (taken at the null history, regardless of the previous history of play) from block $b + 1$ onward. What is the impact of such a deviation on W ? First, the social planner may incur a loss in block b , which is bounded above by $(1 - \delta) \sum_{s=1}^T \delta^{s-1} [B - c(0)] = (1 - \delta^T)B$. Second, she may incur a loss in block $b + 1$ caused by the change in the expected values of the φ_i^{b+1} 's and in the elected candidates' actions. But this loss is also bounded above by $(1 - \delta^T)B$. Third, as we saw in the proof of Lemma 2, the probability that all candidates remain eligible at the end of block b under σ^b is arbitrarily close to one. Hence, the social planner may also obtain a gain by increasing the expected number of candidates who remain eligible at the end of block b . More precisely, for every realization of the number of eligible candidates $C^{b+1} \equiv \{i \in C : e_i^{b+1} = 0\}$ at the end of block

b , this gain is (approximately) bounded below by

$$(1 - \delta^T) \sum_{i \in C \setminus C^{b+1}} \sum_{b=1}^{\lambda} \delta^{(b-1)T} [v_i^\eta - 0] \geq |C \setminus C^{b+1}| (1 - \delta^{\lambda T}) \min_i v_i^\eta$$

for sufficiently large T and δ (recall from Lemma 2(i) that each candidate i 's payoff is arbitrarily close to v_i^η in every block under σ^*). Dividing this gain by $(1 - \delta^T)$ and letting $\delta \rightarrow 1$, we obtain $(1 - \delta^T)^{-1} (1 - \delta^{\lambda T}) |C \setminus C^{b+1}| \min_i v_i^\eta \approx \lambda |C \setminus C^{b+1}| \min_i v_i^\eta$; so that, whenever $C^{b+1} \neq C$, the social planner's gain increases without bound with λ . It follows that if we let $\lambda \rightarrow \infty$ (and $\gamma \rightarrow 0$), the probability that any candidate becomes ineligible at the end of any block b under σ^γ , $\Pr_{\delta^\gamma} \{C^{b+1} \neq C\}$, must converge to zero — otherwise the deviation would give the social planner a payoff greater than $W(\sigma^\gamma) + \gamma$. Hence, the sum of the expected payoffs induced by σ^γ under the eligibility constraints of $\Gamma(\eta, \lambda, T, \delta)$ is arbitrarily close to the sum of the payoffs which σ^γ would induce if these constraints were ignored (i.e., if all candidates always remained eligible with probability one). Lemma 1 (coupled with the principle of optimality) implies that, in the absence of eligibility constraints, the sum of the payoffs would be maximized by choosing σ^* . Therefore, $W(\sigma^\gamma)$ must be arbitrarily close to $W(\sigma^*) \approx \sum_i v_i^\eta$ or, put differently, $W(\sigma^*)$ must be arbitrarily close to the Pareto frontier. Combined with Lemma 2, this implies that each candidate i 's payoff in any PBE of $\Gamma(\eta, \lambda, T, \delta)$ must be arbitrarily close to v_i^η .

Step 2. Take any candidate $i \in C$. We know that for sufficiently large λ, T and δ : (i) $U_i(\sigma)$ must be arbitrarily close to v_i^η for every PBE σ of $\Gamma(\eta, \lambda, T, \delta)$ (step 1); and (ii) i 's equilibrium payoff at the start of every block must be approximately bounded below by v_i^η (otherwise, by Lemma 2(i), she could profitably deviate). Therefore, for every PBE σ , we must have $\mathbb{E}_\sigma [\mathbf{u}_i^b(\sigma)] \approx v_i^\eta$ for all b . By the same logic as in step 1, this implies that the probability that i becomes ineligible in any block b must be close to zero in equilibrium — otherwise, for arbitrarily large λ , the difference between i 's equilibrium payoff (from b onward) and v_i^η would also be arbitrarily large.

Now let $\eta \approx 0$; let $\kappa \equiv \min_{a \neq a^*} [R - c(a) - (R - c(a^*))] = \min_{a \neq a^*} [c(a^*) - c(a)] > 0$; and let $\sigma = (\sigma_1, \dots, \sigma_n)$ be any PBE of $\Gamma(\eta, \lambda, T, \delta)$, so that $U_i(\sigma) \approx v_i^\eta$ for every $i \in C$. For every block $b \in \mathbb{N}$, let $\rho_i^b(\sigma)$ be the expected proportion of the periods in block b , among those in which candidate i is elected, where it fails to choose a^* ; that is, $\rho_i^b(\sigma) \mathbb{E}_\sigma [\sum_{t=(b-1)T+1}^{2T} \pi_i^\eta(\vartheta_i^t)]$ is approximately the expected number of periods in which i is elected and does not choose a^* — recall that the φ_i^b 's

are smaller than η in absolute value. For sufficiently large T and δ , we thus have

$$\begin{aligned}\mathbb{E}_\sigma[\mathbf{u}_i^b(\sigma)] &\approx \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=(b-1)T+1}^{2T} \pi_i^\eta(\vartheta_i^t) [R - c(a_i^t)] \right] \\ &\geq \frac{1}{T} \mathbb{E}_\sigma \left[\sum_{t=(b-1)T+1}^{2T} \pi_i^\eta(\vartheta_i^t) \right] \left[\rho_i^b(\sigma) [\kappa + R - c(a^*)] + [1 - \rho_i^b(\sigma)] [R - c(a^*)] \right] \\ &\approx \mathbb{E}_\theta [\pi_i^\eta(\theta)] [R - c(a^*) + \rho_i^b(\sigma)\kappa] = v_i^\eta + \mathbb{E}[\pi_i^\eta(\theta)] \rho_i^b(\sigma)\kappa \geq v_i^\eta,\end{aligned}$$

where the second approximation follows from Escobar and Toikka's (2013) Lemma 5.1(ii). As candidate i must be eligible arbitrarily often and her expected payoff must be arbitrarily close to v_i^η (step 1), we must have $\rho_i^b(\sigma) \approx 0$ for an arbitrarily large proportion of blocks b with a probability arbitrarily close to one. As we saw in the main text, it follows from Lemma 1 that if candidate i expects to play a^* arbitrarily often in block $b+1$ and is highly likely to remain eligible at the end of block b , then it is optimal for her to play in accordance with σ_i^* arbitrarily often in block b . We conclude that in any PBE of $\Gamma(\eta, \lambda, T, \delta)$, each candidate i must play in accordance with σ_i^* in an arbitrarily large proportion of blocks with an arbitrarily high probability. Finally, Fudenberg and Levine's (1983) existence theorem guarantees that such an equilibrium exists, thus completing the proof of Lemma 3. \square

References

- AUMANN, R.J., M.B. MASCHLER, AND R.E. STEARNS (1995) *Repeated Games with Incomplete Information*. MIT Press. Cambridge, MA.
- ESCOBAR, J.C., AND J. TOIKKA (2013) "Efficiency in Games with Markovian Private Information," *Econometrica* 81, 1887-934.
- FUDENBERG, D., AND D. LEVINE (1983) "Subgame Perfect Equilibria of Finite and Infinite Horizon Games," *Journal of Economic Theory* 31, 251-268.
- LAFFONT, J.-J., AND J. TIROLE (1993) *A Theory of Incentives in Procurement and Regulation*. MIT Press.
- MAILATH, G.J., AND L. SAMUELSON (2006) *Repeated Games and Reputations: Long-run Relationships* Oxford University Press.