WORKING PAPERS

N° 1646

May 2025

"Political Accountability with Outsiders"

Emmanuelle Auriol and Anaïs Dahmani-Scuitti

Toulouse School of Economics

Political Accountability with Outsiders*

Emmanuelle Auriol†, Nicolas Bonneton‡, Mattias Polborn§

May 16, 2025

Abstract

We present a moral hazard model of electoral accountability that challenges the common view of the populist vote as mere frustration with the elite. Rational voters use the threat of electing outsiders to incentivize more competent insiders whose policy preferences diverge from those of voters. Their optimal retention strategy involves differentiated punishment for failing incumbents, replacing them either with other elite politicians or with outsiders. The latter only occurs when the incumbent's policy is both perceived as a failure *and* as benefiting the elite. This strategic voting behavior explains why outsider electoral success is often volatile: rational voters may back an outsider in one election and an establishment candidate in another, without changing their fundamental preferences.

*We have received useful comments from Roland Bénabou, Navin Kartik, Christopher Li, Chirstopher Sandmann, Nicolas Schutz, Konstantin Sonin, Richard van Weelden, and especially Cesar Martinelli, as well as seminar audiences at Michigan State University, Pittsburgh & Carnegie Mellon University, Vanderbilt University, and the Wallis Institute Annual Conference. Emmanuelle Auriol acknowledges TSE-IAST funding from the French National Research Agency (ANR) under the Investments for the Future (Investissements d'Avenir) program, grant ANR-17-EURE-0010. Bonneton gratefully acknowledges financial support from the German Research Foundation (DFG) through CRC TR 224 (Project B04).

†Toulouse School of Economics, emmanuelle.auriol@tse-fr.eu

‡Vanderbilt University, nicolas.bonneton@gmail.com

§Vanderbilt University, mattias.polborn@vanderbilt.edu

1 Introduction

A backlash against the ruling elites has led to a rise of populist parties throughout the Western world in the last two decades. While elections were previously framed as a choice between the traditional left- and right-wing parties, voters now increasingly choose between "insider" and "outsider" candidates. The former are from mildly differentiated parties with a long track record of governing, but frequently perceived as out of touch with regular voters; the latter are often called populists and are characterized by their lack of experience, radically different proposed policies, and antagonism toward the elite. Our paper shows that it can be optimal—even for a *rational* voter with no intrinsic preference for outsider platforms—to occasionally support less competent outsiders, as a way to discipline insiders.

To study how outsider candidates affect insiders' behavior and electoral outcomes, we develop an infinite horizon model of political accountability with insider and outsider challengers, with their type being public information. The framework rests on three main premises: (i) Insiders are more competent than both outsiders and voters, but their preferences may diverge from those of the median voter;[1] (ii) Politicians are both office-motivated and policy-motivated; (iii) While voters observe the chosen policy (hard information), they only get a noisy signal of whether the policy was appropriate and actually benefits them (soft information).[2]

As is well known in political accountability, voters lack control over most contractual terms (e.g., salary) and so the threat of alternation — replacing the incumbent with another candidate — acts as the primary way for voters to incentivize incumbents to follow their interests. Incumbents fear losing the election because they value both the perks of office and their ability to choose policies they prefer. While an incumbent loses the perks of office when he is replaced by another insider politician, he can take comfort in the fact that his replacement will choose relatively similar policies. Voters can increase pressure on insiders by threatening to replace them with a candidate they detest, someone who could undo their policies and shift policy direction in ways insiders do not want. We characterize when this threat is effective and when voters are willing to carry it out.

We show that outsider victories require two events to happen simultaneously: dissatisfaction

[1] In the words of Carnes and Lupu (2023), "Politicians everywhere are significantly better off than the people they govern." The difference in education, wealth, and occupation between voters and the political elite implies that they may disagree with the policy. The insiders are more qualified but cannot necessarily be trusted to use this information to implement policies that are good for the common people (Gilens, 2012).

[2] With imperfect monitoring, punishments must occur on the equilibrium path; as in Green and Porter (1984).

with the incumbent's performance (e.g., a crisis), *and* the suspicion that the elite benefited from the crisis (e.g., the policy implemented seems to align with insider's particularistic interest).[3] Both events occasionally arise on the equilibrium path of our model. The empirical literature often finds that negative economic shocks and a lack of trust in elites are central to turning voters away from political insiders (Berman, 2021; Guriev and Papaioannou, 2022; Rodrik, 2021). Our model suggests that these factors are complementary, as one alone would not suffice. If voters are merely dissatisfied with the incumbent's performance, but are not suspicious of the incumbent's policy choices, they may replace him with another insider candidate, but not with an outsider. This result contrasts with a common view of the outsider vote as an unsophisticated antagonism against the elite.

For the voters in our model, electing a populist after a disappointing performance of an insider incumbent is an "investment;" they recognize that electing an outsider is costly in the short run, but could be worthwhile in order to keep their punishment threat credible to generations of future insider politicians. Such temporary punishment is sequentially rational even if it is not a Nash equilibrium in the stage game (i.e., if the next election would be final), a common feature of repeated games (Mailath and Samuelson, 2006). This result speaks to the question of why voters are sometimes willing to elect outsiders even though they appear economically to be a cure worse than the disease of an out-of-touch insider (e.g., Edwards, 2019, Funke et al., 2021). It is also consistent with the observation that younger voters (relative to older ones) appear more willing to vote for populists.

That said, we find that the threat of outsider replacement can backfire, limiting its effectiveness. In some cases, even if punishing an incumbent by electing an outsider were costless, doing so would be counterproductive. A harsher punishment threat of being replaced by an outsider increases an insider incumbent's incentive to be reelected. However, whether this added pressure improves incentives to behave in a way that benefits voters is ambiguous. On the plus side, in some cases, this threatened harsher punishment may induce the incumbent to pursue what the voters want rather than his own

[3] As an example for the "suspicious" policy and the type of conflict between voters and elite politicians that we have in mind, consider the decision whether the government should give aid to another country facing a financial crisis (e.g., Greece in 2009). If not helping the foreign country would lead to a massive financial crisis with significant contagion, both the elite and (most) voters agree that such an outcome should be avoided by providing aid. Similarly, if it were known that the foreign country would be able to handle the situation without significant contagion, then the elite and voters would agree not to bail out the other country. However, we may also find ourselves in a third state of disagreement in which not providing aid may result in significant losses for some national banks, but without major impact on the national economy. In this case, elite politicians (who are likely closely connected to the owners of the lender banks that a bailout helps) may prefer a bailout, while the representative voter has the opposite preference. Intuitively, the fact that the elite prefers a bailout in more states of the world than the voters makes a bailout a suspicious policy.

agenda. On the other hand, it may deter incumbents from taking actions that are justified by the state of the world but that appear suspicious to the median voter. In effect, a threat of replacement by an outsider then may induce some "faux populism by the elite" in a way that is detrimental to voters.[4]

Another notable result is that lower-quality information makes voters more willing to elect outsiders and also increases the frequency with which they do so. This is consistent with empirical evidence (e.g., Guriev et al., 2021, Allcott and Gentzkow, 2017). One key reason for this is the following mechanism. Consider first a case with very high signal quality so that the incumbent is rarely replaced if he behaves well (and almost always replaced if he does not); thus, threatening to replace the incumbent by another insider is sufficient to incentivize good behavior. Suppose now that the signal quality deteriorates; the incumbent now faces an increased risk of being replaced by mistake. This uncertainty diminishes the incumbent's continuation payoff and could lead him to misbehave. In this case, it may be necessary for voters to threaten a harsher penalty through replacement by an outsider.

Finally, it is important to emphasize that our results do not rely on the presence of exogenously unsophisticated outsiders. We show that voters do not need access to genuine outsider candidates for the disciplinary threat they create to be effective. Reelection incentives alone can induce an insider to credibly commit to an outsider-style platform. The career trajectories of those who implement the outsider-style program differ according to their type. A true outsider disrupts established policies for a single term before ultimately being removed by the voters. By contrast, an insider who mimics an outsider can enjoy a longer political career, by transitioning from a populist platform to an insider platform. As a result, outsider-style politics can emerge even if the political system is shielded from outsider entry.

Section 2 reviews the related literature. Section 3 presents the model. For the analysis, we begin in Section 4 with a benchmark case of perfect monitoring to build intuition, before turning to the general case in Section 5. All proofs of formal statements are in the Appendix.

[4]This effect seems, on the surface, similar to what has been discussed in the pandering literature, but (as we discuss in Section 2) the underlying forces are quite different. In the pandering literature, the fundamental cause for pandering is that the electorate is uncertain about the incumbent's type and uses the incumbent's policy choices to learn about his competence. There is no uncertainty over the politicians' preferences and skills in our model.

2 Related literature

Our paper connects with and contributes to three distinct literatures that are rarely connected: political accountability, pandering and populism.

Political accountability. The literature on political accountability, pioneered by Barro (1973) and Ferejohn (1986) and recently reviewed by Duggan and Martinelli (2017), analyzes how voters can induce politicians to behave well in office, by conditioning reelection on their behavior and the voters' information about the appropriateness of their action. The existing literature typically considers the effects of threatening replacement by one (often exogenously-given) type of opponent. In the classical political agency models, the type of replacement is immaterial for the incumbent's utility.

One exception is Van Weelden (2013), where, similar to our setup, politicians have policy preferences and publicly observable types, making challengers' types crucial for incumbent incentives.[5] In equilibrium, only candidates with policy preferences sufficiently far to the left or right of the median voter are elected. The threat of being replaced by a candidate from the opposing side effectively incentivizes the incumbent to work hard.

There are several important differences between our paper and Van Weelden (2013). First, because signals in our model are imperfect, incumbents are in fact replaced with positive probability along the equilibrium path of our model; moreover, imperfect observability also generates the potential problem of pandering that we describe below. Second, insider and outsider candidates in our model are both vertically and horizontally differentiated (i.e., have different preferences and different capabilities); and while, in van Weelden's model, the main focus is on how polarized the candidates have to be in order to provide effective incentives for each other, our main focus is on the circumstances under which the threat of outsider candidates can improve the incentives for elite politicians and occur on the equilibrium path.

A model that shares with ours the feature that it is sometimes optimal for voters to elect "bad" candidates to punish better ones is Anesi and Buisseret (2022), though their mechanisms are quite different from ours. They analyze a framework of political accountability with both moral hazard

[5] There is also a strand of the accountability literature that focuses on the effects of politician's *unobservable* policy preferences on accountability (Bernhardt et al., 2011, 2009; Duggan, 2000; Van Weelden, 2015) and analyzes how politicians with undesirable private preferences attempt to hide this feature from voters. Note that, when types are private, the incumbent cannot differentiate between challengers.

and adverse selection in which politicians do not care about who replaces them. In their model the number of available politicians is limited, so "bad" candidates serve as tools to rotate others out of power. They show that, for a sufficiently high discount factor, voters can (almost) implement their first-best payoff.

Pandering. The incumbent in our model receives a private signal on the state of the world, and has to choose policy based on this signal, like in the literature on pandering, pioneered by Canes-Wrone et al. (2001).[6] The pandering literature, however, is mostly concerned with the distortions that arise because the politician wants to influence the voters' inference about her unknown type.

From this literature, the closest paper to ours is Maskin and Tirole (2004) where politicians have policy preferences that can differ from the voter and voters observe the policy implemented as well as a signal of its success, like in our model. A key difference is that Maskin and Tirole (2004) assume politicians care only about policy while in office, making the identity of their replacement—central to our paper—inconsequential in theirs. Additional differences include the nature of uncertainty: uncertainty over the incumbent's type in their model versus uncertainty over the state in ours. And, while most pandering models are set in two-period frameworks (in order to remain tractable in the presence of learning), we analyze an infinite-period framework.

Populism. We also contribute to the growing literature on politicians lacking experience, offering simplistic solutions, and opposing the elite—often called outsider, populist, demagogue, or incompetent (see Berman, 2021; Guriev and Papaioannou, 2022; Mudde and Kaltwasser, 2017; Müller, 2017 for reviews).

Our article provides a rational theory of voter support for populists; other papers analyzing different channels and justifications for why populist politicians are sometimes elected include, among others, Acemoglu et al. (2013); Agranov et al. (2020); Bellodi et al. (2023); Gratton and Lee (2023). For example, in Acemoglu et al. (2013), voters worry that some politicians are open to being bribed by right-wing interest groups, and politicians overcompensate in their policy choice to counteract that perception.[7]

[6] See also, e.g., Fox and Stephenson (2011) and Morelli and Van Weelden (2013).

[7] As discussed in the paragraph on pandering similar behavior – mainstream politicians choosing excessively populist actions to ensure reelection – is also sometimes a danger in our framework and limits the voters' ability to provide effective incentives.

In our model, the focus is on how elite politicians react to the presence of populists (or even other insiders who just run as populists). Voters, in turn, choose populists as a strategic punishment tool – an approach we believe is novel in this literature. Indeed, our model assumes voters to be *rational* and only *strategically* supporting outsiders.

In practice, however, outsiders often have sincere supporters. In addition, as stressed by the literature, outsiders may attract "behavioral" voters who are present-biased (e.g., Bernhardt et al., 2022), prefer simplicity (Levy et al., 2022), or are betrayal averse (Di Tella and Rotemberg, 2018). We view these explanations as complementary to ours, as different types of voting behaviors can coexist among various voter groups. Nevertheless, we believe that accounting for strategic support of outsiders is relevant to captures the large swings in populist vote share, as this type of support is inherently more volatile.

Second, we contribute to theories on support for incompetent politicians (e.g., Gailmard, 2022; Mizuno and Okazawa, 2022), showing they can act as a short-term punishment tool. Third, and less related, is the body of work on protest voting, which analyzes settings where voters signal their dissatisfaction with their most-preferred political party (e.g., Kselman and Niou, 2011, Myatt, 2017). In our model, voting for the outsider has no signal value, but rather functions solely as a form of retribution.

3 The model

We consider an infinite-horizon discrete time model with discount factor β, where one period corresponds to the time between two elections (e.g., 4 years). In each period, a representative voter elects, or reelects, a politician from a large pool of insider and outsider candidates, with their type being public information.[8]

The incumbent then chooses between two policies: one typically associated with favoring the "elite", denoted E, and the other with favoring the "masses", denoted M. The actual alignment of these policies with elite or mass interests depends on the state of the world $w \in \{w_M, w_D, w_E\}$ which is drawn from a known distribution (p_M, p_D, p_E). In state w_M, both insiders and voters prefer

[8]Considering a small number of qualified candidates (who therefore have a path to eventually returning to power), or equivalently considering political parties, rather than individuals, would not qualitatively change our results, but rather reinforce the representative voter's desire to elect an outsider. It is indeed harder to motivate politicians to behave well when they can return to power after losing an election, as shown in Online Appendix O.3 where we consider a two-party version of our model in which there are a moderate left party and a moderate right party.

policy M; in state w_E, both prefer policy E. In the disagreement state w_D, insiders prefer policy E, while the voter prefers policy M. The insiders observe the state of the world before making a decision, while the voter remains uninformed at this stage. We assume that $p_E < \frac{1}{2}$, meaning the voter's ex-ante preferred policy is M.

Outsiders share voter preferences and information, and therefore always implement policy M when in power. Alternatively, as we will show in Section 4.3, insiders in our model can credibly commit through reelection incentives to an outsider platform— playing M no matter what. As such, our analysis carries through even when only insider candidates are available, allowing for two interpretations of politicians offering simplistic and anti-elite policies.[9]

After the policy is implemented and prior to the next election, the representative voter (imperfectly) observes whether the policy matches the true state of the world. Formally, before the next election, the voter receives a binary signal regarding whether the implemented policy aligns with their preferred policy. If the policy matches the state of the world, the signal is "good" with probability α and "bad" with probability $(1-\alpha)$. If the wrong policy is chosen, the signal is always bad. As such, receiving "good" is definitive good news, simplifying the voter's decision to re-elect the incumbent, while receiving a bad signal introduces ambiguity. This asymmetry simplifies the model and is akin to the learning literature (e.g., Bolton and Harris, 1999, Keller et al., 2005), where agents receive binary signals— either conclusive, or ambiguous.[10]

Voters and politicians receive a policy payoff, π, when their preferred policy is implemented, and payoff 0 otherwise. However, for voters, this payoff does not materialize immediately; it becomes apparent only after a delay.[11] Consequently, and consistently with models of informational asymmetry between voters and incumbents, it cannot serve as a variable in the re-election strategy. While in office, the politicians also earn ego-rents ϕ, and zero otherwise. These rents reflect benefits like control over public resources, prestige, and influence. Lastly, in line with most of the literature, we

[9] Yet another interpretation of the outsiders' behavior is that they may be well-informed but hold preferences different from the median voter such that they prefer policy M independent of the state of the world.

[10] Even with more general signal structures, we conjecture that our key qualitative results remain robust, as the voter would still re-elect the incumbent in the most optimistic case, which occurs when they receive a good signal.

[11] Alternatively, we can reinterpret the signal in the model in the following, mathematically equivalent, way. At the end of each period, the voter gets an intrinsic payoff, $\Pi_v > 0$ with probability α; and zero otherwise. Implicit in this formulation is that policies are complex and their effects uncertain, i.e. even if the politicians dutifully implement the preferred voter policy, there is no guarantee that the voter will get a high payoff. In this case, $\pi = \alpha\Pi_v$ can be interpreted as the voter's expected payoff from her preferred policy. In this interpretation of the model, we essentially assume that the incorrect policy is never successful, while even a correct policy may fail to deliver a success with probability $1-\alpha$.

consider that losing incumbents cannot subsequently return to office.[12]

Within each period: (i) the state is drawn, (ii) the elected incumbent chooses a policy, (iii) a public signal is realized, (iv) a new election occurs. A strategy for the voter maps the public history (the sequence of past policies, incumbent types, and observed signals) into a re-election decision after each period. A strategy for the politician specifies which policy to implement in each period, as a function of the state and public history. As is standard in the literature (e.g., Barro, 1973, Ferejohn, 1986), we restrict our attention to pure strategies,[13] and to the voter's *preferred* perfect Bayesian equilibrium (PBE)—specifically, the equilibrium in which their desired policy is implemented most frequently.[14]

The parameter $\alpha \in [0, 1]$ plays a central role in our analysis of (re-)election strategies, as higher values of α correspond to greater voter information quality (or, equivalently, lower policy complexity). To highlight its significance, we begin by examining the case in which the voter *perfectly* observes whether the implemented policy aligns with the true state.

4 Full Information

We examine the limit case in which the signal is fully informative: $\alpha = 1$. This benchmark helps build intuition about the pros and cons of making outsiders part of the voter's retention strategy.

4.1 First-best: Efficient political alternation

We begin by identifying the maximal set of parameters for which the voter obtains their preferred policy in equilibrium, without relying on outsiders, neither on- nor off-path. This is a natural starting point, as the literature typically assumes a single type of challenger and models incentives as stemming mainly from the perks of holding office. Such an equilibrium is sustained (only) by the following *mainstream retention strategy*: reelect if the policy matches the state; vote for another

[12] See Ferejohn (1986) and Anesi and Buisseret (2022) for notable exceptions.

[13] Voters cannot commit to random threats credibly. Voting is private and decentralized, with no mechanism to bind voters to punish with a specific probability, nor can politicians observe or verify the randomization process. The temptation to claim they will randomly punish but then not follow through is too strong. Furthermore, while our model formally describes the electoral process as the decision of a representative voter, she can be seen as a proxy for a multitude of moderate voters who are decisive in choosing the winning candidate. Coordinating a random electoral strategy among such a large group is complex, and unrealistic.

[14] Beliefs are unproblematic in our model. All action/signal combinations occur on the equilibrium path, so Bayes rule can be used to calculate all beliefs. More importantly, though, the incumbent's type (i.e., which state the incumbent saw last period) is irrelevant for the future as the state is drawn anew in each period.

insider otherwise. Observe that punishing insiders with one another is not costly for the voter.

To determine when the mainstream retention strategy constitutes an equilibrium, we first compute a well-behaved incumbent's utility before observing the state of the world. A well-behaved incumbent is always reelected and thus receives an ego-rent ϕ in each period. Additionally, he receives a policy-payoff π only if the state is w_E or w_M. As such, his ex-ante utility is

$$U^* = \phi \frac{1}{1-\beta} + \pi \frac{1-p_D}{1-\beta}.$$

To check if this is an equilibrium, we must compute the incumbent's payoff when deviating, while holding other politicians' strategies constant. The incumbent has no incentive to misbehave in states w_E or w_M, as his policy preferences align with the voter's, and choosing the correct policy ensures reelection. In the disagreement state w_D, if the incumbent chooses policy M, his continuation utility is: $U^*(w_D) = \phi + 0 + \beta U^*$. If the incumbent deviates and chooses E, he gains a policy-payoff π in the current period, at the cost of not being reelected. The continuation utility from deviating is thus: $U^d(w_D) = \phi + \pi + \beta\pi\frac{1-p_D}{1-\beta}$. Well-behaving is an equilibrium if and only if $U^*(w_D) \geq U^d(w_D)$. Substituting and simplifying, this condition is equivalent to

$$\frac{\pi}{\phi} \leq \frac{\beta}{1-\beta}. \tag{1}$$

Thus, implementing the voter's preferred policy in every period can be sustained in equilibrium through political alternation among insiders if and only if the discounted benefit of holding office, $\frac{\beta}{1-\beta}\phi$, exceeds π, the immediate payoff from implementing the preferred policy. When politicians care more about policy, i.e., when π is high, it is harder to align the incumbent's incentives with the voter's preferences. Conversely, larger perks of holding office ϕ encourage good behavior.

4.2 The outsider threat

Whether an insider incumbent is reelected or replaced by another incumbent does not change future policy. As such, when insiders value policy sufficiently strongly relative to the personal benefits of holding office, the threat of alternation may not be enough to make them prioritize voters' preferences. In this case, electing an outsider—who poses a threat to the insider's political legacy and privileges—can serve as a more impactful, albeit *costly*, form of punishment. Voters are, under cer-

tain conditions, willing to execute their threat because, if they fail to do so, their credibility is lost, rendering future threats ineffective (Mailath and Samuelson, 2006).[15]

To make this point, we consider the least costly (off-path) form of outsider punishment for the voter, i.e., replacing an insider with an outsider for exactly one period before returning the helm of government to another insider. We refer to this retention strategy as *the outsider threat*. While longer punishments can expand the set of parameters under which insiders are disciplined (see the online Appendix O.1),[16] focusing on this one-period strategy is sufficient to understand the main trade-off between the cost of outsider punishment and the incentives it provides.

To determine when there is an equilibrium in which insider incumbents always act in the interest of voters under the outsider threat, we need to examine two questions: First, if an insider incumbent believes in the threat, under which conditions is he sufficiently incentivized to always deliver the voter's preferred policy? And second, if the voter observes a misbehaving incumbent, is she willing to punish the incumbent by electing an outsider?

Proposition 1 summarizes the results.

Proposition 1 *The implementation of the voters' preferred policy in every period can be sustained in equilibrium with the outsider threat if the following two conditions hold simultaneously:*

1. *The representative voter is willing to elect an outsider in response to a deviation:*

$$p_E\,(1-\beta) \leq p_D \tag{2}$$

2. *Given this threat, the insider incumbent is well-behaved:*

$$\frac{\pi}{\phi} \leq \frac{\beta}{1-\beta}\frac{1}{1-\beta p_E} \tag{3}$$

Proof: See Appendix A.1. □

[15] To characterize the maximal set of parameters for which the voter's preferred equilibrium can be sustained, we follow a standard approach in the literature and construct a continuation equilibrium in which any failure to carry out a prescribed punishment leads to the voter's worst credible continuation. In our setting, this amounts to the voter choosing, in every period, between two fixed outcomes: (i) a misbehaving insider or (ii) an outsider in power. If $p_E \geq p_D$, the voter selects the insider (either never replacing them, or occasionally doing so with another insider). Conversely, if $p_D \geq p_E$, the voter strictly prefers the outsider and replaces the incumbent in every period.

[16] Online Appendix O.1 provides a characterization of the maximal set of parameters sustaining the voter-preferred outcome when she is implementing a $n-$period punishment.

Condition (3) is weaker than Condition (1) (because $\frac{1}{1-\beta p_E} > 1$), implying that the outsider threat is more effective in incentivizing compliant behavior than simple alternation between insiders. As a result, the outsider threat expands the set of parameters under which the voter obtains her preferred policy in equilibrium. The size of this expansion increases with p_E, since the incumbent is more easily disciplined if the expected cost of electing an outsider is high from his perspective (since the outsider makes the wrong decision in state E).

The main takeaway from Proposition 1 is that, in some cases, the voter's first-best equilibrium can be achieved *only* if the outsider is part of her retention strategy. This can even happen if the voter would prefer to be ruled by a misbehaving insider rather than an outsider.

To be more specific, if state w_E is more likely than state w_D (i.e., if $p_D < p_E$), the voter prefers a misbehaving insider to an outsider. In a one-shot election, voting for an outsider would not constitute an equilibrium. The key question, therefore, is whether it is optimal to vote for an outsider in the context of repeated interactions. Intuitivelly, voting for an outsider here entails an immediate cost but yields a long-term benefit by disciplining insiders. As a result, the discount factor β plays a crucial role in (2): more patient voters are more willing to incur the short-term cost of electing an outsider. In the limit, when $\beta = 0$, equation (2) simplifies to $p_E \leq p_D$—since the voter does not value the future benefits of disciplining insiders, she simply votes for whoever performs better in the current period.[17]

If electing an outsider for one period is insufficient to discipline insiders (i.e., Condition (3) is not satisfied), the voter must consider a harsher punishment. The harshest punishment is the grim trigger strategy, in which any deviation by the incumbent results in the permanent election of outsiders.[18] This strategy constitutes an equilibrium only if $p_E \leq p_D$, i.e., electing an outsider is a Nash equilibrium of the stage game. Clearly, sustaining the grim trigger punishment is more challenging for the voter. In Online Appendix O.1, we analyze the case with multi-period punishments to highlight, more generally, that longer punishments provide better incentives for insiders. However, they also come at a higher cost for the voter.

[17] If $p_D > p_E$, misbehaving insiders are more costly to voters than outsiders. Thus, even in the absence of long-term reputation concerns, at the start of a punishment subgame, the voter prefers electing the outsider to retaining a misbehaving insider in that period. Furthermore, by deploying the outsider threat, the representative voter ensures a return to well-behaved insiders after the current punishment phase.

[18] In which case implementing the voter's ideal program is incentive-compatible for an insider if and only if $\frac{\pi}{\phi} \leq \frac{\beta}{1-\beta-p_E\beta}$ (see Online Appendix O.1 for details of the computations).

Patient vs. impatient voters. Proposition 1 illustrates that stronger punishment acts as a long-term investment. Indeed, we can see in conditions (2) and (3) that if the future is sufficiently important (β is large), then each condition is more likely to be satisfied. For this investment to be justified, individuals must exhibit sufficient patience. This finding contrasts with the conclusions of Bernhardt et al. (2022), which suggest that populism arises when voters focus predominantly on short-term preferences. This raises important questions about how different age groups evaluate political trade-offs, the role of patience and intertemporal preferences in political decision-making. Our results offer an intriguing interpretation into the behavior of younger voters who often are more likely to support populist candidates than older voters. Younger voters could be inherently more concerned with long-term outcomes and be willing to endure the short-term costs of populist policies to strengthen the incentives for insiders to act responsibly. This perspective challenges conventional narratives about populism being driven purely by short-term thinking and immediate gratification in their decision-making. Instead, it highlights the possibility that younger voters may view populism as a tool to enforce accountability and reshape insider behavior over the long term.

4.3 Insiders offering an outsider program

If the voter needs to incentivize incumbents by threatening them with outsider replacement, do they need access to "true" outsiders (i.e., agents who are actually less competent than insiders)? The answer is no. We now show that an insider can credibly commit through reelection incentives to an outsider program.

Consider the following voter retention strategy that uses only insiders: "If an insider is elected with a mandate to implement M and successfully fulfills this promise, reelect him in the following period on an insider platform—that is, to implement the voter's preferred policy. Otherwise, replace him with another insider who receives the same arrangement: namely, the new insider promises to implement M and, if he delivers, is reelected on the insider platform."

Following this strategy, the voter can incentivize an insider to mimic a populist for the punishment phase, provided that the condition from Proposition 1 are met.

To see this, observe that the trade-off for the "fake outsider" in his first period in office is the following when the state is w_E or w_D.[19] By deviating to policy E, the fake incumbent gets a policy

[19]In state w_M, it is obvious that there is no incentive to misbehave as the insiders policy and career objectives are

payoff of π this period, loses an expected policy payoff of $p_E\pi$ in the next period,[20] and loses the entire stream of office benefits from tomorrow on. But this is the same intertemporal trade off as for the insider incumbent in state w_E in Proposition 1. We deduce easily the next result.

Proposition 2 *Suppose the voter's preferred equilibrium is sustained by the outsider threat, where a misbehaving insider is punished by electing an outsider. Then, this equilibrium can also be implemented by replacing a deviating insider with another insider who credibly commits to an outsider program in the punishment phase.*

Proposition 2 shows that even if the political system is shielded from outsider entry (i.e., a model with only insiders), under certain conditions, the voter's preferred equilibrium can be sustained when insiders can credibly promise to a "populist-like" behavior for one period.

In many cases, the voter is thus indifferent between electing a genuine outsider or a "fake" outsider. However, the career trajectories of those who implement the outsider-style program differ according to their type. A true outsider disrupts established policies for a single term before ultimately being removed by the voters. By contrast, an insider who mimics an outsider can enjoy a longer political career, by transitioning from a populist platform to an insider platform.

The voter's indifference between a true and a "fake" outsider arises from the model's parsimony. In reality, several factors not captured here could break this indifference. For example, electing a true outsider may involve additional risks not explicitly modeled, such as a small probability of permanent institutional breakdown. In such cases, voters would no longer remain indifferent. Electing a fake outsider—an opportunistic insider—offers a controlled form of punishment: because it is administered by an experienced and familiar insider, voters can expect greater adherence to institutional norms. However, the drawback is that such punishment may be perceived as too weak to effectively discipline insiders. In particular, there may be a small probability that the opportunistic insider reneges on the promise to implement M, especially in states of nature where fulfilling it would be very costly.

perfectly aligned with the voter's objectives.

[20] This is because, after a deviation, another fake populist will be in office next period, while if the current incumbent chose M today, he would be able to choose policy E in state w_E

5 Imperfect Policy Evaluation

In the baseline model where $\alpha = 1$, voters can perfectly monitor policy outcomes ex post. Under this ideal condition, alternation never occurs along the equilibrium path. Moreover, voters' decisions depend solely on the actual performance of incumbents—that is, whether the correct policy was chosen—and are unaffected by whether a policy appears to favor the elite or the masses.

However, in practice, voters face significant challenges in assessing the causal impact of policies on their payoffs. These challenges arise due to various confounding factors, such as economic complexity, incomplete information, and external shocks (e.g., technological disruptions, globalization, or geopolitical events) that obscure the true relationship between policies and outcomes. Moreover, the proliferation of misinformation or "fake news" further diminishes voters' ability to attribute outcomes accurately, weakening the signal upon which electoral decisions rely.

We model such a noisier information environment by assuming $\alpha < 1$. In this setting, voters cannot perfectly observe the causal effects of policies on their payoffs. Instead, they must rely on noisy signals that may misrepresent whether a policy was appropriate. As we will see, the introduction of imperfect policy evaluation yields several important insights that depart from the baseline case.

To ease the analysis, in this section we restrict attention to equilibria where strategies depend only on the most recent period and a single state variable, *credibility*, in line with standard models of electoral accountability. Credibility in $\{0, 1\}$ remains at 1 as long as the voter has always punished candidates according to her strategy, while failing to punish a deviation causes credibility to permanently drop to 0. Once credibility is lost, it cannot be regained. As shown in Online Appendix O.2, relaxing this restriction allows for more complex punishment schemes—such as voters alternating between different strategies across periods—that enlarge the set of parameters under which the voter's preferred policy can be implemented in every period in equilibrium. While these more sophisticated strategies may be viewed as less realistic,[21] our main result remains valid: the outsider threat still features, under certain conditions, in the preferred voter equilibrium.

[21] One period in our model typically represents a legislative or executive term of four to five years, making it less realistic for voters to condition their vote today on actions that occurred a decade ago.

5.1 First-best: efficient political alternation

As in the baseline model, we begin by identifying the set of parameters and the strategies for which the representative voter's ideal policy can be implemented in every period (i.e., the first best). With noisy information, such an equilibrium can only be sustained by well-incentivized insiders, since outsiders occasionally make mistakes and punishment occurs with positive probability even along the equilibrium path.

A key difference with perfect monitoring is that the voter's optimal retention strategy now depends not only on the incumbent's (perceived) performance but also on policy choices. There are two strategies to consider.[22] The first one, referred to as F_1, reelects the incumbent if and only if the signal is good. Thus, F_1 is quite harsh because any negative signal results in the insider's replacement. The second strategy, referred to as F_2, is more forgiving in that it reelects the incumbent also if the signal is bad, but the policy choice was M. Proposition 3 summarizes the results concerning the conditions under which all insiders who implement the representative voter's preferred policy form an equilibrium.[23]

Proposition 3 *The implementation of the voters' preferred policy in every period can be sustained in equilibrium if and only if*

1. *the representative voter plays F_1; and the following condition holds*

$$\frac{\pi}{\phi} \leq \frac{\alpha\beta}{1-\alpha\beta} \tag{f1}$$

2. *or, the representative voter plays F_2, and the following condition holds*

$$\frac{\beta(1-\alpha)}{1-(p_E\alpha+1-p_E)\beta} \leq \frac{\pi}{\phi} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta} \tag{f2}$$

[22] For a detailed discussion of why other strategies are not relevant to the voter, see the preamble of the proof of Proposition 3 in Appendix B.1.

[23] Even if the conditions of Proposition 3 are met, other equilibria may nevertheless exist. For instance, if $\frac{\beta\alpha}{1-(1-p_D)\alpha\beta} \leq \frac{\pi}{\phi} \leq \frac{\beta\alpha}{1-\alpha\beta}$, one can show that under the strategy F_1 there are two equilibria: the "good" equilibrium where the incumbent implement the voter's preferred policy, and a "bad" equilibrium where he implements his preferred policy (computations are available from the authors upon request). Politicians' decisions hinge on what they believe their successors will do. The multiplicity of equilibria in our framework is due to the strategic complementarities of politicians' strategies: If future incumbents misbehave, this will increase the ruling incumbent's payoff when losing the election, raising his willingness to misbehave.

Proof: See Appendix B.1. □

Figure 1 illustrates Proposition 3, showing when the first-best is an equilibrium. There are four regions in the parameter space: (i) where only F_1 sustains the first-best, (ii) where only F_2 does, (iii) where both do, and (iv) where neither does. Note that when $\alpha = 1$, conditions *(f1)* and *(f2)* coincide: conditioning reelection on the policy choice rather than the outcome only matters under imperfect monitoring.

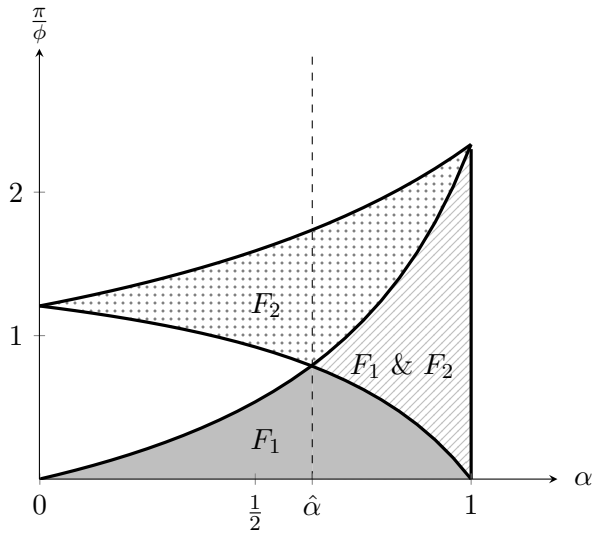

Figure 1: For $p_E = 0.4$ and $\beta = 0.7$, this figure depicts the parameters for which the incumbent can be incentivized to always implement the voter's preferred policy. Dotted area: Only F_2 incentivizes. Dark gray area: Only F_1 incentivizes. Dashed gray area: Either F_1 or F_2 can be used. Elsewhere, no equilibrium that incentivizes correct behavior is sustainable.

Condition *(f1)* ensures that the incumbent behaves appropriately under the reelection strategy F_1. It requires that the discounted value of the ego-rent for an always well-behaved incumbent, $\frac{\beta\alpha}{1-\alpha\beta}\phi$, is greater than π, the incumbent's instantaneous benefit from implementing his preferred policy.[24] Compared to the baseline model with perfect monitoring, the left-hand side of the condition in *(f1)* increases in α. As signal accuracy decreases, we expect political turnover to rise, since well-behaved incumbents are eliminated with probability $1 - \alpha$ due to an incorrect signal. For sufficiently small α—specifically, when α is smaller than the critical value $\frac{\pi}{\beta(\phi+\pi)}$ such that the condition in *(f1)* binds—the $F1$ retention strategy no longer induces insiders to behave well.

[24]Similar to the baseline model with perfect monitoring, a higher policy payoff π makes it more difficult to incentivize the politician to go against his own policy preference. On the other hand, a higher ego rent/salary ϕ, or a higher time discount factor β, make it easier to induce good behavior because it increases the importance of future office benefits.

Condition *(f2)* ensures that the incumbent behaves appropriately under the reelection strategy F_2. By guaranteeing reelection whenever policy M is implemented—and thereby offering a higher discounted value of holding office—strategy F_2 provides stronger incentives than F_1 when behavior in the disagreement state is the central concern. For this mechanism to be effective, however, the ratio π/ϕ can be neither too high nor too low.

Specifically, if π/ϕ is too high,[25] the incentive to implement policy E in state w_D would be too strong, even if the voter's strategy is F_2. When $\frac{\pi}{\phi}$ exceeds the upper threshold $\frac{\beta}{1-(p_E\alpha+(1-p_E))\beta}$ from condition *(f2)*, it becomes impossible to implement the first-best outcome. Since this threshold decreases as p_E increases, it is less likely that the condition in *(f2)* holds (and thus that F_2 works) when p_E is very small.[26] This is because the frequency with which it is necessary to implement M in state E to remain in power (and thus the expected policy cost from deviating) is small.

Interestingly, π/ϕ cannot be too low either,[27] as is typically the case when the ego rent ϕ is large relative to the policy payoff π (e.g., when politicians' salaries are excessively high). In such cases, the more forgiving strategy F_2 fails to implement the first-best outcome, and only the harsher strategy F_1 may remain effective. The underlying problem is that F_2 becomes vulnerable to opportunistic "faux populism": when ϕ is high relative to π, the incumbent may choose to implement policy M not only in the disagreement state (w_D), but also in the consensus state (w_E), solely to secure reelection. In contrast, F_1 deters such behavior by punishing all unfavorable signals, thereby avoiding the pitfalls associated with excessive leniency under F_2.

Therefore, the theory predicts that greater privileges for elected officials should be paired with higher performance standards and more frequent turnover. Moreover, the result that π/ϕ cannot be too low under F_2 sheds light on the striking disparity between the compensation of politicians and that of private-sector executives. Whereas top executives often receive a significant share of a firm's profits, the salaries of senior politicians remain modest relative to the size of the economy. This discrepancy is surprising given the high stakes involved in political decision-making. However, our model indicates that to incentivize politicians to act on their private information, it is not optimal for them to be *excessively* concerned with keeping their job.

As seen in Figure 1, when $\alpha < \hat{\alpha}$, there exists a non-empty range of values for $\frac{\pi}{\phi}$ where im-

[25] That is, when $\frac{\pi}{\phi} \geq \max\left\{\frac{\alpha\beta}{1-\alpha\beta}, \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}\right\}$.

[26] As $\frac{\beta}{1-(p_E\alpha+(1-p_E))\beta}$ decreases in p_E, deviations to M in state w_E become more attractive when p_E is small.

[27] That is, when $\frac{\pi}{\phi} \leq \min\left\{\frac{\alpha\beta}{1-\alpha\beta}, \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}\right\}$.

plementing the voter's ideal program is not an equilibrium, but –either– increasing or decreasing $\frac{\pi}{\phi}$ makes it one. Corollary 1 establishes this result.

Corollary 1 *There exists a unique* $\hat{\alpha} \in \left(\frac{1}{2}, 1\right)$ *such that always implementing the voters' preferred policy can be sustained in equilibrium if and only if:*

- $\frac{\pi}{\phi} \in [0, \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}]$ *if* $\alpha \geq \hat{\alpha}$,
- $\frac{\pi}{\phi} \in [0, \frac{\alpha\beta}{1-\alpha\beta}] \cup [\frac{\beta(1-\alpha)}{1-(p_E\alpha+1-p_E)\beta}, \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}]$ *if* $\alpha < \hat{\alpha}$.[28]

Proof: See Appendix B.2. □

Intuitively, in the area sandwiched between areas F_1 and F_2, the incumbent's continuation utility under F_1 given the rather harsh punishment probability is not high enough to induce good behavior in state w_D, while, under F_2, the incumbent opportunistically chooses M in all states, including in state w_E, to stay in power. This is a case where the politician's salary, ego-rents, and perks are either too low or, more surprisingly, too high to properly incentivize them. To restore optimality, for instance, the politician's salary would have to substantially increase so as to decentralize the first best with the harsh strategy F_1, or, alternatively, to substantially decrease so that staying in power at all costs is no longer attractive under the forgiving strategy F_2.

This sandwiched area results from restricting attention to pure strategies that are conditioned only on the most recent period of play. While we view these restrictions as realistic, we stress that relaxing them expands the set of parameters under which the first-best can be achieved (see Online Appendix O.2). If voters could collectively randomize over retention decisions, they could fully convexify the set of parameters that sustain the first-best.[29] Even without randomization, voters could alternate deterministically between F_1 and F_2 across periods. This switching may partially expand the implementable region by averaging the incentives associated with each strategy. Even so, this only coarsely convexifies the set and remains less effective than mixing.

[28] Since $\frac{\alpha\beta}{1-\alpha\beta} < \frac{\beta(1-\alpha)}{1-(p_E\alpha+1-p_E)\beta}$, the two intervals in the second case are disjoint.

[29] Consider the following mixed retention strategy: if the signal is bad and the policy is E, replace the incumbent; if the signal is bad and the policy is M, reelect the incumbent with probability r. When $r = 0$, this coincides with F_1; when $r = 1$, with F_2. The idea is to choose r high enough to deter deviation in state w_D by raising continuation values, but low enough to prevent pandering in state w_E.

5.2 Electing Outsiders on the Equilibrium Path

If neither condition *(f1)* nor *(f2)* of Proposition 3 is satisfied, there is no equilibrium where the voter receives her preferred policy in every period. We show in this section that occasionally voting for an outsider who is harmful to the insider's policy interests can be desirable for the voters.

As previously discussed, we restrict attention to the mildest form of outsider punishment, i.e., replacing by an outsider for exactly one period before returning the helms of government to another insider. Analyzing this limited form of punishment allows us to derive the key insights about the benefits and costs of electing outsiders. Furthermore, when this punishment effectively deters insiders from misbehaving, the voter always prefers the shortest punishment duration, as her objective is to minimize the frequency at which outsiders hold power. In this case, the one-term assumption is not restrictive.[30]

Roadmap. To show when the voter's preferred equilibrium involves the election of outsiders, we proceed in three steps. First, we establish that, if outsiders are ever elected in equilibrium, this can only occur under a specific set of circumstances (Subsection 5.2.1). Second, we show when the outsider threat is effective in disciplining insiders (Subsection 5.2.2). Finally, we show when the voter is willing to carry out their threat (Subsection 5.2.3). Together, these steps identify when outsider punishment arises in the voter's preferred equilibrium.

5.2.1 Threats as differentiated punishments

Recall that a good signal implies that the incumbent chose the policy preferred by the representative voter. In equilibrium, an insider incumbent is then reelected. Hence, any optimal election strategy differs only in what happens after a bad signal. Let P represent the election of an outsider, A the election of another insider, and R reelection of the incumbent.

There are five conceivable strategies (in case of a bad signal) that involve replacing an insider incumbent with an outsider.[31] Lemma 1 shows that only two of these strategies can be part of an equilibrium, as each of the other ones are dominated by (at least) one of the first two.

[30] If electing a one-term outsider is insufficient to discipline insiders, the representative voter may consider reelecting the outsider for multiple periods. Given our assumptions, the only possible long-term punishment is a grim trigger strategy, where the voter permanently replaces insiders with outsiders.

[31] The five strategies are $(E, M) \to (P, A)$, $(E, M) \to (P, R)$, $(E, M) \to (P, P)$, $(E, M) \to (A, P)$, $(E, M) \to (R, P)$.

Lemma 1 *The only electoral strategies that result in the election of an outsider and may be used to discipline insiders in an equilibrium are* $S_1 : (E, M) \to (P, A)$ *or* $S_2 : (E, M) \to (P, R)$.

Proof: See Appendix C.1. □

Under S_1, a bad signal leads to replacement by another insider if the implemented policy was M, and replacement by an outsider if the implemented policy was E. In contrast, under S_2, the incumbent is only replaced by an outsider after choosing E and the voter receiving a bad signal, while an incumbent who chose M is always reelected. Note that conditioning retention on policy choice—as in S_1 and S_2—only matters under imperfect monitoring. When information is perfect ($\alpha = 1$), reelection decisions depend solely on whether the correct policy was implemented.

A surprising result is that no other strategy than these two can arise in equilibrium. In particular, this is true for the harshest form of punishment, $(E, M) \to (P, P)$, where any unsuccessful incumbent is replaced by an outsider. Intuitively, the voter does not want to punish an incumbent who chose M, but failed to generate a good signal. Replacing him with an outsider would just reduce the continuation utility of a well-behaved incumbent without generating any benefit.

Lemma 1 thus pushes back against a common view of the populist vote as an unmitigated rejection of establishment politicians. The electoral strategies detailed in Lemma 1 are more nuanced in that they eliminate an incumbent only if he has failed after choosing a suspicious policy.

Illustration: The French yellow vests movement illustrates this idea well. The movement, which consisted of a series of populist and grassroots weekly protests over several months, was triggered by a *small* gas environmental-tax increase in November 2018.[32] In contrast, a *larger* increases in fuel prices following the start of the war in Ukraine in 2022 did not trigger similar unrest.[33] The 2018 environmental-tax increase was seen as unfair and pro-elite (because it is only the working and middle classes in rural and suburban areas who drive to work long distance), while the later price hike was viewed as beyond government control. In this case, simpler and pro-masses agenda found more traction when the policy seemed pro-elite, rather than with an external economic shock. In

[32] At the start of the movement, more than 70 percent of the French population supported the protests (https://www.ifop.com/wp-content/uploads/2019/11/116084-Rapport-CN-SR-N80.pdf) Indeed, although over-represented among the far left and far right, the yellow vests were supported by a large proportion of French people, across the political spectrum (Douenne and Fabre, 2022).

[33] In November 2018, fuel prices were below 1.5 euros per liter, with planned increases of 3 cents for gasoline and 6 cents for diesel. After the war in Ukraine started, prices rose by more than 7 cents for gasoline and 14 cents for diesel, reaching nearly 1.9 euros, but did not provoke widespread protest.

response of the yellow vests, the government increased the minimum monthly salary by 100 euros, and employers were encouraged to pay a year-end bonus that would be tax-exempt.

5.2.2 The outsider threat

In this subsection, we analyze for which parameters S_1 and S_2 incentivize insiders to implement the voter's ideal policy, and compare it to F_1 and F_2. In contrast to the baseline model, we show that with imperfect information, the threat of replacement by another elite politician may, in some cases, provide strictly better incentives to insiders than the outsider threat. Recall that both S_1 and S_2 involve replacing the insider incumbent only following a bad signal; what differs is whether replacement depends on the policy choice. Under S_1, the incumbent is never reelected after a bad signal: he is replaced by an outsider if he chose E, and by another insider if he chose M. Under S_2, the incumbent is replaced by an outsider only if he chose E; otherwise, he is reelected.

The advantage of the strategy S_2 over strategy S_1 is analogous to the benefit of F_2 over F_1: the (ex-ante) continuation utility of a well-behaved incumbent is greater because the incumbent is re-elected after the w_M and w_D states with probability 1, rather than with probability α. However, similar to the case of F_2, S_2 might prompt the incumbent to choose M even in state w_E to ensure reelection. Moreover, this issue is further intensified compared to F_2 due to the harsher punishment.

Lemma 2 summarizes when the outsider threat generates good behavior by insiders.

Lemma 2 *Always implementing the voter's ideal policy is the best response for insiders to others' insiders following the voter's ideal policy, and*

1. *to voter playing S_1 if and only if*

$$\frac{\pi}{\phi} \leq \frac{\alpha\beta}{1-\alpha\beta}\frac{1+p_E\beta(1-\alpha)}{1-p_E\alpha\beta} \tag{s1}$$

2. *and, to voter playing S_2 if and only if*

$$\frac{1+p_E\beta(1-\alpha)}{1+p_E\beta(1-\alpha)-\beta}\,(1-\alpha)\beta \;\leq \frac{\pi}{\phi} \leq \frac{1+p_E\beta(1-\alpha)}{1+p_E\beta(1-\alpha)-\beta}\,\frac{\beta}{1-p_E\beta\alpha} \tag{s2}$$

Proof: See Appendix C.2. □

Since $\frac{1+p_E\beta(1-\alpha)}{1-p_E\alpha\beta} = 1 + \frac{p_E\beta}{1-p_E\alpha\beta} \geq 1$, the right-hand side of the condition *(s1)* of Lemma 2 is larger than the right-hand side of the condition *(f1)* of Proposition 3, implying that strategy S_1 is less susceptible to deviations in state w_D than F_1. Intuitively, the incumbent's punishment is more severe under S_1, which decreases the expected continuation policy payoff after a deviation. Similarly, observe also that the right hand side of *(s2)* is larger than the right hand side of *(f2)* in Proposition 3.[34] Thus, if the incumbent's temptation to implement policy E in state w_D is slightly too large under F_2 (because $\frac{\phi}{\pi}$ is relatively small), then the outsider threat may prompt him to behave. Misbehavior leads to outsider replacement, and thus to a reduced probability that the incumbent's preferred policy is implemented in the next period.

On the other hand, there are cases where the outsider threat (S_1 or S_2) does not enhance the insider's incentives compared to F_2, as summarized in the following corollary.

Corollary 2 *There exists a non-empty set of parameters for which efficient political alternation (F_2) sustains an equilibrium in which all insiders choose the voter's preferred policy, while the outsider threat (S_1 nor S_2) cannot.*

Proof: See Appendix C.3. □

The proof follows immediately from comparing conditions *(f2)*, with conditions *(s1)* and *(s2)*. Intuitively, the outsider threat is not helpful if the critical problem under F_2 is that the incumbent is too tempted to ensure reelection in state w_E by pandering to voters and implementing policy M. The threat of replacement by an outsider exacerbates this pandering problem because the insider incumbent is now even more worried about losing his job after implementing the correct policy in state w_E.

Another potential downside of the outsider threat is that the incumbent might seek to avoid punishment, i.e., the election of an outsider, at all costs. In particular, one can wonder whether it is now harder to incentivize the incumbent to implement E in state w_E: The insider could always avoid being replaced by an outsider by never playing E. It turns out that this is not the case because choosing M over E when the state of the world is w_E is just as bad as being governed by an outsider.[35]

[34] Since $\frac{1+p_E\beta(1-\alpha)}{1-p_E\beta\alpha}$ is decreasing in p_E, α, and β, the difference between F_2 and S_2 and the difference between F_1 and S_1 are smaller when these variables are larger.

[35] Intuitively, if an insider incumbent in state w_E chooses E, he receives a positive policy payoff in the current period, and has some chance of being reelected, but also runs the risk of being replaced by an outsider. Choosing M instead removes the risk of outsider replacement, but at the cost of losing any chance of reelection. Furthermore, the incumbent

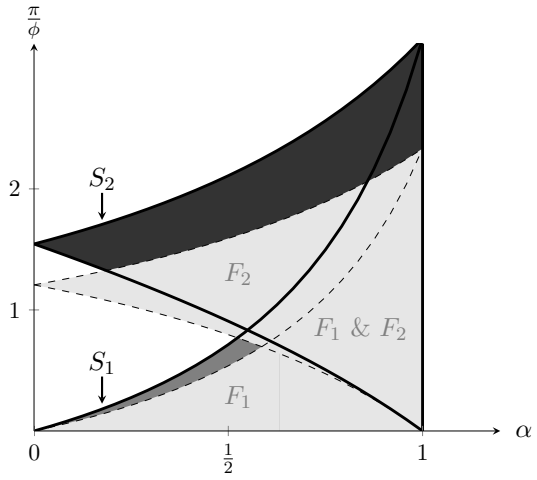

Figure 2: The light gray area (as in Figure 1) shows parameters for which F_1 or F_2 implement the voter's ideal program. The darker gray and black areas, respectively, show parameters for which strategies S_1 and S_2, respectively, deter misbehavior among insiders, while F_1 and F_2 do not. As in Figure 1, $p_E = 0.4$ and $\beta = 0.7$.

Figure 2 illustrates our results. It shows that it is possible for the outsider threat to induce the incumbent to behave well in cases where the electoral strategies F_1 and F_2 fail to do so. The light gray area is one in which the threat of alternation and replacement by another insider is sufficient to induce the incumbent to always choose the voter's preferred policy; this is the same area as the Figure 1. The dark gray and black areas show the parameters for which the outsider threat (S_1 or S_2) deters opportunistic behavior by insiders, while F_1 or F_2 do not. Figure 2 also illustrates Corollary 2: selecting an outsider instead of another insider is not always the most effective way to incentivize insiders. This is shown by the light gray area for $\alpha \leq 0.5$, which falls outside both S_1 (dark gray) and S_2 (black). Finally, in the white areas of Figure 2, none of the considered strategies incentivize good behavior among insiders.

Last, recall that Proposition 2 established that insiders can credibly promise to govern like an outsider for one period (and then being reelected on an insider mandate). This result still holds under imperfect information,[36] and in fact, commit to an outsider program is relatively easier in this setting. Since reelection depends only on the implemented policy (M) rather than the realization

loses the policy benefit in the present period for sure, while having an insider rather than outsider successor is only beneficial if tomorrow's state is w_E or w_D.

[36]Recall that we focus here on equilibria in which strategies depend only on the most recent period. However, implementing "fake populism" to incentivize an insider to behave like an outsider requires voters to recall not only the last period but also the one preceding it.

of the signal, an insider elected on a mandate to behave like an outsider can ensure reelection with probability one. Paradoxically, this means an insider is more tempted to deviate when ruling as a well-behaved insider than following an outsider mandate, as the latter guarantees staying in power.

The fact that in some cases S_1 or S_2 can succeed to incentivize insiders, while F_1 and F_2 fail to do so, does not necessarily mean that voters should adopt these strategies, as they require to sometimes elect an outsider, which is costly for them. We analyze this question in the next section.

5.2.3 From threats to votes

The following proposition builds on the previous two lemmata, and shows when the voter's preferred equilibrium features the election of outsiders on the equilibrium path.

Proposition 4 *The voter's preferred equilibrium involves voters following S_1 or S_2, and insiders implementing the voter's preferred policy in every period if and only if:*

1. *The representative voter is willing to elect an outsider in response to a deviation:*

$$p_E \left(1 - \frac{\beta}{1 + p_E \beta (1 - \alpha)}\right) \leq p_D. \tag{4}$$

2. *The outsider threat is necessary and sufficient: (s1) or (s2) hold, while (f1) and (f2) fail.*

In this case, an insider is replaced by an outsider with probability $(1 - \alpha)p_E$ following a correct but "'suspicious" policy choice in state w_E.

Proof: See Appendix C.4. □

Proposition 4 is a central result of our paper. It characterizes the conditions under which a rational voter without intrinsic preference for outsider platforms finds it optimal to support them on the equilibrium path as a means to discipline insiders. In the resulting equilibrium, insiders choose the voter's preferred policy in every period but are occasionally replaced by outsiders when two events happen simultaneously: dissatisfaction with the incumbent's performance (a bad signal) and the suspicion that policy choice favored elite interests.

Condition (4) tells us when the outsider threat is incentive-compatible for the voter. While, in a punishment subgame under complete information, the voter expects to carry out the punishment

only once,[37] with imperfect information, the threat must be implemented repeatedly whenever a bad signal follows a suspicious policy. This results in higher long-term costs. Hence, the noisier the information, the less appealing the punishment strategy becomes (the left-hand side of (4) decreases in α). In particular, while any sufficiently patient voter ($\beta \to 1$) is willing to implement the threat in the complete information case, this is not the case here, even when $\beta = 1$. Lastly, the outsider threat is more attractive when p_D is large and p_E is small, as insiders require incentives in the disagreement state and the cost of mistakenly replacing a well-aligned insider is lower.

Proposition 4 follows from Lemmata 1 and 2, and a characterization of when the voter is willing to implement his punishment strategy (Condition 4), but requires additional arguments. Since the outsider threat occasionally needs to be implemented, incentivizing via outsiders does not constitute a first-best equilibrium. Thus, we also need to check whether using any alternative strategy not previously considered in our analysis could improve the voters' welfare. For example, a voter could elect an insider who promises to implement policy E for one period as a way to punish insiders. In the proof of Proposition 4, we show that such deviations never improve the voter's welfare.

How often outsiders are in power. Proposition 4 also shows that the probability that an insider is replaced by an outsider is

$$(1-\alpha)p_E.$$

This provides two insights. First, a seemingly paradoxical result is that the probability of voting for an outsider does *not* depend on the likelihood of the disagreement state (p_D). While p_D determines how much insiders need to be disciplined, it does not directly affect how often outsiders are elected. On the equilibrium path, insiders are well-behaved but still face punishment when they implement E and the signal is bad. As such, an outsider vote follows insiders choosing E when it is optimal—both for themselves and for the voter—but then facing misplaced suspicion due to a bad signal, which leads to outsider replacement.

Second, lower information accuracy (α) increases the frequency of outsider replacements. This complements our earlier results in Propositions 3 and Lemma 2, where we showed that poor information quality weakens first-best retention strategies F_1 and F_2, making outsider-based incentives more appealing. Here, we further show that low α not only makes outsider threats more attractive

[37] This is because incumbents, on the equilibrium path, never misbehave under complete information.

but also increases their actual implementation rate. This aligns with empirical findings suggesting that misinformation and reduced information quality contribute to outsider electoral success.

6 Conclusion

In this paper, we have developed a model in which voters elect outsiders to discipline elite politicians. They do so only after the chosen policy is both perceived as a failure and as benefiting the elites. For voters, electing an outsider is a rational investment: despite lowering expected utility in the next period, it deters elites from acting against their interests in the future. Our model also clarifies how information shapes outsider voting. A more opaque environment makes voters both more inclined to support outsiders and to do so more frequently.

In contrast to our model, the most common explanation for the outsider vote (both in the popular press and in the academic literature) is that voters turn to them when frustrated by the failure of traditional politicians to improve their well-being. Our theory offers a complementary explanation that is based on rational and strategic voting by a Downsian median voter, and provides insights about voters' behavior and the often very volatile electoral success of populist parties. For example, in Austria, the right-wing FPÖ in the last three federal elections, received 26% (2017, in government); 16% (2019, in opposition); and 29% (2024, in opposition).

Our model can be understood as a short-hand in which "sincere" voters (who are convinced that outsiders candidates will always implement a better policy than insiders) coexist with rational voters who sometimes, but not always, support outsiders for the reasons outlined in our model.[38] The presence of strategic voters explains sudden and dramatic swings in outsider support, as rational voters may collectively back an outsider in one election and an establishment candidate in another, without changing their political preferences.[39]

[38] Observe that, as long as the former group does not constitute a majority, rational voters remain decisive in election outcomes.

[39] In contrast, if all support for outsiders were purely sincere, we would expect their electoral success to remain relatively constant across elections, or at least more closely correlated with the material success of voters.

References

Acemoglu, D., Egorov, G., and Sonin, K. (2013). A political theory of populism. *The Quarterly Journal of Economics*, 128(2):771–805.

Agranov, M., Eilat, R., and Sonin, K. (2020). A political model of trust.

Allcott, H. and Gentzkow, M. (2017). Social media and fake news in the 2016 election. *Journal of economic perspectives*, 31(2):211–236.

Anesi, V. and Buisseret, P. (2022). Making elections work: Accountability with selection and control. *American Economic Journal: Microeconomics*, 14(4):616–644.

Barro, R. J. (1973). The control of politicians: an economic model. *Public choice*, pages 19–42.

Bellodi, L., Morelli, M., Nicolo, A., and Roberti, P. (2023). The shift to commitment politics and populism: Theory and evidence. *BAFFI CAREFIN Centre Research Paper*, (204).

Berman, S. (2021). The causes of populism in the west. *Annual Review of Political Science*, 24:71–88.

Bernhardt, D., Câmara, O., and Squintani, F. (2011). Competence and ideology. *The Review of Economic Studies*, 78(2):487–522.

Bernhardt, D., Campuzano, L., Squintani, F., and Câmara, O. (2009). On the benefits of party competition. *Games and Economic Behavior*, 66(2):685–707.

Bernhardt, D., Krasa, S., and Shadmehr, M. (2022). Demagogues and the economic fragility of democracies. *American Economic Review*, 112(10):3331–66.

Bolton, P. and Harris, C. (1999). Strategic experimentation. *Econometrica*, 67(2):349–374.

Canes-Wrone, B., Herron, M. C., and Shotts, K. W. (2001). Leadership and pandering: A theory of executive policymaking. *American Journal of Political Science*, pages 532–550.

Carnes, N. and Lupu, N. (2023). The economic backgrounds of politicians. *Annual Review of Political Science*, 26.

Di Tella, R. and Rotemberg, J. J. (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.

Douenne, T. and Fabre, A. (2022). Yellow vests, pessimistic beliefs, and carbon tax aversion. *American Economic Journal: Economic Policy*, 14(1):81–110.

Duggan, J. (2000). Repeated elections with asymmetric information. *Economics & Politics*, 12(2):109–135.

Duggan, J. and Martinelli, C. (2017). The political economy of dynamic elections: Accountability, commitment, and responsiveness. *Journal of Economic Literature*, 55(3):916–984.

Edwards, S. (2019). On latin american populism, and its echoes around the world. *Journal of Economic Perspectives*, 33(4):76–99.

Ferejohn, J. (1986). Incumbent performance and electoral control. *Public choice*, pages 5–25.

Fox, J. and Stephenson, M. C. (2011). Judicial review as a response to political posturing. *American Political Science Review*, 105(2):397–414.

Funke, M., Schularick, M., and Trebesch, C. (2021). The cost of populism: Evidence from history. Technical report, ECONtribute Policy Brief.

Gailmard, L. (2022). Electoral accountability and political competence. *Journal of Theoretical Politics*, 34(2):236–261.

Gilens, M. (2012). Affluence and influence. In *Affluence and Influence*. Princeton University Press.

Gratton, G. and Lee, B. E. (2023). Drain the swamp: A theory of anti-elite populism.

Green, E. J. and Porter, R. H. (1984). Noncooperative collusion under imperfect price information. *Econometrica: Journal of the Econometric Society*, pages 87–100.

Guriev, S., Melnikov, N., and Zhuravskaya, E. (2021). 3g internet and confidence in government. *The Quarterly Journal of Economics*, 136(4):2533–2613.

Guriev, S. and Papaioannou, E. (2022). The political economy of populism. *Journal of Economic Literature*, 60(3):753–832.

Keller, G., Rady, S., and Cripps, M. (2005). Strategic experimentation with exponential bandits. *Econometrica*, 73(1):39–68.

Kselman, D. and Niou, E. (2011). Protest voting in plurality elections: a theory of voter signaling. *Public Choice*, 148:395–418.

Levy, G., Razin, R., and Young, A. (2022). Misspecified politics and the recurrence of populism. *American Economic Review*, 112(3):928–962.

Mailath, G. J. and Samuelson, L. (2006). *Repeated games and reputations: long-run relationships*. Oxford university press.

Maskin, E. and Tirole, J. (2004). The politician and the judge: Accountability in government. *American Economic Review*, 94(4):1034–1054.

Mizuno, N. and Okazawa, R. (2022). Why do voters elect less qualified candidates? *Journal of Theoretical Politics*, 34(3):443–477.

Morelli, M. and Van Weelden, R. (2013). Ideology and information in policymaking. *Journal of Theoretical Politics*, 25(3):412–439.

Mudde, C. and Kaltwasser, C. R. (2017). *Populism: A very short introduction*. Oxford University Press.

Müller, J.-W. (2017). *What is populism?* Penguin UK.

Myatt, D. P. (2017). A theory of protest voting. *The Economic Journal*, 127(603):1527–1567.

Rodrik, D. (2021). Why does globalization fuel populism? economics, culture, and the rise of right-wing populism. *Annual Review of Economics*, 13:133–170.

Van Weelden, R. (2013). Candidates, credibility, and re-election incentives. *Review of Economic Studies*, 80(4):1622–1651.

Van Weelden, R. (2015). The welfare implications of electoral polarization. *Social Choice and Welfare*, 45:653–686.

Appendix

A Full Information

A.1 Proof of Proposition 1

Point 1: If state w_D is more likely than state w_E (i.e., if $p_D > p_E$), poorly incentivized insiders are more costly to voters than outsiders. In this case, the threat is not costly to implement. In contrast, when $p_E > p_D$, the voter prefers a misbehaving insider to an outsider. If the voter carries out the threat, she receives her preferred policy with probability $1-p_E$ for one period, and thereafter continues to receive her preferred policy in all future periods, yielding a total utility of

$$\pi\left((1-p_E)+\frac{\beta}{1-\beta}\right).$$

Conversely, if she tolerates the self-serving insider, she will not receive her preferred policy in every future disagreement state, resulting in a utility of $\pi\frac{1-p_D}{1-\beta}$. Which leads to

$$p_E\,(1-\beta) \leq p_D. \tag{5}$$

Point 2: We start by noting that the utility of a well-behaved incumbent remains the same as in the first-best scenario, U^*. What changes is the incumbent's utility from a one-shot deviation (while holding other politicians' strategies constant). He receives both policy payoff and perks today; tomorrow, he is replaced by an outsider, and the day after tomorrow, the outsider is replaced by another insider:

$$U_p^d(w_D) = \phi + \pi + \beta(\underbrace{\pi p_M + 0}_{\text{an outsider}}) + \underbrace{\frac{\beta^2}{1-\beta}(1-p_D)\pi}_{\text{another insider}}$$

Well-behaving is incentive-compatible for insider politicians if and only if $U^*(w_D) \geq U_p^d(w_D)$. Recall that the payoff $U^*(w_D)$ is given by:

$$U^*(w_D) = \phi + 0 + \beta U^* = \phi + \beta(\frac{\phi}{1-\beta} + \pi\frac{1-p_D}{1-\beta}) \ .$$

Replacing the payoffs by their expression, the inequality $U^*(w_D) \geq U_p^d(w_D)$ becomes:

$$\beta(\frac{\phi}{1-\beta}+\pi\frac{1-p_D}{1-\beta})\geq\pi+\pi p_M\beta+\beta^2\pi\frac{1-p_D}{1-\beta}$$

which simplifies to

$$\frac{\pi}{\phi}\leq\frac{\beta}{(1-\beta)(1-\beta p_E)}.$$

B First-Best with Imperfect Policy Evaluation

Notation. We denote the utility of an insider's incumbent if every insider – including himself – always selects the preferred policy of the voter, by

$$U^*(\sigma)=\phi\lambda_\phi(\sigma)+\pi\lambda_\pi(\sigma) \tag{6}$$

where $\lambda_k(\sigma)$, $k\in\{\phi,\pi\}$, denotes how the politician's payoff depends on the electoral strategy σ used by the voter. Let $U^*(\sigma|w)$ denote the same object, but conditional on the realization of the state of the world w in the current period. Finally, we denote by $U^d(\sigma|w)$ the continuation utility of a deviating incumbent in state w, i.e., one who selects his best one-shot deviation, given the state w.

B.1 Proof of Proposition 3

Preamble: For the voter's ideal policy to always be implemented, each incumbent must be well-incentivized. Given our assumptions on information, a good signal implies that the voter's preferred policy has been implemented. In this case, the incumbent should (and in equilibrium, will) be reelected. The relevant decision for the voter, therefore, concerns what to do when the signal is bad. The voter has three possible retention strategies in response to a *bad signal*:

1. $F_1:(M,E)\to(A,A)$ – Replace the incumbent regardless of policy choice.
2. $F_2:(M,E)\to(R,A)$ – Reelect if the incumbent chose M, but replace if they chose E.
3. $(M,E)\to(A,R)$ – Replace if the incumbent chose M, but reelect if they chose E.

where R (Reelect) means the incumbent remains in office, and A (Alternation) means the incumbent is replaced by another insider. The third strategy, $(M,E)\to(A,R)$, is not incentive-compatible.

If insiders know that choosing E still results in reelection, then in disagreement states w_D, they will always choose E, as they receive their preferred policy with no electoral consequences. This undermines the voter's ability to enforce the choice of M in w_D.

As a result, the only viable retention strategies are F_1 and F_2.

Point *(f1)*: We start by considering an equilibrium in which F_1 is played and all incumbents implement the voter's preferred policy. By virtue of equation (6), before observing the state of the world, the discounted expected utility of a well-behaved incumbent (i.e., one who implements E in state w_E, and M otherwise) is $U^*(F_1) = \phi\lambda_\phi(F_1) + \pi\lambda_\pi(F_1)$. Because a well-behaved incumbent is reelected with probability α and has a discount factor β, we have that

$$\lambda_\phi(F_1) = \sum_{i=0}^{\infty}(\alpha\beta)^i = \frac{1}{1-\alpha\beta}. \tag{7}$$

Furthermore, because all incumbents behave in the same way, the probability that the current incumbent's preferred policy is implemented in any period solely depends on the realization of the state of the world, independent of who is in power. Specifically, in any period, a politician gets his favorite policy with probability $1 - p_D$. Thus, $\lambda_\pi(F_1) = \sum_{i=0}^{\infty}(1-p_D)(\beta)^i = \frac{1-p_D}{1-\beta}$. The incumbent's discounted expected utility is therefore $U^*(F_1) = \phi\frac{1}{1-\alpha\beta} + \pi\frac{1-p_D}{1-\beta}$. To check whether the implementation of the voter's ideal program is incentive-compatible under F_1, we need to compute the incumbent's payoff when deviating, holding constant the other politicians' strategies.

First, observe that the incumbent has no incentive to misbehave in states w_E or w_M under election strategy F_1, as his policy preferences are aligned with the voter's, and choosing the correct policy also increases the incumbent's reelection probability.

Second, if the state of the world is w_D, and he is well-behaved (i.e., chooses policy M), his continuation utility is $U^*(F_1|w_D) = \phi + 0 + \alpha\beta U^*(F_1) + (1-\alpha)\beta\pi\lambda_\pi(F_1)$, where the last term captures the discounted expected payoff of not being in power.

In contrast, if the incumbent deviates and chooses E, he gets policy utility π in the present period, but is not reelected. Therefore, this deviation generates a continuation utility of $U^d(F_1|w_D) = \phi + \pi + \beta\pi\lambda_\pi(F_1)$. It is optimal for the incumbent to behave well if and only if $U^*(F_1|w_D) \geq U^d(F_1|w_D)$,[40]

[40] For the if-direction, observe that an incumbent can only deviate in one period because he is replaced with probability 1 after a deviation.

which, upon replacing $\lambda_\phi(F_1)$ by its value from equation (7) and simplifying, is equivalent to

$$\frac{\pi}{\phi} \leq \frac{\alpha\beta}{1-\alpha\beta}. \tag{8}$$

Point *(f2)*: Strategy F_2 differs from F_1 in that the incumbent is also reelected if he chose M, even if the signal is bad. Intuitively, this strategy has advantages and disadvantages for incentivizing the incumbent relative to F_1. The advantage is that it affords a well-behaved incumbent a higher continuation utility because it reduces the risk of being replaced (i.e., if the state is w_D, a well-behaved incumbent is now guaranteed to be reelected), and this makes it more attractive for the incumbent to choose policy M in state w_D. On the other hand, it opens the door for some form of opportunistic faux populism, i.e., an incumbent might choose policy M in state w_E because it guarantees reelection, even though both voter and politician would be better off with policy E.

Observe first that, if F_2 works so that all politicians are well-behaved along the equilibrium path, the policy payoffs remain unchanged relative to F_1, i.e., $\lambda_\pi(F_2) = \lambda_\pi(F_1) = \frac{1-p_D}{1-\beta}$. In contrast, a well-behaved incumbent is reelected unless the state is w_E and the signal is incorrect, so that, relative to F_1, the ex-ante probability of reelection (i.e., before the state of the world is known) increases from α to $p_E\alpha + (1-p_E)$.

Under F_2, the discounted expected payoff from ego rents is[41]

$$\lambda_\phi(F_2) = \frac{1}{1-(p_E\alpha + 1 - p_E)\beta}. \tag{9}$$

The ex-ante expected discounted payoff for the incumbent (before the state of the world is known) if he behaves well is $U^*(F_2)$ defined equation (6) with $\sigma = F_2$. When the state of the world is w_D and the incumbent is well-behaved, he is reelected with probability one, so that his expected discounted utility under election strategy F_2 can be written as $U^*(F_2|w_D) = \phi + \beta U^*(F_2)$. In contrast, if he deviates to E, his expected discounted payoff is $U^d(F_2|w_D) = \phi + \pi + \beta\pi\lambda_\pi(F_2)$. Deviating is not attractive if and only if $U^*(F_2|w_D) \geq U^d(F_2|w_D)$, which is equivalent to $\beta\phi\lambda_\phi(F_2) \geq \pi$. In other words, and rewritting this condition, the incumbent is well-behaved in the disagreement state w_D with strategy F_2 if and only if:

[41] Note that $\lambda_\phi(F_2) > \lambda_\phi(F_1) = \frac{1}{1-\alpha\beta}$.

$$\frac{\pi}{\phi} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}. \tag{10}$$

Since $\frac{\alpha\beta}{1-\alpha\beta} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$, condition (10) is weaker than condition (8) for strategy F_1. This captures the advantage of punishing the incumbent less often. A well-behaved politician can expect to stay in office for longer, and thus has a higher continuation utility, giving him a stronger incentive to choose M when the state of the world is w_D.

However, we now have to check, in addition, that an incumbent in state E does not want to deviate to play M. When the state of world is w_E, if all politicians are well-behaved, the incumbent's expected utility is $U^*(F_2|w_E) = \phi + \pi + \alpha\beta U^*(F_2) + (1-\alpha)\beta\lambda_\pi(F_2)\pi$. If the incumbent deviates, only for one period, by choosing M, he ensures his reelection in the next period and gets utility $U^d(F_2|w_E) = \phi + \beta U^*(F_2)$. This one-shot deviation is unattractive for the incumbent if and only if $U^*(F_2|w_E) \geq U^d(F_2|w_E)$, hence if

$$\frac{\beta(1-\alpha)}{1-(p_E\alpha+1-p_E)\beta} \leq \frac{\pi}{\phi}. \tag{11}$$

Observe that the left-hand side of condition (11) is always smaller than the right-hand side of condition (10), so that there are always some values of π/ϕ such both condition (10) and condition (11) are satisfied.

B.2 Proof of Corollary 1

Let's first recap the results. All the politicians are well-behaved constitutes an equilibrium if and only if

$$\frac{\alpha\beta}{1-\alpha\beta} > \frac{\pi}{\phi},$$

or,

$$\frac{(1-\alpha)\beta}{1-(p_E\alpha+(1-p_E))\beta} \leq \frac{\pi}{\phi} \leq \frac{\beta}{1-(p_E\alpha+(1-p_E))\beta}$$

This is point 1 and 2 in Proposition 3. It is next easy to check that, $\frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$ since $1-\alpha \leq 1$. It is also easy to check that $\frac{\alpha\beta}{1-\alpha\beta} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$ since $\alpha\beta p_E \leq 1$. Let

$$a_1 = min\left\{\frac{\alpha\beta}{1-\alpha\beta}, \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}\right\}$$

$$a_2 = max\left\{\frac{\alpha\beta}{1-\alpha\beta}, \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}\right\}$$

If $\frac{\pi}{\phi} < a_1$ then only F_1 can achieve the first-best equilibrium. If $\frac{\pi}{\phi} > a_2$ then only F_2 can potentially achieve the first-best equilibrium. It must however be the case that $\frac{\pi}{\phi} < \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$ to avoid having the politician adopting a populist behavior. Now if $\frac{\pi}{\phi} \in (a_1, a_2)$ two cases can occur. If $\frac{\alpha\beta}{1-\alpha\beta} \leq \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}$ then $a_1 = \frac{\alpha\beta}{1-\alpha\beta}$ and $a_2 = \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}$ so that with $\frac{\pi}{\phi} \in (a_1, a_2)$ neither F_1, nor F_2 works: the first best cannot be implemented. If $\frac{\alpha\beta}{1-\alpha\beta} > \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}$ then $a_1 = \frac{1-\alpha}{1-(p_E\alpha+(1-p_E))\beta}$ and $a_2 = \frac{\alpha\beta}{1-\alpha\beta}$. Since $\frac{\alpha\beta}{1-\alpha\beta} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$ we deduce that $\frac{1-\alpha}{1-(p_E\alpha+(1-p_E))\beta} < \frac{\pi}{\phi} < \frac{\alpha}{1-\alpha\beta} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$ so that both F_1 and F_2 can implement the first-best.

Let $\hat{\alpha}$ be so that $\frac{\alpha}{1-\alpha\beta} = \frac{1-\alpha}{1-(p_E\alpha+(1-p_E))\beta}$ (i.e., so that $a_1 = a_2$). Solving for this second degree equation, one root is larger than 1, and the other one, which is our solution yields:

$$\hat{\alpha} = \frac{2+\beta p_E - \sqrt{4(1-\beta)+(\beta p_E)^2}}{2\beta(1+p_E)}. \tag{12}$$

It is easy to check that under our assumptions $\hat{\alpha} \in (\frac{1}{2}, 1)$. The condition $\frac{\alpha\beta}{1-\alpha\beta} \geq \frac{(1-\alpha)\beta}{1-\beta(p_E\alpha+1-p_E)}$ is equivalent to $\alpha \geq \hat{\alpha}$. The set of $\frac{\pi}{\phi}$ values over which F_1 and/or F_2 implements the first best is compact (i.e. they overlap) if and only if $\alpha \geq \hat{\alpha}$. Symmetrically, if $\alpha < \hat{\alpha}$ then there is an interval of values for π/ϕ for which neither F_1 nor F_2 work.

C Electing Outsiders on the Equilibrium Path

C.1 Proof of Lemma 1

There are five strategies, conditional on having an insider politician in power and receiving a bad signal, that entails voting for an outsider: S_1: $(E, M) \to (P, A)$, S_2: $(E, M) \to (P, R)$, from Lemma 1, and in addition, $S_3 : (E, M) \to (P, P)$, $S_4 : (E, M) \to (A, P)$ and $S_5 : (E, M) \to (R, P)$ where (as previously introduced) P represents the election of an outsider, A the election of another insider, and R reelection of the incumbent. First, note that S_5 is never incentive compatible because it would be optimal for the incumbent to choose E when the state of the world is w_D as it will secure reelection while implementing the preferred elite policy over the voter's one. More interestingly, in what follow we show that S_3 is dominated by S_1 in that whenever the former give good incentive

to the incumbent the latter does as well; and the latter involved voting for an outsider at a lower frequency. Then we are going to show that S_4 is dominated by F_1.

Step 1: S_3 ***is dominated by*** S_1**.** If the voter follows S_3 then only when the state of the world is w_D the incumbent has an incentive to misbehave (by choosing E over M). If the state of the world is w_D and the politician is well-behaved, he chooses M so that an outsider is then elected in case of negative signal. It implies that with S_3 he gets: $U^*(S_3|w_D) = \phi + 0 + \alpha\beta U^*(S_3) + (1-\alpha)\beta\pi(p_M + \beta\lambda_\pi(S_3))$, where $U^*(S_3)$ is defined equation (6) with $\sigma = S_3$. If however the incumbent deviates by playing E in this period, because an outsider is going to be elected next period with probability one, he gets: $U^d(S_3|w_D) = \phi + \pi + \beta\pi(p_M + \beta\lambda_\pi(S_3))$. The conjecture equilibrium (well-behaved insiders) is indeed an equilibrium if and only if $U^*(S_3|w_D) \geq U^d(S_3|w_D)$, which is equivalent to $\frac{\alpha\beta}{1-\alpha\beta}\phi + \beta\pi\alpha(\lambda_\pi(S_3)(1-\beta) - p_M) > \pi$. Let's compare this condition with condition (15), which is the condition for S_1. First, the second term is here multiplied by $\alpha < 1$. Second, since outsiders are elected more often with strategy S_3, $\lambda_\pi(S_3)$ is smaller than $\lambda_\pi(S_1)$. Hence, for any π which satisfy the condition (15) in S_1, it also satisfies the one in S_3: S_1 is more effective than S_3. Finally, the frequency at which the preferred voter's policy is obtained is higher in S_1 than in S_3.

Step 2: S_4 ***is dominated by*** F_1**.** As in step 1, we focus on the conflicting state of the world w_D. If the incumbent is well-behaved in state w_D, he gets $U^*(S_4|w_D) = \phi + 0 + \alpha\beta U^*(S_4) + (1-\alpha)\beta\pi(p_M + \beta\lambda_\pi(S_4))$, where $U^*(S_4)$ is defined equation (6) with $\sigma = S_4$. Observe that when he is well-behaved, he chooses M, and so an outsider will not be elected in case of negative signal. If, however, the incumbent deviates, he gets $U^d(S_4|w_D) = \phi + \pi + \beta\pi\lambda_\pi(S_4)$. The conjectured equilibrium is indeed an equilibrium if and only if $U^*(S_4|w_D) \geq U^d(S_4|w_D)$, which is equivalent to

$$\alpha\beta\phi\lambda_\phi(S_4) + \underbrace{\lambda_\pi(S_4)\pi\beta(1-\alpha)(\beta-1)}_{<0} \geq \pi \tag{13}$$

Note that, under S_4, a well-behaved incumbent is reelected with probability α. Thus, we have that $\lambda_\phi(S_4) = \lambda_\phi(F_1)$ defined in equation (7). Therefore if the condition (13) is satisfied, then the condition (8) is also satisfied. And F_1 restores the first best, while S_4 does not.

C.2 Proof of Lemma 2

Point *(s1):* We first consider election strategy S_1, where a bad signal always leads to the incumbent's replacement. Because a well-behaved incumbent is reelected with probability α as under F_1, we have $\lambda_\phi(S_1) = \lambda_\phi(F_1)$, as defined in equation (7). To calculate the incumbent's continuation utility from policy, we presume that all other insiders are well behaved so that, when an insider incumbent is in power, each insider gets their preferred policy with probability $1 - p_D$ in any given period. In contrast, when an outsider is in power, each insider gets their preferred policy only with probability p_M. We can therefore express $\lambda_\pi(S_1)$ recursively:

$$\lambda_\pi(S_1) = (1 - p_D) + \beta[\underbrace{\alpha\lambda_\pi(S_1)}_{\text{reelection}} + \underbrace{(1-\alpha)(1-p_E)\lambda_\pi(S_1)}_{\text{other insider}} + \underbrace{(1-\alpha)p_E(p_M + \beta\lambda_\pi(S_1))}_{\text{outsider}}]$$

Solving for $\lambda_\pi(S_1)$, we obtain:

$$\lambda_\pi(S_1) = \frac{1 - p_D + \beta(1-\alpha)p_E p_M}{(1-\beta)[1 + (1-\alpha)p_E\beta]}. \tag{14}$$

To check if the incumbent is willing to implement the voter's preferred policy in each state, we need to compute the best profitable deviation in state w_D, holding constant the other politicians' strategy. If he chooses M in state w_D, a bad signal leads to replacement by another insider. Thus, the expected discounted utility in equilibrium is $U^*(S_1|w_D) = \phi + 0 + \alpha\beta U^*(S_1) + (1-\alpha)\beta\pi\lambda_\pi(S_1)$, where $U^*(S_1)$ is defined by equation (6) for $\sigma = S_1$.

If, instead, the incumbent deviates to E in state w_D,[42] an outsider will be elected in the next period, before being replaced by another insider in the following period. Thus, the current incumbent's expected discounted payoff is $U^d(S_1|w_D) = \phi + \pi + \beta\pi(p_M + \beta\lambda_\pi(S_1))$. For an insider incumbent, choosing M in state w_D is optimal if and only if $U^*(S_1|w_D) \geq U^d(S_1|w_D)$, which is equivalent to:

$$\frac{\pi}{\phi} \leq \frac{\alpha\beta}{1-\alpha\beta} \frac{1 + p_E\beta(1-\alpha)}{1 - p_E\alpha\beta} \tag{15}$$

The only part missing from the proof is to check that if the state of the world is w_E, the incumbent is willing to play E. If he does, he gets: $U^*(S_1|w_E) = \phi + \pi + \alpha\beta U^*(S_1) + (1-\alpha)\beta\pi(p_M +$

[42] In principle, this deviation is for one period only, even though this does not matter as, following a deviation, the incumbent is not reelected anyway.

$\beta\lambda_\pi(S_1))$, where the last term captures the fact that, if the incumbent chooses E and the signal is bad, he is replaced by an outsider. If, instead, the incumbent deviates (in the present period only), then he gets $U^d(S_1|w_E) = \phi + 0 + \beta\pi\lambda_\pi(S_1)$, as he will lose the election with certainty when choosing M and will then be replaced by another insider. The deviation is not profitable if and only if $U^*(S_1|w_E) \geq U^d(S_1|w_E)$ which, after substituting $U^*(S_1)$ defined equation (6) for $\sigma = S_1$, is equivalent to $\pi + \phi\alpha\beta\lambda_\phi(S_1) + \beta\pi(1-\alpha)(p_M - \lambda_\pi(S_1)(1-\beta)) > 0$. Substituting $\lambda_\phi(S_1) = \lambda_\phi(F_1)$ defined equation (7), $\lambda_\pi(S_1)$ defined equation (14), and simplifying, we obtain: $\frac{\pi}{1+p_E\beta(1-\alpha)} + \frac{\phi\alpha\beta}{1-\alpha\beta} > 0$. This condition is always satisfied, even if π, the rent from ruling, is set to zero.

Point *(s2):* Observe first that on the equilibrium path the incumbent's expected policy payoff from playing S_2 does not change relative to S_1 because all insiders choose the same policy, and the probability that an outsider is in charge is exactly the same as under S_1 (as the outsider is only elected after policy E and a bad signal). Thus, we have $\lambda_\pi(S_2) = \lambda_\pi(S_1)$ defined equation (14). Second, the incumbent's continuation utility from holding office is the same as under F_2, because the incumbent is replaced in the same states. Thus we have $\lambda_\phi(S_2) = \lambda_\phi(F_2)$ defined equation (22). To check when this is an equilibrium, we need to compute the best profitable deviation, given the strategies of the other politicians. First, consider the state of the world w_D, where the incumbent may have an incentive to misbehave by choosing E over M. If he is well-behaved in state w_D, he gets: $U^*(S_2|w_D) = \phi + \beta U^*(S_2)$. If, instead, he chooses to play E, the continuation utility is $U^d(S_2|w_D) = \phi + \pi + \beta\pi(p_M + \beta\lambda_\pi(S_2))$. The incumbent has no incentive to deviate from S_2 if and only if $U^*(S_2|w_D) \geq U^d(S_2|w_D)$, which, after substituting $U^*(S_2)$ defined equation (6) for $\sigma = S_2$, is equivalent to $\beta\phi\lambda_\phi(S_2) + \beta\pi((1-\beta)\lambda_\pi(S_2) - p_M) \geq \pi$. Substituting $\lambda_\pi(S_2) = \lambda_\pi(S_1)$ from equation (7), $\lambda_\phi(S_2) = \lambda_\phi(F_2)$ from equation (22), and simplifying yields $\pi \leq \frac{\beta\phi}{1-(p_E\alpha+1-p_E)\beta} + \frac{\pi\beta p_E}{1+p_E\beta(1-\alpha)}$. This is equivalent to

$$\frac{\pi}{\phi} \leq \frac{1 - p_E\beta\alpha + p_E\beta}{1 - p_E\beta\alpha} \frac{\beta}{1-(p_E\alpha+1-p_E)\beta} \tag{16}$$

which is the RHS of equation *(s2)*.

Second, as with strategy F_2, there is a danger with strategy S_2 that the incumbent implements false populism in order to ensure reelection. Given the state of the world w_E, a well-behaved incumbent's continuation payoff is $U^*(S_2|w_E) = \phi + \pi + \alpha\beta U^*(S_2) + (1-\alpha)\beta\pi(p_M + \beta\lambda_\pi(S_2))$. If, instead, the incumbent deviates by choosing M in order to ensure reelection (only for one period), then the

expected continuation payoff is $U^d(S_2|w_E) = \phi + 0 + \beta U^*(S_2)$, where $U^*(S_2)$ is defined equation (6) for $\sigma = S_2$. The equilibrium strategy is at least as good as the one shot deviation if and only if $U^*(S_2|w_E) \geq U^d(S_2|w_E)$, which is equivalent to $\pi \geq \beta(1-\alpha)\phi\lambda_\phi(S_2)+\beta(1-\alpha)\pi(\lambda_\pi(S_2)(1-\beta)-p_M)$. Substituting $\lambda_\pi(S_2) = \lambda_\pi(S_1)$ from equation (7), $\lambda_\phi(S_2) = \lambda_\phi(F_2)$ from equation (22), and simplifying yields $\pi \geq \frac{(1-\alpha)\beta\phi}{1-\beta(1-p_E(1-\alpha))} + \pi\frac{\beta p_E(1-\alpha)}{1+p_E\beta(1-\alpha)}$. This can be rewritten as

$$\frac{\pi}{\phi} \geq \frac{(1-\alpha)\beta}{1-\beta(1-p_E(1-\alpha))}(1+p_E\beta(1-\alpha)) \tag{17}$$

which is the LHS of equation $(s2)$. Putting the two conditions yields:

$$\frac{1+p_E\beta(1-\alpha)}{1+p_E\beta(1-\alpha)-\beta}\,(1-\alpha)\beta \;\leq \frac{\pi}{\phi} \leq \frac{1+p_E\beta(1-\alpha)}{1+p_E\beta(1-\alpha)-\beta}\,\frac{\beta}{1-p_E\beta\alpha} \tag{18}$$

C.3 Proof of Corollary 2

We show that there exists a non-empty set of parameters for which F_2 sustains an equilibrium in which all insiders are well behaved, which requires that conditions (10) and (11) hold, while condition (15) and condition (18) do not hold so that S_1 and S_2 fail to incentivize them. This case occurs if and only if there exists C_1, with $1 < C_1 \leq \frac{1}{1-\alpha}$, such that:

$$\frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta} \leq \frac{\pi}{\phi} \leq C_1\frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta},$$

and

$$\frac{\alpha\beta}{1-\alpha\beta}\frac{1+p_E\beta(1-\alpha)}{1-p_E\alpha\beta} < \frac{\pi}{\phi}.$$

The first condition guarantees that F_2's conditions (10) and (11) are satisfied while S_2's condition (18) is not. The second condition guarantees that S_1's condition (15) is not satisfied. It is straightforward that, for any $C_1 > 1$, the set of $\frac{\pi}{\phi}$ for which the first condition is satisfied is non-empty. So the set of interest is non-empty if and only if $1 + p_E\beta(1-\alpha) \leq \frac{1}{1-\alpha}$, which is equivalent to $\beta p_E \leq \frac{\alpha}{(1-\alpha)^2}$, and the following condition is satisfied:

$$\frac{\alpha\beta}{1-\alpha\beta}\frac{1+p_E\beta(1-\alpha)}{1-p_E\alpha\beta} \leq \frac{(1-\alpha)\beta}{1-(p_E\alpha+1-p_E)\beta}.$$

It is easy to check that if p_E is small, or if β is small, there are many values of α so that both conditions hold. To see this point consider the limit case $p_E = 0$ then the first condition is always true, while the second is equivalent to $\frac{\alpha}{1-\alpha\beta} \leq \frac{1-\alpha}{1-\beta}$, which is true for all $\alpha \leq \frac{1}{2}$. By continuity there are many cases so that both conditions hold.

C.4 Proof of Proposition 4

Part 0: Lemmata 1 and 2, together with Proposition 3, establish that, when *(s1)* or *(s2)* hold and *(f1)* and *(f2)* fail, the outsider threat is effective: insiders implement the voter's preferred policy in equilibrium. To complete the proof of Proposition 4, two additional steps are required. First, we verify when the voter finds it optimal to implement the outsider threat (Part 1), i.e., when Condition (4) holds. Second, we show that no alternative retention strategy—outside those considered in the main analysis—yields a higher expected utility for the voter (Part 2).

Part 1: We begin by characterizing Condition (4):

Remark 1 *To reduce notational burden, we have so far assumed that the policy payoff, for both politicians and voters, is denoted by π. Each time π appeared in a condition, it has represented the politician's policy payoff. However, our arguments do not require that the voter's policy payoff be the same as the politician's. To clarify this distinction, in the following proof, we introduce a separate notation: ν represents the voter's payoff from obtaining her preferred policy.*

Denote ν the voter's payoff from her preferred policy, and let λ_ν denote the discounted expected probability that the voter obtains her favorite policy, which depends on the electoral strategy, and on the policy and signal from the previous period (as that influences whether an outsider is in power in the current period).

The expected discounted payoff from electing an outsider is $\lambda_\nu(S_1|bad, E) = \lambda_\nu(S_2|bad, E)$. That is, it does not matter for this calculation whether S_1 or S_2 is played since both policies, when effective, lead to the same rate of outsiders in power when the politician chose to implement E and the signal

is bad. In order to compute $\lambda_\nu(S_1|bad, E)$, we express it recursively:

$$\lambda_\nu(S_1|bad,E) = \underbrace{1-p_E}_{\text{t: outsider}} + \underbrace{\beta}_{\text{t+1: insider}} + \underbrace{\beta^2((1-\alpha)p_E\lambda_\nu(S_1|bad,E) + (1-(1-\alpha)p_E)\lambda_\nu(S_1|\text{good or } M))}_{\text{t+2 onward}}$$

$$\lambda_\nu(S_1|\text{good or } M) = 1 + \beta((1-\alpha)p_E\lambda_\nu(S_1|bad,E) + (1-(1-\alpha)p_E)\lambda_\nu(S_1|good\ or\ M)).$$

Solving the system we obtain: $\lambda_\nu(S_1|bad, E) = \frac{1-p_E^2(1-\alpha)\beta-p_E(1-(2-\alpha)\beta)}{(1-\beta)(1+p_E(1-\alpha)\beta)}$. The voter's expected discounted utility from electing an outsider after policy E and a bad signal is $U^*(S_1|bad, E) = \lambda_\nu(S_1|bad, E)\nu$. If, instead, the voter deviates and "forgives" the incumbent (or replaces him with another insider), then credibility is lost, and we switched to the voter's worst equilibrium: Insiders choose their preferred (policy E in all future periods in which the state is w_D) and voter always or never reelected the incumbent.[43]

In this case, the voter's expected discounted payoff from keeping the incumbent after policy E and a bad signal is given by $\frac{1-p_D}{1-\beta}\nu$. The representative voter finds it beneficial to implement the threat if and only if $U^*(S_1|bad, E) > \frac{1-p_D}{1-\beta}\nu$, which is equivalent to $\frac{1+p_E^2(\alpha-1)\beta-p_E(1+(\alpha-2)\beta)}{(1-p_E(\alpha-1)\beta)} \geq 1-p_D$. After some simplification it boils down to $\frac{p_E(1-\beta+p_E\beta(1-\alpha))}{1+p_E\beta(1-\alpha)} \leq p_D$, which is condition (4) in Proposition 4.

Part 2: So far, we have considered two mandates — that is, two policy rules mapping states of the world to the action voters expect politicians to take.

-$\mathcal{M}_{mainstr}$ $(w_D, w_E, w_M) \longrightarrow (M, E, M)$

-$\mathcal{M}_{outsider}$ $(w_D, w_E, w_M) \longrightarrow (M, M, M)$

which are assigned to incumbents, either because they align with the voter's preferences (whenever incentive-compatible) or as a form of punishment. Note that there exist six alternative mandates:

-$\mathcal{M}1$ $(w_D, w_E, w_M) \longrightarrow (E, E, M)$,

-$\mathcal{M}2$ $(w_D, w_E, w_M) \longrightarrow (E, E, E)$,

-$\mathcal{M}3$ $(w_D, w_E, w_M) \longrightarrow (M, M, E)$,

-$\mathcal{M}4$ $(w_D, w_E, w_M) \longrightarrow (E, M, M)$,

[43] In the worst case, the voter is left with only two possible outcomes in every period: (i) a misbehaving outsider or (ii) an outsider in power. If $p_D \geq p_E$, the worst equilibrium for the voter is one in which an outsider remains in power indefinitely. But in this case, it is costless for the voters to punish the incumbent, so we focus on the other case.

-$\mathcal{M}5$ $(w_D, w_E, w_M) \longrightarrow (E, M, E)$,

-$\mathcal{M}6$ $(w_D, w_E, w_M) \longrightarrow (M, E, E)$.

These were excluded from our analysis because choosing any of these mandates—either to maximize welfare or to punish insiders—never increases the voter's expected welfare. Hence, a rational voter will never find it optimal to adopt them.

Part 2a: $\mathcal{M}1-5$ yield a lower expected per-period payoff than the outsider threat retention strategy.

Suppose first that each alternative mandate is incentive-compatible. We must compare the voter's welfare under the outsider threat retention strategy, where well-behaved insiders and an outsiders alternate in power, with the welfare under $\mathcal{M}1-5$. Since condition (4) holds, the voter prefers implementing the outsider threat over allowing a misbehaving insider to remain in power, i.e., over $\mathcal{M}1$. Thus, this mandate leads to a lower expected per-period payoff.

Since $\mathcal{M}2$ results in even worse outcomes than $\mathcal{M}1$, it follows that the voter would also reject it. Similarly, $\mathcal{M}3, \mathcal{M}4$, and $\mathcal{M}5$ all trivially yield lower voter payoffs than a ruling outsider, which itself generates a lower expected payoff than the alternation between well-behaved insiders and outsiders.

As a result, if $\mathcal{M}1-5$ are incentive-compatible, they lead to a lower expected per-period payoff than the outsider threat retention strategy. If they are not incentive-compatible, then these mandates leads to insiders following $\mathcal{M}1$, which still generates a lower expected payoff than the outsider threat retention strategy.

Part 2b: $\mathcal{M}1-5$ are a less effective punishment than voting for an outsider.

First, note that $\mathcal{M}1$ is not appealing as a punishment because it allows an insider to implement their preferred policy. Similarly, $\mathcal{M}3, \mathcal{M}4$, and $\mathcal{M}5$ can also be trivially discarded as punishment strategies.

Now consider $\mathcal{M}2$ as a punishment. For it to be preferred over the outsider threat, it must impose a lower cost on the voter, which would require:

$$1 - p_E \geq 1 - p_M - p_D,$$

which simplifies to $p_E \geq \frac{1}{2}$. However, this contradicts our assumption that $p_E < \frac{1}{2}$, ensuring that M is the voter's ex-ante preferred policy. Thus, $\mathcal{M}2$ cannot serve as a better punishment than the outsider threat.

Part 2c: $\mathcal{M}6$.

Disregarding $\mathcal{M}6$ is less immediate because, for sufficiently small p_M, this mandate, when incentive-compatible, yields a higher expected per-period payoff than the outsider threat.

We now show that whenever $\mathcal{M}6$ is an equilibrium, the first-best outcome is attainable. For $\mathcal{M}6$ to be implemented, it must be incentivized through reelection concerns. Since the politician has an incentive to deviate in both w_D and w_M, only a strategy analogous to F_1 can work: replacing a politician following $\mathcal{M}6$ with another politician following $\mathcal{M}6$ if and only if the signal is bad. The range of parameters where this equilibrium holds is given by condition (8).[44] However, when this condition holds, the first-best is implementable, so the voter would be worse off if politicians were to follow $\mathcal{M}6$.

Finally, we show that punishing an insider following $\mathcal{M}_{mainstr}$ with $\mathcal{M}6$ is less effective than using the outsider threat. This is because the incumbent following $\mathcal{M}6$ also requires incentives/punishment. If condition (8) is not satisfied, a politician under $\mathcal{M}6$ would actually prefer not being reelected, since being replaced by a well-behaved insider grants them their favorite policy more frequently. Thus, implementing $\mathcal{M}6$ as a punishment requires an outsider threat, making $\mathcal{M}6$ unapealing punishment scheme for the voter.

[44] The only difference from F_1's computation is that the frequency at which a politician gets their preferred policy changes from $1-p_D$ to $1-p_D-p_M$. As shown earlier, this term cancels out in the analysis because, on the equilibrium path, the rate at which a politician receives their preferred policy does not depend on whether they are in power from the next period onward, making it irrelevant.

Online Appendix

O.1 Multi-Period Punishments

Set $\alpha = 1$, and consider the following $n-$outsider threat: "if betrayed by the incumbent, elect an outsider for exactly n periods before returning to another insider." The next proposition is analogous to Proposition 1, but with the $n-$outsider threat:

Proposition 5 *Always implementing the voter's ideal policy can be sustained in equilibrium with the $n-$outsider threat if the following two conditions hold simultaneously:*

1. *The representative voter is willing to elect an outsider for n terms in response to a deviation:*

$$p_E\left(1-\beta^n\right) \leq p_D. \tag{19}$$

2. *Given this threat, the insider incumbent is well-behaved:*

$$\frac{\pi}{\phi} \leq \bar{\rho}(p_E, n) \equiv \frac{\beta}{1-\beta-p_E\beta(1-\beta^n)}. \tag{20}$$

The union of the parameter sets satisfying Conditions (19)–(20) for all n defines the maximal set of parameters under which the voter obtains her preferred policy in equilibrium. Note that $\overline{\rho}(p_E, n)$ increases with n: the longer the punishment phase, the stronger the incentives for an insider to adopt the voter's preferred policy. However, the longer the punishment phase, the higher the cost for the voter, making it less likely that she will find it beneficial to implement large n-outsider threats.

Proof:

Point 1: If state w_D is more likely than state w_E (i.e., if $p_D > p_E$), poorly incentivized insiders are more costly to voters than outsiders. In this case, n-outsider threat is not costly to implement. In contrast, when $p_E > p_D$, the voter prefers a misbehaving insider to an outsider. If the voter carries out the n-outsider threat, she receives her preferred policy with probability $1-p_E$ for n periods, and

thereafter continues to receive her preferred policy in all future periods, yielding a total utility of

$$\pi\left(\frac{1-\beta^n}{1-\beta}(1-p_E)+\frac{\beta^n}{1-\beta}\right).$$

Conversely, if she tolerates the self-serving insider, she will not receive her preferred policy in every future disagreement state, resulting in a utility of $\pi\frac{1-p_D}{1-\beta}$. Which leads to

$$p_E\left(1-\beta^n\right)\leq p_D \tag{21}$$

Point 2: To determine when the $n-$outsider threat constitutes an equilibrium, we begin by noting that the utility of a well-behaved incumbent remains the same as in the first-best scenario, U^*. What changes is the utility from a one-shot deviation (while holding other politicians' strategies constant), since the incumbent will be replaced by an outsider, which reduces the next n period policy-payoffs and perks:

$$U_p^d(w_D)=\phi+\pi+\beta\frac{1-\beta^n}{1-\beta}(\underbrace{\pi p_M+0}_{\text{an outsider}})+\beta^{n+1}\quad\underbrace{\pi\frac{1-p_D}{1-\beta}}_{\text{an other insider}}$$

Well-behaving is an equilibrium for insider politicians if and only if $U^*(w_D)\geq U_p^d(w_D)$ where the payoffs are given by:

$$U_p^d(w_D)=\phi+\pi+\sum_{i=1}^{n}\beta^i(\pi p_M+0)+\beta^{n+1}\pi\frac{1-p_D}{1-\beta}=\phi+\pi+\pi p_M\beta\frac{1-\beta^n}{1-\beta}+\beta^{n+1}\pi\frac{1-p_D}{1-\beta}$$

and

$$U^*(w_D)=\phi+0+\beta U^*=\phi+\beta(\frac{\phi}{1-\beta}+\pi\frac{1-p_D}{1-\beta}).$$

The inequality thus becomes:

$$\beta(\frac{\phi}{1-\beta}+\pi\frac{1-p_D}{1-\beta})\geq\pi+\pi p_M\beta\frac{(1-\beta^n)}{1-\beta}+\beta^{n+1}\pi\frac{1-p_D}{1-\beta}$$

which simplifies to

$$\frac{\pi}{\phi}\leq\frac{\beta}{(1-\beta)+(p_M-1+p_D)\beta(1-\beta^n)}$$

$$\frac{\pi}{\phi} \leq \frac{\beta}{1-\beta-p_E\beta(1-\beta^n)}.$$

□

O.2 Relaxing restrictions on strategies

Proposition 6 *If voters can randomize (which is not the case in the main text), then the implementation of the voters' preferred policy in every period can be sustained in equilibrium if and only if*

$$\frac{\pi}{\phi} \leq \frac{\beta}{1-(p_E\alpha+1-p_E)\beta}$$

Before proving this proposition, we explain how it relates to the results in the main text. Up to now, we restricted attention to pure retention strategies conditioned only on the most recent period of play. Proposition 6 shows that if voters were able to mix, they could convexify the set of parameters for which the voter's preferred policy is sustainable in equilibrium. In particular, mixing allows voters to fill the "sandwiched" region, where neither F_1 nor F_2 alone can provide sufficient incentives for the incumbent to behave in both states.

Importantly, however, mixing does not eliminate the role of the outsider threat. For instance, parameters in regions S_1 and S_2 remain inaccessible even under optimal mixing. Hence, outsiders can still be elected in the voter's preferred equilibrium.

Finally, setting aside randomization, voters might attempt to mimic some of its benefits by alternating between pure strategies over time (which again is not possible in the main text)–e.g., using F_1 in even periods and F_2 in odd ones. Such switching may partially expand the set of parameters for the first-best by averaging the incentives associated with each strategy.[45] However, this is only a coarse way to convexify the set of parameters, and it remains less effective than mixing. In particular, if $\alpha = 0$ or is close to zero, such strategies provide no additional power to discipline

[45] Consider, for instance, playing F_1 in even periods and F_2 in odd periods. In periods when F_1 is played, incumbents are less tempted to deviate than if F_1 were played in every period, because the prospect of future periods with F_2 increases expected discounted office rents. In periods when F_2 is played, pandering is also less of a concern, since future punishments under F_1 reduce the incumbent's continuation value relative to always playing F_2 (and pandering happens only when future office rents are sufficiently high).

incumbents.

Proof: In Proposition 3, we characterized the maximal set of parameters for which the voter's preferred policy can always be sustained in equilibrium when the voter adopts a pure strategy. As made shown in Corollary 1, and illustrated in Figure 1, there exists a region that is excluded by both F_1 and F_2: under F_1, the incumbent's continuation value is too low to induce good behavior in state w_D, while under F_2, the incumbent chooses M in all states, including in state w_E, in order to stay in power. We now show that, within this region, the voter's ideal policy can be sustained in equilibrium if the voter is allowed to randomize. Specifically, consider the following mixed retention strategy, denoted F_{mix}: *if the signal is bad and the policy is E, replace the incumbent; if the signal is bad and the policy is M, reelect the incumbent with probability r.* When $r = 0$, this is F_1; when $r = 1$, it is F_2. The idea is to choose r high enough to make continuation values sufficiently high to deter deviation in state w_D, but low enough to prevent pandering in state w_E.

Condition for F_{mix} Observe first that, if F_{mix} works so that all politicians are well-behaved along the equilibrium path, the policy payoffs remain unchanged relative to F_1, i.e., $\lambda_\pi(F_{mix}) = \lambda_\pi(F_1) = \frac{1-p_D}{1-\beta}$. In contrast, the ex-ante probability of reelection (i.e., before the state of the world is known) increases from α (with F_1) to $1-(1-\alpha)(p_E+(1-p_E)(1-r))$ (with F_{mix}). Under F_{mix}, the discounted expected payoff from ego rents is

$$\lambda_\phi(F_{mix}) = \frac{1}{1-(1-(1-\alpha)(p_E+(1-p_E)(1-r)))\beta}. \tag{22}$$

The ex-ante expected discounted payoff for the incumbent (before the state of the world is known) if he behaves well is $U^*(F_{mix})$ defined equation (6) with $\sigma = F_{mix}$. When the state of the world is w_D and the incumbent is well-behaved, he is reelected with probability $\alpha + (1-\alpha)r$, so that his expected discounted utility under election strategy F_{mix} can be written as $U^*(F_{mix}|w_D) = \phi+\beta[(\alpha+(1-\alpha)r)U^*(F_{mix}) + (1-(\alpha+(1-\alpha)r))\pi\lambda_\pi(F_{mix})$. In contrast, if he deviates to E, his expected discounted payoff is $U^d(F_{mix}|w_D) = \phi + \pi + \beta\pi\lambda_\pi(F_{mix})$. Deviating is not attractive if and only if $U^*(F_{mix}|w_D) \geq U^d(F_{mix}|w_D)$, which is equivalent to $\beta\phi(\alpha+(1-\alpha)r)\lambda_\phi(F_{mix}) \geq \pi$. Rewriting this condition, the incumbent is well-behaved in the disagreement state w_D with strategy F_2 if and only if:

$$\frac{\pi}{\phi} \leq \frac{\beta(r+\alpha-r\alpha)}{1-(1-p_E)r(1-\alpha)\beta-\alpha\beta} \tag{23}$$

Note that, if $r=0$, this condition is identical to $(f1)$, and if $r=1$, this condition is identical to the RHS of $(f2)$.

We now have to check, in addition, that an incumbent in state E does not want to deviate to play M. When the state of world is w_E, if all politicians are well-behaved, the incumbent's expected utility is $U^*(F_{mix}|w_E) = \phi + \pi + \alpha\beta U^*(F_{mix}) + (1-\alpha)\beta\lambda_\pi(F_{mix})\pi$. If the incumbent deviates, only for one period, by choosing M, he ensures his reelection with probability r in the next period and gets utility $U^d(F_{mix}|w_E) = \phi + r\beta U^*(F_{mix}) + (1-r)\beta\lambda_\pi(F_{mix})\pi$. This one-shot deviation is unattractive for the incumbent if and only if $U^*(F_{mix}|w_E) \geq U^d(F_{mix}|w_E)$, hence if

$$\frac{\pi}{\phi} \geq (r-\alpha)\beta\lambda_\phi(F_{mix})$$

Clearly, for any $r \leq \alpha$, this condition is always satisfied. Otherwise, rearranging the inequality, we obtain the following condition:

$$\frac{\pi}{\phi} \geq \frac{\beta(r-\alpha)}{1-(1-p_E)r(1-\alpha)\beta-\alpha\beta} \tag{24}$$

Conclusion: Fix α, and let $\frac{\pi}{\phi}$ be a value in the region where both conditions $(f1)$ and the lower bound in $(f2)$ fail—that is, in the parameter range sandwiched between those for which F_1 and F_2 sustain the first-best. We now show that such a value can be supported in equilibrium under F_{mix}.

The conditions (23) and (24) define an interval of parameters that can be supported by F_{mix}, and this interval depends continuously on r. For $\alpha = 0$, this interval collapses to a singleton:

$$\frac{\pi}{\phi} = \frac{r\beta}{1+(1-p_E)r\beta}.$$

Moreover, this expression is always contained in the interval defined by conditions (23) and (24) for any $\alpha > 0$. For $r=0$, the expression equals zero; for $r=1$, it matches the lower bound of the F_2 condition (i.e., the pandering constraint). It follows from the Intermediate Value Theorem that there exists $r \in (0,1)$ such that both incentive constraints are satisfied. Thus, F_{mix}, for an appropriately chosen r, sustains an equilibrium with well-behaved insiders.

O.3 Two mainstream parties

In the main text, we examined the incentives of individual politicians. A similar analysis can be conducted by focusing on the incentives of political parties instead. The key distinction between politicians and parties lies in their long-term considerations. While an incumbent politician who fails to get re-elected typically cannot regain power, a party, particularly in a two-party system like that of the United States, has the potential to return to power in the future. Including a two-party system in our model does not alter the main findings but introduces additional effects. Importantly, in a two-party system, it becomes more challenging to incentivize incumbents due to their reduced fear of losing elections, meaning that the first-best is harder to sustain, and the outsider's threat is even more appealing.

To make this point formal, we revisit the first-best with imperfect information (Section 5), by considering a variant of the model in which, instead of a large pool of identical insiders, there are only two. The next proposition is analogous to Proposition 3, but in a setting with two insiders.

Proposition 7 *Consider a pool of candidates of two insiders. The implementation of the voters' preferred policy in every period can be sustained in equilibrium if and only if*

1. *the representative voter plays F_1; and the following condition holds*

$$\frac{\alpha\beta}{1+\beta-2\alpha\beta} \geq \frac{\pi}{\phi} \qquad (f1\text{-}two)$$

2. *or, the representative voter plays F_2, and the following condition holds*

$$\frac{\beta(1-\alpha)}{1+\beta-2(\alpha p_E+1-p_E)\beta} \leq \frac{\pi}{\phi} \leq \frac{\beta(\alpha p_E+1-p_E)}{1+\beta-2(\alpha p_E+1-p_E)\beta} \qquad (f2\text{-}two)$$

Simple inspection reveals that Condition *(f1-two)* is less often satisfied than Condition *(f1)*: When politicians or parties have a high chance to get back in power after losing an election, it is harder to incentivize them. Under retention S_2, the pattern is less straightforward. As in F_1, the incumbent faces a stronger temptation to deviate by choosing their preferred policy. However, the reduced fear of losing elections also means pandering is less problematic in a two-party system, making F_2 more appealing for some parameter values.

Proof:

Point *(f1-two)*: Although the intrinsic payoff does not change ($\lambda_\pi(F_1) = \frac{1-p_D}{1-\beta}$ does not change), the sum of discounted probability of being in power does change due to the possibility of coming back to power. Formally, we can express $\lambda_\phi(F_1)$ recursively:

$$\lambda_\phi(F_1) = 1 + \beta\alpha\lambda_\phi(F_1) + \beta(1-\alpha)\lambda_\phi^{op}(F_1)$$
$$\lambda_\phi^{op}(F_1) = 0 + \beta\alpha\lambda_\phi^{op}(F_1) + \beta(1-\alpha)\lambda_\phi(F_1)$$

where $\lambda_\phi^{op}(F_1)$ represents the discounted expected payoff from ruling, while currently being in the opposition. Solving the system we obtain:

$$\lambda_\phi(F_1) = \frac{(1-\alpha\beta)}{(1-\beta)(1+\beta-2\alpha\beta)} \tag{25}$$
$$\lambda_\phi^{op}(F_1) = \frac{(1-\alpha)\beta}{(1-\beta)(1+\beta-2\alpha\beta)} \tag{26}$$

Provided that the representative voter plays F_1, all politicians being well-behaved is an equilibrium if and only if $\phi+\beta\alpha U^*(F_1)+\beta(1-\alpha)U^{op*}(F_1) \geq \phi+\pi+\beta U^{op*}(F_1)$, which boils down to, $\phi\beta\alpha(\lambda_\phi(F_1)-\lambda_\phi^{op}(F_1)) \geq \pi$. Replacing by their values from equations (25) and (26) and simplifying we obtain:

$$\frac{\alpha\beta}{1+\beta-2\alpha\beta} \geq \frac{\pi}{\phi}. \tag{27}$$

Point *(f2-two)*: We now turn to strategy F_2. Compared to F_1, the ex-ante probability of reelection (i.e., before the state of the world is known) increases from α to $\bar{\alpha} \equiv p_E\alpha + 1 - p_E \geq \alpha$. Using an identical reasoning as above we obtain:

$$\lambda_\phi(F_2) = \frac{(1-\bar{\alpha}\beta)}{(1-\beta)(1+\beta-2\bar{\alpha}\beta)},$$
$$\lambda_\phi^{op}(F_2) = \frac{(1-\bar{\alpha})\beta}{(1-\beta)(1+\beta-2\bar{\alpha}\beta)}.$$

Provided that the representative voter plays F_2, all politicians being well-behaved is an equilibrium if and only if $\phi+\beta\bar{\alpha}U^*(F_1)+\beta(1-\bar{\alpha})U^{op*}(F_2) \geq \phi+\pi+\beta U^{op*}(F_2)$, which boils down to, $\phi\beta\bar{\alpha}(\lambda_\phi(F_2)-$

$\lambda_\phi^{op}(F_2)) \geq \pi$. Replacing by their values and simplifying we obtain:

$$\frac{\beta\bar{\alpha}}{1+\beta-2\bar{\alpha}\beta} = \frac{\beta(\alpha p_E + 1 - p_E)}{1+\beta-2(\alpha p_E + 1 - p_E)\beta} \geq \frac{\pi}{\phi} \tag{28}$$

Finally, we have to check, that an incumbent in state E does not want to deviate to play M. When the state of world is w_E, if all politicians are well-behaved, the expected utility of the incumbent is $U^*(F_2|w_E) = \phi+\pi+\alpha\beta U^*(F_2)+(1-\alpha)\beta U^{op*}(F_2)$. If the incumbent deviates, only for one period, by choosing M, he ensures his reelection in the next period and gets utility $U^d(F_2|w_E) = \phi + \beta U^*(F_2)$. This one-shot deviation is unattractive for the incumbent if and only if $U^*(F_2|w_E) \geq U^d(F_2|w_E)$, hence if $\pi \geq \beta(1-\alpha)\phi(\lambda_\phi(F_2) - \lambda_\phi^{op}(F_2))$, or

$$\frac{\pi}{\phi} \geq \frac{\beta(1-\alpha)}{1+\beta-2(\alpha p_E + 1 - p_E)\beta}. \tag{29}$$