Flip-flopping and Electoral Concerns *

Giovanni Andreottola ⋆

January 2018

Please download the latest version here: www.giovanniandreottola.com/research

Abstract

Politicians who change their mind on a policy issue are often confronted with the accusation of being flip-floppers. However, a changing environment sometimes makes policy revisions necessary. The present analysis suggests that flip-flopping on a policy issue is detrimental to a politician's reputation because it sends a bad signal on the accuracy of his information. As a result, electorally concerned politicians can have the incentive to stick to their initial policy choice, despite it being inefficient, in order to avoid the stigma of flip-flopping. This distorted behaviour is not only damaging in terms of policy welfare, but also in terms of a worse selection of competent politicians through elections. I also discuss benefits and costs of several possible ways to address the unwillingness of politicians to respond to information: these include a single-term limit rule, the effect of media varying the transparency of actions and consequences as well as delegation of one of the actions to an independent agent.

*This paper is I am grateful to my advisors Andrea Mattozzi and David Levine for their great supervision. The paper also benefitted from useful discussions with Ran Eilat, Piero Gottardi, Andrea Galeotti, Bard Harstad, Helios Herrera, Giacomo Ponzetto, Massimo Morelli, Santiago Oliveros, Alessandro Riboni and several others. I would also like to thank participants at the EUI Microeconomics Working Group, the EUI Workshop on Media Effects, the IOEA in Cargese and the NICEP Workshop in Nottingham and the Leitner Seminar in Yale for their questions and comments.

⋆Department of Economics, European University Institute.
Email: Giovanni.Andreottola@eui.eu Website: www.giovanniandreottola.com

But with Kerry the charge isn't that he's inconstant. It's that in his inconstancy he flips wrong – the far more serious charge of bad judgment.
Mickey Kaus (Slate)

Bush's decision-making style was based on his gut instincts [...] Bush was quick to reach decisions, and, once reached, he saw change as a sign of weakness.
(Newsweek)

1 Introduction

Consistency is one of the qualities that voters value the most in a politician. As the political scientist Fearon (1999) wrote: "*If I think of elections as a problem of choosing a competent, like-minded type not easily bought by special interests, then it makes perfect sense to be highly concerned with principledness and consistency*". As a result, voters tend to dislike politicians who change their mind on a policy issue, which is disparagingly denoted as *flip-flopping* [1]. As a matter of fact, one of the most frequent attacks used in electoral races is the allegation of being a flip-flopper. Two famous cases of presidential candidates that have considerably suffered from being viewed as flip-floppers are John Kerry and Mitt Romney. The fact that voters tend to punish politicians who flip-flop is also to a large extent confirmed by survey evidence collected by political scientists[2].

In a changing world in which politicians are constantly exposed to new information, however, changing one's mind on an issue is natural and actually the optimal thing to do in many situations: as Keynes put it, *"When the facts change, I change my mind"*. In this respect, the reputational stigma of flip-flopping can seem puzzling. In this paper I show how flip-flopping can be detrimental to a politician's reputation even in cases where ideological or private-interest related concerns are absent, i.e. situations in which changing one's mind simply reflects a change in the information available to the politician. In particular, flip-flopping is rationally punished by voters if i) the optimal policy choice is persistent ii) politicians have private information which cannot be credibly revealed to the public and iii) voters are not (fully) capable of judging the validity of a policy choice. The reason for this penalization is that policy shifts are associated with poorly informed politicians; therefore, voters trying to select well-informed (competent) politicians assign

[1] Use of this expression dates back to at least 1890 according to The New York Times archives. Other terms used to shed negative light on the change of course of a politician are *u-turn* and *backflip*.

[2] See Tomz and Houweling (2012) and Doherty, Dowling, and Miller (2015): I will elaborate more about these and other studies in the related literature section.

a better reputation to politicians who do not flip-flop. There is indeed evidence of such a rationale in political commentary: Jack Shafer of the media outlet Politico, for example, wrote the following [3] referring to Hillary Clinton: "*So if new or better information has been the impetus for her policy shifts, she must concede that she has a fat history of taking the wrong position in the early going and then requiring a re-do*".

If one part of the story has to do with how voters perceive policy shifts, the flip side of it has to do with how the reputational stigma attached to flip-flopping affects politicians' behaviour. Since flip-flopping is bad for a politician's electoral prospects, a strongly office-motivated politician will have the incentive to distort his behaviour in order to avoid flip-flops: my model therefore describes a form of electoral pandering which is endogenously induced by the previous action of the politician. In other words, politicians display an excess of conservatism, or postured consistency, with respect to their previous decisions.

One historical example in which such a logic seems to have played an important role is the decision to start the Iraq war: George Bush and Tony Blair believed Saddam Hussein had (or was in the process of developping) weapons of mass destruction and that waging war on Iraq was necessary to stop him. When evidence pointing in the opposite direction was revealed (including intelligence reports and the work of UN inspectors), the two presidents should have realized that the threat posed by Saddam was after all not so large and reconsidered their plans to invade Iraq. However, they never performed this (beneficial) *flip-flop*. The publication of the Chilcot report in the United Kingdom has recently provided further evidence that Tony Blair's decision to maintain his support for the Iraq invasion was not in accordance with the information he had received from his intelligence sources. As The Guardian put it: "*That was the point at which the UK government could and should have said the US must count the UK out. Blair should have admitted that this was a line in the sand. But he didn't call a halt*".

After establishing that reputational concerns lead to an insufficient amount of flip-flopping, the second part of the paper discusses some institutional design approaches to tackle the problem: the first-one is a single term limit policy. By forbidding the re-election of incumbents, such a policy eliminates all policy distortions: as a matter of fact, in this model electoral rents are the only source of misalignment between voters and politicians. At the same time, however, the single term limit forces voters to forgo all the potential gains of learning about a politician's type from his track-record. Not least, such policies require commitment, for example through constitutions.

[3] The full article can be found at this link: `http://www.politico.com/magazine/story/2015/10/democratic-debate-hillary-clinton-flip-flop-213247`

The other institutional feature that I discuss is the media. The importance of the media in informing citizens and evaluating policies goes without explanation, but how does that affect the issue of distorted consistency highlighted in my model? I focus my attention on the two main roles played by the media: the reporting of politicians' policy choices (reporting media) and the evaluation of policy (commentator media).

Two results from the analysis of the media model stand out: first of all, fully accurate reporting of policy choices is never optimal. A noisy media can insulate the politician from the reputational stigma of flip-flopping: lies are thus crowded out by noise. Moreover, the net effect of this substitution is such that the performance of elections in selecting competent types improves. If the noise is too large, however, selection eventually worsens (no learning is possible when the media is fully noisy).

The non-monotonic effects of the media reporting accuracy suggest that the significantly simpler and broader access to politicians' track records made possible by improvements in technology and the rise of social media might have led to an increase in policy distortions. This result is related to the idea that transparency can be damaging for political accountability: adding a noisy reporting media to the model is in fact similar to relaxing the assumption of full action transparency.

Relatedly, I also show that delegating the first action to an independent agent (for example a bureaucrat or a committee) can be beneficial even if this agent is incompetent: flip-flopping still signals incompetence but the incentives for the incumbent to distort his actions are diminished. This result also sheds some light on the comparison between my model and a single action pandering model.

The other type of media that I analyze is what I define commentator media: instead of reporting on the actions of the politician, which are now again assumed to be commonly known, the commentator media sends a signal on the state of the world. Whereas this usually acts towards disciplining politicians to follow their signal, there are situations in which increasing the informativeness of the commentator media results in a more distortive behaviour by politicians [4]. The reason is that when the payoff from receiving a media endorsement increases, being endorsed as a consistent policy-maker can be so profitable that politicians are willing to further distort their actions and gamble on this favourable outcome.

The main takeaway is that the returns to having more accurate media might be low if not even negative. This could have relevant implications, for example with respect to the public subsidization of media outlets.

[4] This can happen if i) persistence is high ii) incompetent types have a sufficiently informative signal and iii) most politicians are competent.

Finally, I also study a variation of the model in which a public signal about the state of the world is observed before the politician chooses the second action. In this setup, I show that, unlike in the benchmark model, flip-flopping can be motivated by electoral opportunism and the equilibrium level of flip-flopping can be larger than the socially optimal one. The reason is that (under some parameter restrictions) when a public signal is available flip-flopping only occurs to match the public signal (one can think of it as an opinion poll): therefore, politicians have an incentive to flip-flop and match the poll even if their information does not support such a move.

2 Related Literature

This paper is related to several streams of literature. The electoral concerns model that I consider builds on Canes-Wrone, Herron, and Shotts (2001), Prat (2005) and Ashworth and Shotts (2010). Besides some differences in how the model is built, of secondary importance for the results, the main novelty of my model is to introduce an additional period in which the incumbent takes an action. This feature is central for the contribution of the paper, since it allows me to delve into the intrinsically dynamic nature of flip-flopping. Moreover, by adding a *stage-zero* to an electoral concerns model, my analysis also allows for a simple endogenous interpretation of the concept of pandering towards a popular action described in electoral concerns models. In my model, as a matter of fact, the popular action is simply the previous policy choice. In this sense, my work contributes to the political economy literature on conformity (sometimes also denoted as pandering or herding): along with the aforementioned Prat (2005), one of the seminal contributions is by Maskin and Tirole (2004) - who compare the welfare properties of representative democracy, direct democracy and judicial power; Levy (2007) considers a committee of career concerned decision makers and Frisell (2009) shows that voters' beliefs about the politician can have self-fulfilling consequences. Levy (2004), on the other hand, shows that in a similar setting also anti-herding (anti-conformism) can take place.

Connected to the idea of pandering towards a popular action is that of status-quo bias and propensity to reform: examples of models with endogenously status-quo biased politicians include Fu and Li (2014), where career-concerned policy makers undertake reform with lower than optimal probability, and Dewan and Hortala-Vallve (2014), in whose model voters learn through either the success of a reform or the information provided by a rival candidate. My work, therefore, links the concepts of pandering and status quo-bias: the past actions of the politician influence the voters' beliefs to which the politician has an incentive to conform.

As far as the discussion on media is concerned, Prat (2005) and Ashworth and Shotts (2010) are the closest references. The introduction of a commentator media sector in the model draws mostly on Ashworth and Shotts (2010): my contribution lies in the characterization of partially truthful equilibria (not just finding conditions for a truthful equilibrium) and in the comparative statics analysis which shows how increasing the accuracy of the commentator media need not decrease the distortion of politicians' behaviour. These comparative statics also innovates with respect to the paper of Gentzkow and Shapiro (2006): in their paper, the decision maker is a media outlet, which panders towards the prior belief of citizens in order to form a reputation of accuracy. Whereas Gentzkow and Shapiro show that the more likely it is for voters to learn the true state of the world, the lower the pandering by the media, I show that introducing an informative media might actually lead politicians to act in a more distorted way.

The analysis of the reporting media, on the other hand, is related to the discussion of action versus consequence transparency in Prat (2005). Since for many policy choices the assumption of action secrecy is not a realistic alternative, I show that the accuracy of the reporting media plays a similar role, and I demonstrate that the optimal arrangement is to have some but not perfect information on the action taken by the politician.

The single term limit rule I mention in the institutional design section, on the other hand, is related to the comparison between representative democracy and judicial power carried out by Maskin and Tirole (2004). A single term limit rule has similar effects to those of delegating decision making to a judicial power not subject to elections.

The consequences of reputation concerns on expert behaviour have also been studied outside the electoral environment. In particular, repeated action by experts has been analyzed by Prendergast and Stole (1996), Li (2007) and in a recent paper by Aghion and Jackson (2016): in addition to many differences in the modelling strategy, the main conceptual difference between my analysis and that of Prendergast and Stole has to do with the fact that I consider a changing environment rather than a fixed state. Aghion and Jackson, on the other hand, consider repeated action over a sequence of independent and identically distributed states (whereas I allow for correlation) and with action consequences being observable to the principal (whereas they are unobservable in my baseline model). The defining feature of Li's paper is to assume that decision makers' information becomes more accurate in the second period. In her model, a (fully reputation motivated) agent has to do make two reports before being evaluated by the principal. The state of the world is fixed and observable before the evaluation takes place, but the agent can be paid a wage that only depends on beliefs about her ability. In such a setup it is possible for agents to improve their reputation by changing their mind, unlike in my model. In other situations, the premium for consistency leads low competence agents to gamble on being proved consistently right, similarly to what happens in my model with the commentator media. The idea of gambling on a policy likely to be proved wrong is also at the heart of the paper by Majumdar and Mukand (2004): theirs is a model of experimentation, in which an incumbent can choose to implement a risky policy and has to decide whether to continue the project after a potentially unsuccessful trial. In their model, low type governments are inefficiently reluctant to abandon bad projects, gambling on the small probability of success that would boost their reputation.

Another paper which shares part of the mechanism with my work is Patacconi and Vikander (2015), in which a policy maker receives two signals from an agency and a conflict between voters and policy-maker can arise when the signals are mixed, similarly to what happens in my model after a mixed history of signals. However, whereas in their study voters protest to get an unbiased policy enforced, in the environment I describe there are no protests but elections to select the most competent candidate.

Since my paper explains the reputational stigma of flip-flopping, it is also worth mentioning some political economy works on this topic. Whereas my paper is the first, to the best of my knowledge, to provide a theory of flip-flopping by a politician taking repeated decisions over an issue, there exists some theoretical work by Agranov (2016) and Hummel (2010) dealing with flip-flops between primary and general elections, hence with a completely different objective than that of my analysis.

On the empirical side, finally, there are several papers assessing how voters react to politicians flip-flopping. Doherty, Dowling, and Miller (2015), for example, find that flip-flopping affects the perception voters have of a politician. They show that voters are more forgiving of flip-flops on complex issues or issues which are far away in time. These predictions are consistent with my model, in which flip-flops are a bad signal the higher the state persistence (and we can think that persistence is lower over long periods) and the worse is the signal that voters have. In the same paper, the authors also show that the reputational cost of a flip-flop is compensated by the fact that the new position taken by the politician will be seen favourably by some voters, creating a trade-off for the politician. Similarly, Tomz and Van Houweling (2012) conduct survey experiments showing that candidates repositioning affects their support not only in terms of committment to an ideological issue, but also in terms of perceived valence. The valence aspect is important in particular for issues that are not too salient for voters.

Another issue that has been studied is whether the effects of a flip-flop differ between issues of principle (as abortion or gay marriage) versus pragmatic issues related to a specific policy. Tavits (2007) shows that flip-flopping on pragmatic issues is seen less badly than flip-flopping on ideological issues. On the other hand, Tomz and Houweling (2012) do not find any difference among issues: independent of the issue, candidates who reposition perform worse. In my model there is no difference between types of issues: I assume that one decision is correct depending on the state of the world and that citizens would all agree if perfectly informed. Tomz and Houweling (2012) also argue that the bad perception of flip-flopping will deter candidates away from it: this is exactly what I happens in the formal model I present, where flip-flopping is not bad per se, but politicians avoid it because it carries a bad signal for their reputations.

Levendusky and Horowitz (2012) show experimentally that in the context of international relations, leaders who make threats and subsequently back down pay a cost in terms of electoral support and reputation (which in the international relations literature is called *audience cost*): one of the main reasons for this effect is that a leader changing his mind is seen as less competent than one who stays coherent. What is more, they show that partisanship does not play a significant role in the determination of audience costs. This evidence seems to capture a mechanism lying very close to the one I present in my model. As a matter of fact, given the presence of asymmetric information between politicians and voters and the fact that politicians are often evaluated for their foreign policy conduct before the consequences of their actions are fully known, foreign policy issues are among those where the theory I develop should have more bite.

3 The Model

There is an incumbent politician, a set of identical voters and two periods $t \in \{1, 2\}$. The politician's job is to take an action $a_t \in \{0, 1\}$ in each period. The action's payoff depends on the underlying state of the world, which can take two values $\omega \in \{0, 1\}$. The initial probability of each state is equal to $p_1 = \frac{1}{2}$; moreover, the state is persistent so that for $j \in \{0, 1\}$, $Pr(\omega_2 = j|\omega_1 = j) = \gamma > \frac{1}{2}$. The utility for taking the right action is normalized to 1 and it applies to both voters and politicians. At the beginning of period one, an incumbent is randomly drawn. Incumbents can be of two types: competent and incompetent. In each period, both types receive an informative signal $s_t \in \{0, 1\}$ on the state of the world, but the accuracy of the signal, $Pr(s_t = j|\omega_t = j) = q_\theta$ depends on the politician's type $\theta \in \{H, L\}$: $\frac{1}{2} < q_L < q_H \leq 1$. To simplify exposition I fix $q_H = 1$

(competent politicians perfectly observe the state) and therefore I drop the subscript from q_L, which will be simply denoted by q. Notice that since the initial prior is $\frac{1}{2}$, for all $\gamma < 1$ the signal received by any politician is always decision-relevant. A signal is decision-relevant when the probability of matching the action to the state is maximized by $a_t = s_t$. I denote by $\rho_t = Pr(\omega_t = s_t|s_t, \theta, s_{t-1})$ the posterior probability that the state is equal to the signal after observing realization s_t. Since the posterior of the perfectly informed competent politician is always equal to 1, ρ will denote, when not further specified, the posterior of the incompetent politician. In particular, ρ_2 will be used as short form for $\rho_2(\omega_2 = s_2|s_1, s_2 \neq s_1, \theta = L)$, whereas $\bar{\rho}_2$ will indicate $\rho_2(\omega_2 = s_2|s_1, s_2 = s_1, \theta = L)$.

Politicians know their competence, and the signals they receive are private information: competent politicians represent a fraction λ of incumbents. At the end of period $t = 2$ there is an election, in which the representative voter decides whether to retain (r) or fire (f) the incumbent politician (denote the decision $e \in \{r, f\}$). Voters know the statistical process governing the economy and they observe the track-record of the incumbent politician, i.e. the actions taken over the two periods, denoted by $\tau = (a_1, a_2)$. Track-records are used to form beliefs $\mu(a_1, a_2) = Pr(\theta = H|a_1, a_2)$ over a politician's competence, which I call reputation. Before the election takes place, a challenger appears. The challenger is competent with probability λ_O. Moreover, as I will describe shortly, a draw from a uniform distribution in $v \in [-b, b]$, observed before the election takes place and independent of the competence of incumbent and challenger, determines the relative valence of the challenger versus the incumbent. The representative voter's utility depends on whether the politician's action matches the state of the world (plus the valence draw in case the challenger is elected); moreover, electing a competent politician gives voters a utility of b (denote by θ_e the type of the election winner). Hence, $U_c = \sum_{t=1}^{2} \mathbb{1}_{a_t=\omega_t} + \mathbb{1}_{e=f} z + \mathbb{1}_{\theta_e=H} b$. Politicians derive utility both from taking the right action while in office and winning the election, with $\frac{\phi}{2}$ denoting the additional utility received if re-elected. Formally: $U_p = \sum_{t=1}^{2} \mathbb{1}_{a_t=\omega_t} + \mathbb{1}_{e=r} \frac{\phi}{2}$.

As far as players' strategies are concerned, I'll start with politicians. A politician's information set, or history, denoted by h_t, includes all actions up to $t-1$ and signals up to t. The strategy of the incumbent is a mapping Ψ from any history h_t to any probability distribution of actions $a_t \in \{0, 1\}$. In particular, since signals are decision-relevant, it is useful to express the incumbent's strategy as the probability $\sigma(s_t)$ of choosing a policy a_t in accordance to the realization at time t of the signal, denoted by s_t.

As far as the voter's strategy is concerned, it implies retaining or firing the politician in the election at the end of $t = 2$. The voter's decision is $e \in \{r, f\}$ and it depends on the incumbent's track record $\tau = (a_1, a_2)$. Voters choose each politician with probability $\frac{1}{2}$ when indifferent [5].

The equilibrium concept I use is Perfect Bayesian (PBE), but I do not consider pooling equilibria where at each time t, both types play the same action with probability 1. These equilibria are ruled out by a simple trembling-hand perfection refinement [6].

4 Results

I start the analysis from the election in which the voter chooses between the incumbent and the challenger. Just before the election, the valence draw v is realized. As a result, voting for the challenger gives the voter a utility of $v + \lambda_O b$. Voting for the incumbent gives instead the voter a utility of $\mu(a_1, a_2)b$. Therefore, the incumbent is re-elected if $v \leq (\mu(a_1, a_2) - \lambda_O)b$, and $Pr(v \leq (\mu(a_1, a_2) - \lambda_O)b) = \frac{(\mu(a_1,a_2)-\lambda_O)b-(-b)}{2b} = \frac{1}{2} + \frac{\mu(a_1,a_2)-\lambda_O}{2}$. In other words, the re-election probability is linearly increasing in reputation.

Having described the election stage, go back one step and consider the decisions of the incumbent in periods $t = 1$ and $t = 2$. Signals are always decision-relevant, so maximizing the probability to match the action to the state requires that politicians follow their signals. If that happens in equilibrium, then I call the equilibrium truthful.

Definition 1. *An equilibrium is truthful if and only if $\sigma(s_t) = 1$ for each s_t at any t and for each type θ.*

Let's start with a simple observation about the reputation of different track records under truthful play, which I will denote by $\mu^T(a_1, a_2)$ in order to stress the fact that $a_t = s_t$. There are four possible track records, $\tau \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. Given the symmetric initial prior, however, the probability of obtaining each of the two consistent and flip-flopping signal sequence is the same, and therefore, as the next claim will prove, $\mu^T(0, 0) = \mu^T(1, 1)$ and $\mu^T(0, 1) = \mu^T(1, 0)$. In other words, all that matters is whether the politician played the same action over the two periods or changed his mind. In particular, I will call the former *consistent* track-records and the latter *flip-flopping* track-records, and indicate them as $\tau = C$ and $\tau = F$. I will therefore indicate with μ^T_C and μ^T_F consistent and flip-flopping track-records under truthful play.

Fact 1. *Under truthful play, the reputation of a consistent track-record is strictly larger than that of a flip-flopping track-record: $\mu^T_C > \mu^T_F$.*

[5] Notice, however, that this is a zero probability event, since challengers are distributed according to an atomless distribution.

[6] For an example and the explanation of how the trembling-hand refinement eliminates these equilibria, refer to the Appendix.

Proof. Take any flip-flopping track record. Under truthful play, $a_t = s_t$, so that $Pr(a_2, a_1|\theta) = Pr(s_2, s_1|\theta)$. Denote by $\frac{A(q_\theta)}{2}$ the probability that a politician receives a flip-flopping sequence of signals:

$$\frac{A(q_\theta)}{2} = \frac{1-\gamma}{2}(q_\theta^2 + (1-q_\theta)^2) + \frac{\gamma}{2}2q_\theta(1-q_\theta).$$

First notice that since $p_1 = \frac{1}{2}$, the above expression holds for both types of flip-flopping sequences of signals $(0,1)$ and $(1,0)$. It can be noted that $\frac{A(q_\theta)}{2} = \frac{1-\gamma}{2} + \frac{2\gamma-1}{2}2q_\theta(1-q_\theta)$ and since $\gamma > \frac{1}{2}$, $2\gamma - 1 > 0$ and $A(q)$ is decreasing in q. Hence, $A(1) = 1-\gamma < A(q)$. Finally, let's construct the reputations from the 4 possible track-records. Since $\frac{A(q)}{2}$ only depends on whether the politician received a consistent or flip-flopping sequence of signals, there are only two possible levels of reputation: one from flip-flopping and one from being consistent. To see that reputation from flip-flopping is lower than from consistent play consider the following inequality:

$$\mu_F^T = \frac{\lambda A(1)}{\lambda A(1) + (1-\lambda)A(q)} < \frac{\lambda(1-A(1))}{\lambda(1-A(1)) + (1-\lambda)(1-A(q))} = \mu_C^T$$

From now on, I will simply write $1-\gamma$ for $A(1)$ and denote by $A \equiv A(q)$. $\square$

The result I just stated is very important for the development of the whole paper since it highlights the reason that leads incumbents to distort their actions: the bad reputation associated with flip-flopping. The stigma of flip-flopping, as a matter of fact, puts an office motivated incumbent in front of a trade off. Whenever receiving a second signal that contradicts the first one, i.e. $s_2 \neq s_1$, a politician knows that doing the right thing for society will get him a worse reputation. If the flip-flopping stigma is large enough compared to the benefit from following his signal, the politician will choose to act contradicting his information.

In light of this, it can be shown that a necessary and sufficient condition for the existence of a truthful equilibrium is that incumbents have the incentive to follow their signals after receiving a non-constant stream of signals. Moreover, a single crossing property makes it sufficient to simply look at the incentives of incompetent incumbents. The reason is that whereas both competent and incompetent politicians enjoy the same benefit from avoiding a flip-flop, they do not sustain the same costs: having a less accurate signal, as a matter of fact, means having a larger probability of matching the state of the world when playing $a_t \neq s_t$. Lying is therefore cheaper for incompetent politicians. This property is very useful for the characterization of the equilibrium.

Fact 2 (**Single Crossing Property**). *The cost of acting against one's signal is* $\rho_t(\theta) - (1 - \rho_t(\theta)) = 2\rho_t(\theta) - 1$*, and since* $\rho_t(\theta = L) < \rho_t(\theta = H) = 1$*, contradicting the signal is more costly for the competent politician.*

It follows that a truthful equilibrium is only sustainable as long as incompetent politicians are willing to follow their signal at $t = 2$ after receiving a stream of non-constant signals. This in turn requires the office motivation parameter to be low enough, because as ϕ grows larger, the benefits from holding office progressively dwarf the utility from matching the action to the state of the world.

Proposition 1. *A truthful equilibrium is sustainable as long as* $\phi \leq \bar{\phi}$*, where*

$$\bar{\phi} = \frac{2\rho_2 - 1}{\mu_C^T - \mu_F^T}$$

Proof. See Appendix. □

The parameter $\bar{\phi}$ is therefore the upper-bound on electoral rents[7] under which a truthful equilibrium is sustainable. The question now is: what happens when ϕ is larger than $\bar{\phi}$? Moreover, does the game allow for a unique equilibrium? The answer to these questions is given by the following theorem:

Theorem 1. *The game always has a unique non-pooling Perfect Bayesian Equilibrium. For* $\phi \leq \bar{\phi}$*, the unique equilibrium is the truthful equilibrium. For* $\phi > \bar{\phi}$*, the unique equilibrium is partially truthful, meaning that:*
$\sigma(s_1) = \sigma(s_2|\theta = H) = \sigma(s_2 = a_1|\theta = L) = 1$ *and* $\sigma(s_2 \neq a_1|\theta = L) = \sigma^* < 1$.

Proof. See Appendix. □

This theorem proves that the game always has a unique non-pooling equilibrium. It must therefore be the case that when $\phi \leq \bar{\phi}$, the unique equilibrium is truthful. When $\phi > \bar{\phi}$, on the other hand, the unique equilibrium is only partially truthful, since whenever the signal received by politicians in the second period suggests flip-flopping, incompetent politicians mix between following their signal and pandering towards their previous action.

Theorem 1 has two interesting implications: the first is that for any level of ϕ, flip-flopping always decreases the incumbent's reputation (since the two reputations average at λ, a bad reputation is always below λ). However, the larger the flip-flopping avoidance

[7]The actual rents for the politician are $\frac{\phi}{2}$, but that only serves the purpose of simplifying calculations and has no other consequence on the model.

distortion caused by incompetent politicians not following their signal, the smaller is the reputation gap between consistent play and flip-flopping, i.e. the less bad is the reputation from flip-flopping. As ϕ tends to infinity and politicians only care about re-election, the reputation gap between consistency and flip-flopping approaches zero.

The second (and related) implication is that when the equilibrium is partially truthful, there is an insufficient amount of flip-flopping compared to the truthful equilibrium. In other words, voters stigmatize flip-flopping but at the same time they would be better off if more flip-flopping took place. As a matter of fact, it could even happen that voters are more confident about the policy being correct after seeing a flip-flop rather than consistent policy: a flip-flop sends a bad signal on the incumbent's type but is always earnest, whereas a consistent policy is a good sign on the politician's type but not necessarily earnest. I summarize these insights in the following corollary:

Corollary 1. *In equilibrium, flip-flopping gives a bad reputation compared to truthful play: $\mu_C^* - \mu_F^* > 0$ for any ϕ. Equivalently, $\mu_C^* > \lambda > \mu_F^*$. In a partially truthful equilibrium, there is insufficient flip-flopping compared to the truthful equilibrium.*

Notice that another interesting implication of the model is that change hurts incumbents: when the state of the world changes, the reputation of incumbents is likely to fall (because of flip-flopping) and this means the incumbent is more likely to get replaced. At the same time, this also means that conditional on a change in leader, leaders who rise to power after a change in the state are worse than those who get in office after a period of stability. On the other hand, conditional on having a bad leader in office, a change in state increases the chances of having a better leader in the following period. In other words, improvements in leadership are more likely after a change in the state of the world.

Fact 3. *A leadership change is more likely after the state of the world changed. Conditional on a leadership change, leaders who rise to power after a change in the state of the world are worse; however, a change in the state of the world increases the chance of having a better leader in the following period.*

Proof. See Appendix. □

4.1 Comparative Statics and Welfare

First of all, the definition of welfare in this model is based on the expected utility of voters, calculated as of time $t = 0$ (i.e. before randomly picking the incumbent); as such, welfare does not account for the utility of politicians [8].

Definition 2. *Social welfare W is defined as:*

$$W = \mathbb{E}_0\left[\sum_{t=1}^{2} \mathbb{1}_{a_t=\omega_t}\right] + \mathbb{E}_0[v|e=f] + \mathbb{1}_{\theta_e=H}b$$

Welfare can be decomposed in two parts, accountability and selection. Accountability indicates whether the incumbent acts in the best possible manner for society, which in this case means to follow the signal: accountability welfare is therefore the probability that the incumbent chooses the optimal policy, formally $\sum_{t=1}^{2} \mathbb{1}_{a_t=\omega_t}$. Selection-welfare, on the other hand, indicates the utility the voter derives from the election winner, also accounting for the valence shock: formally $\mathbb{1}_{\theta_e=H}b + \mathbb{1}_{e=f}v$.

Let's denote by $d \equiv \lambda - \lambda_O$ and by $\tilde{q} = q - A(1-\sigma^*)(2\rho_2 - 1)$ the accuracy of the incompetent's politician policy choice, taking into account that not following the signal changes its accuracy across the two states of the world. The expression for welfare can be rewritten in the following way:

$$W = [\lambda + (1-\lambda)q] + [\lambda + (1-\lambda)\tilde{q}] + b\left[\frac{1}{4} + \frac{\lambda_O^2}{4} - \frac{\lambda\lambda_O}{2} + \frac{\lambda+\lambda_O}{2} + \frac{\mathbb{E}\mu^2}{4}\right] \quad (1)$$

Proof. As far as $t=1$ and $t=2$ are concerned, since the competent politician follows his signal, which is perfectly accurate, he always takes the right decision. The incompetent politician, instead, follows the signal in the first period, taking the right decision with probability q, but in the second period, if the signal indicates flip-flopping as the optimal action, he contradicts it with probability $1-\sigma^*$. As a result, the accuracy of the incompetent incumbent's signal is $(1-A)\bar{\rho}_2 + A(\sigma^*\rho_2 + (1-\sigma^*)(1-\rho_2)) = q - (1-\sigma^*)(2\rho_2-1)A$, to get which I used the fact that $(1-A)\bar{\rho}_2 + A\rho_2 = q$, since $q = Pr(\omega = s_1|s_1)$. As far as the valence draw is concerned, the expression for expected welfare is the following: $\mathbb{E}[v|e=f] = \sum_{\tau\in\{C,F\}} Pr(\tau)Pr(e=f|\tau)\mathbb{E}[v|v\, b(\mu_\tau - \lambda_O)]$. Now, $\mathbb{E}[v|v\, b(\mu - \lambda_O)] = \frac{b}{2}[1 + (\mu_F - \lambda_0)]$ and $Pr(e=f|\mu) = \frac{1}{2} - \frac{\mu_F - \lambda_0}{2}$. So $\mathbb{E}[v|e=f] = \mathbb{E}_\mu \frac{b}{4}[1-(\mu-\lambda_O)^2] = \mathbb{E}_\mu \frac{b}{4}[1-d^2]$. Finally, $Pr(e=f) = \mathbb{E}_\mu\left[\frac{1}{2} - \frac{\mu-\lambda_O}{2}\right]$ whereas $Pr(e=r|\theta=H) = \frac{1}{2} - \frac{\lambda_O}{2} + \frac{Pr(\tau=C|H)\mu_C + Pr(\tau=F|H)\mu_F}{2}$, where $Pr(\tau=C|H) = \gamma$ and $Pr(\tau=F|H) = 1-\gamma$. Taking expectations with respect to μ where necessary, and considering that $\mathbb{E}\mu = Pr(\tau =$

[8] This is standard in political economy, as we consider politicians too small a fraction of the population for their welfare to matter in the aggregate.

$C)\mu_C + Pr(\tau = F)\mu_F = \lambda$, the welfare expression can be rewritten as

$$W = [\lambda + (1-\lambda)q] + [\lambda + (1-\lambda)\tilde{q}] + \frac{b}{4}[1 - \mathbb{E}_\mu(\mu - \lambda_O)^2]$$
$$+ b\lambda_O\left(\frac{1}{2} - \frac{\lambda - \lambda_O}{2}\right) + b\lambda\left[\frac{1}{2} - \frac{\lambda_O}{2} + \frac{\gamma\mu_C + (1-\gamma)\mu_F}{2}\right].$$

Consider now the expression $\mathbb{E}_\mu\mu^2$, where the expectation is taken over the distribution of μ:

$$\mathbb{E}\mu^2 = Pr(\tau = C)\mu_C^2 + Pr(\tau = F)\mu_F^2$$

Using the fact that $Pr(\tau = C) = \frac{\lambda\gamma}{\mu_C}$ and $Pr(\tau = F) = \frac{\lambda(1-\gamma)}{\mu_F}$ I can rewrite the former expression to get:

$$\mathbb{E}\mu^2 = \lambda(\gamma\mu_C + (1-\gamma)\mu_F).$$

This is very useful to simplify the result since we can see that the last term is in fact equal to: $\frac{b}{2}\mathbb{E}\mu^2$. Doing this substitution and multiplying out all the terms, the welfare expression can finally be expressed as (1).

□

Since from the analysis of the equilibrium I get a condition on $\mu_C - \mu_F$, however, it is useful to know whether the gap between the good reputation from consistent play and the bad reputation from flip-flopping is sufficient to conclude something on selection welfare. The following lemma tackles this question:

Lemma 1. *Selection welfare increases when μ_C and the reputation gap $\mu_C - \mu_F$ increase.*

Proof. See Appendix. □

Lemma 1 will be helpful to prove some of the welfare results, since it tells us that whenever the reputation gap gets larger, knowing that the reputation from consistent play increased is enough to tell that selection welfare increased. As a result, let's now move on to stating how variations in the model parameters affect welfare.

Fact 4. *Increasing ϕ weakly decreases welfare.*

Proof. When ϕ increases, the benefits from office increase and this increases the accountability distortion. This happens because $\mu_C - \mu_F$ is an increasing function of σ; when ϕ increases, therefore, σ^* decreases up to the point where $\mu_C - \mu_F$ is equal to the new costs of deviating for the signal. This reasoning can be easily verified from the equilibrium condition, noting that the right-hand side increases as σ^* increases:

$$2\rho_2 - 1\phi = \frac{\lambda\gamma}{\lambda\gamma + (1-\lambda)(1 - A\sigma^*)} - \frac{\lambda(1-\gamma)}{\lambda(1-\gamma) + (1-\lambda)A\sigma^*}$$

The decrease in σ^* negatively impacts welfare in the second period. In terms of selection of politicians, since a lower σ^* decreases μ_C and increases μ_F, then from Lemma 1 we know that this means that the second moment will decrease, hence also selection welfare worsens. □

It has to be kept in mind, of course, that this model abstracts from all those reasons why it might be a good idea to offer electoral incentives to politicians (for example to improve the share of competent politicians in the pool); as a matter of fact, in this model politicians' interests are aligned with those of citizens except for electoral incentives.

Another conclusion from the analysis is that having a wide competence gap between the two politicians' types is bad for accountability: as a result, increasing q always increases welfare, both because politicians pander less (i.e. σ^* increases) and because incompetent politicians have better information when they choose a policy.

Fact 5. *As long as $\sigma^* < 1$, increasing q strictly increases welfare.*

Proof. Assume that q increases to q'. An increase in q moves $A = (1-\gamma)(q^2 + (1-q)^2) + 2\gamma q(1-q)$ down and $\frac{2\rho_2-1}{\phi}$ up. As long as both $\sigma^*(q)$ and $\sigma^*(q')$ are strictly less than 1, then in equilibrium $\frac{2\rho_2-1}{\phi} = \mu_C - \mu_F$ and therefore an increase of q to q' leads to a larger equilibrium value of $\mu_C - \mu_F$. Moreover, since a decrease in A lowers μ_C and raises μ_F ceteris paribus, σ^* has to increase in order to make $\mu_C - \mu_F$ larger. In particular, it has to be the case that $A\sigma^*$ increases. It follows that $\mu_C(q') > \mu_C(q)$, and as a result from Lemma 1 we know that selection welfare improves. In terms of accountability welfare, $\sigma^*(q') > \sigma^*(q)$ and $q' > q$, so not only incompetent politicians are better, but they also act in a less distorted way. Hence, accountability welfare and therefore total welfare increases. Notice that once in a truthful equilibrium, an increase in q decreases $\mu_C - \mu_F$, since A keeps decreasing but σ cannot increase any further. As a result, the second moment decreases and the probability of having a competent politician in office in the second period decreases. However, bad politicians are better so the effect on selection welfare is ambiguous. □

Notice that when q is high, i.e. politicians are in general competent, then μ_F is lower, i.e. flip-flopping hurts more. This might be one of the reasons why flip-flopping can hurt candidates in a race such as the US presidential election, despite the fact that electoral incentives are high. In other words, flip-flopping is worse for a candidate's reputation with more homogeneity between competent and incompetent candidates [9].

[9]Notice that a bad reputation in this setup is just the probability of being of the incompetent type, no matter how bad the incompetence is.

When λ increases, both μ_C and μ_F increase. However, $\mu_C - \mu_F$ can either increase or decrease. As a result, σ^* can also move in either direction. This means that there exist cases in which having a better pool of politicians increases the distortion generated by incompetent politicians avoiding flip-flops. In terms of welfare from the selection of politicians, however, an increase in λ is always beneficial, at least whenever the starting point is a partially truthful equilibrium. The reason is that the equilibrium level of $\mu_C - \mu_F$ remains constant. This means that both μ_C and μ_F increase, and from Lemma 1 we know that this means the selection of politicians improves.

Fact 6. *Increasing λ in a partially truthful equilibrium improves selection welfare, but it can either increase or decrease accountability welfare.*

Proof. Assume that the game has a partially truthful equilibrium. The change in λ will therefore lead to either a truthful equilibrium or the game will remain in a partially truthful equilibrium. In the first scenario, by definition the policy distortion decreases, because σ^* increases to 1. Moreover, given that $\mathbb{E}\mu$ and since in a truthful equilibrium (compared to a partially truthful one) $\mu_C - \mu_F$ is larger, then welfare improves, given that also μ_C increases for sure, given that $\mathbb{E}\mu$ has to increase.

Let's now consider the case in which the game remains in a partially truthful equilibrium. Accountability welfare can either increase or decrease. The reason is that, denoting by $D_C \equiv \lambda\gamma + (1-\lambda)(1-A\sigma)$ and $D_F = \lambda(1-\gamma) + (1-\lambda)A\sigma$, the expression $\frac{\partial(\mu_C - \mu_F))}{\partial\lambda} = \frac{A\sigma^*}{1-\gamma}\frac{1}{D_F^2} - \frac{1-A\sigma^*}{\gamma}\frac{1}{D_C^2}$ can move in both directions following an increase in λ. As a consequence, σ^* can either increase or decrease in order to move $\mu_C - \mu_F$ back to the equilibrium level. In other words, despite more competent politicians being there, it is possible for the increasingly distortive behaviour of incompetent politicians to decrease accountability welfare.

In terms of selection welfare, on the other hand, the average μ increases despite $\mu_C - \mu_F$ remaining constant. This means that either μ_C and μ_F increase, or that at least $Pr(\tau = C)$ has to increase. However, we can check using $Pr(\tau = C) = \frac{\lambda\gamma}{\mu_C}$ that μ_C has to increase, because otherwise both $Pr(\tau = C)$ and $Pr(\tau = F)$ would increase, which is a contradiction. It follows by Lemma 1 that selection welfare increases.

□

Things get more complicated when evaluating a change in the persistence parameter γ. With more persistence there are several possible cases: first of all, σ^* can be either increasing or decreasing in γ. In the former case, accountability improves with higher persistence, whereas in the latter it could go either way. In terms of selection, however, it can be shown that a more persistent state of the world decreases the effectiveness of elections in selecting competent politicians:

Fact 7. *The probability of selecting a competent politician through elections decreases as γ increases.*

Proof. When γ increases, the equilibrium value of $\mu_C - \mu_F$ decreases. From Lemma 1, we know that in such a situation selection only improves when both μ_C and μ_F increase. However, since it has to be that $Pr(\tau = C)\mu_C + Pr(\tau = F)\mu_F = \lambda$, then if both μ_C and μ_F were to increase, $Pr(\tau = F)$ would have to increase, too. However, this is not possible, because if γ increases and $Pr(\tau = F)$ also increases, then μ_F decreases. This means that selection always worsens when γ increases. □

This result is interesting since it tells us that elections perform worse in a less variable world. This can seem surprising, especially given that conditional on a leadership change, elected leaders are worse after a change in the state. However, the result is driven by the fact that whenever a bad leader is in office, a change in state increases the chances of having a competent one after the election.

Lastly, although I do not include politicians in the calculation of social welfare, it is nonetheless important to understand whether they are made better or worse off by their strategic behaviour.

Fact 8. *Incompetent politicians can be worse off in a partially truthful equilibrium in which they distort their actions compared to the truthful benchmark.*

Proof. Writing the utility of incompetent politicians under the two scenarios and using the envelope theorem we get to the following expression to evaluate:

$$(1-A)\left[\frac{\lambda\gamma}{\lambda\gamma+(1-\lambda)(1-A)} - \frac{\lambda\gamma}{\lambda\gamma+(1-\lambda)(1-A\sigma^*)}\right] + \\ -A\left[\frac{\lambda(1-\gamma)}{\lambda(1-\gamma)+(1-\lambda)A\sigma^*} - \frac{\lambda(1-\gamma)}{\lambda(1-\gamma)+(1-\lambda)A}\right]$$

Since the result can go in both directions, it is enough to provide two examples. First of all, let's consider the case in which incompetent politicians are better off under the equilibrium with distortions. This just requires a small enough γ: it can be checked numerically that if γ is sufficiently low, then for any $\sigma^* \in [0, 1)$, as well as λ and q, the expression is negative. In order to show that it is possible for the expression to be positive, check numerically that for sufficiently large γ and sufficiently low q and λ there is an interval $\sigma \in [\underline{s}, 1)$ in which the expression is strictly positive. Then take ϕ such that a truthful equilibrium is sustainable and increase it: since σ^* is continuous in ϕ, then for any $\hat{\sigma}$ in an ϵ-interval around 1 there exists a $\hat{\phi}$ such that $\sigma^*(\hat{\phi}) = \hat{\sigma}$. Since we know that the expression is positive at $\hat{\sigma}$, then there necessarily exist values of γ, q and λ such that the incompetent incumbent would be better off if he were able to commit to not distorting his actions. $\square$

Another way to describe this result is that an institution forbidding incumbents to distort their actions (which would always benefit society as a whole) would also be supported by incompetent politicians when they make up a large enough share of the political class, their information is very bad and the world is more stable. This last part is particularly interesting, since it might suggest that such a reform could be easier to achieve when the environment is more stable.

5 Institutional Design

So far I have shown how the bad reputation from flip-flopping can give (incompetent) politicians an incentive to distort their actions. In the baseline model I have presented, there are three fundamental ingredients leading to the result: the first is the fact that politicians face an election at the end of $t = 2$, because being re-elected gives them a utility of $\frac{\phi}{2}$; the second is the fact that voters have no information on what constitutes the appropriate policy in each period and are therefore only able to evaluate incumbents based on their track record in office; third, voters are perfectly able to observe the action taken by the politician, in such a way that a flip-flop is immediately caught and used to form reputations. My aim in this section is to relax these assumptions by introducing new institutional features to the baseline model.

5.1 Single Term Limit

In the baseline model I analyzed above, the implicit assumption I make is that the incumbent can serve up to two terms in office, whereas the challenger can at most serve once, since the game ends after the third period. As a result, it is interesting to see what would happen under a single term limit rule, in which all politicians can serve only one term in office [10]. The single term limit for the President is an institution in

[10]Notice that my baseline model is not exactly a model of a two-term limit economy; if the challenger was able to also serve two terms, then voters would expect him to be subject to electoral concerns in his first period in office, meaning that he could avoid flip-flopping. As a result, they will always prefer the incumbent when their reputations are the same: in other words, such a model would feature an endogenous incumbency advantage.

several Latin American countries [11], Israel and South Korea, as well as for the head of the European Central Bank, among others. In the setup I describe in this model, the single term limit is a blunt yet effective instrument to eliminate all distortions due to politicians' willingness to avoid flip-flopping. At the same time, however, having a single term limit also means forgoing the possibility to condition the reelection decision on the beliefs about the incumbent's type. It follows that banning reelections is only welfare improving if the accountability distortion from the flip-flop avoidance is large and the upside from retaining incumbents with a good reputation is low.

Proposition 2. *A single term limit is welfare improving if and only if:*

$$b \leq 2\frac{(1-\lambda)A(1-\sigma^*)(2\rho_2-1)}{\frac{1+\lambda_O^2+\mathbb{E}\mu^2}{2}+\lambda-\lambda_O-\lambda\lambda_O}$$

Intuitively, the condition states that the single-term limit rule is beneficial if the benefit from selecting a competent politician is low enough. Moreover, notice that if $\lambda \to 1$, the right-hand side goes to 0, meaning that if the incumbent is competent with a sufficiently high probability, even a very small benefit b is sufficient to prefer the re-electability of the incumbent. The same happens if $q \to \frac{1}{2}$ or as $\phi \to +\infty$, because in these cases lying will be so prevalent to offset most of the learning about the incumbent.

Finally, notice that implementing a single term limit rule requires the ability to commit (for example through a constitution) not to re-elect an incumbent thought to be competent with a high probability [12].

5.2 Transparency of Actions: Reporting Media

Voters usually rely on the media to learn the policies chosen by politicians. In some circumstances, for example when bills containing multiple prescriptions are voted [13], it is not so straightforward to understand whether the incumbent politician flip-flopped or played consistently. The same might happen when voting in a committee is secret and

[11] El Salvador, Mexico, Honduras and a few others.

[12] The idea of commitment to a single term limit is developped in a rather hyperbolic form by Calvino (1969) in a short story in which he describes a hypothetical society in which leaders are beheaded at the end of their term in office.

[13] An example of a situation in which it was not simple to label a policy choice as a flip-flop is the vote by Bernie Sanders against the auto-industry bailout in January 2009: Sanders had actually supported the bailout previously and supported it afterwards, but after having voted in favour of it, he voted against the release of that tranche of aid since it also contained financial aid for the banking sector, which Sanders was not in favour of bailing out. In the recent presidential primary election, Hillary Clinton used this alleged flip-flop to attack Sanders.

knowing the result only enables to make probabilistic statements on whether a member voted in a certain direction.

In this section I therefore relax the assumption of full observability of the incumbent's actions and evaluate its impact on social welfare. The fact that action transparency is not always beneficial is well known in the literature [14]. What I am going to show here is that an imperfectly accurate reporting media is always optimal (or, in terms of transparency, partial action transparency is always optimal), unless the equilibrium is truthful with a fully accurate media in the first place.

Assume that the incumbent's track record is only observable through the report of a media company, and that the media company's reporting technology is not perfectly accurate: given a true incumbent track record τ that displays flip-flopping, with probability $1-g$ the media company sends out a report indicating that the politician acted consistently: the same happens when the true track record is consistent [15]. For example, a voter observing a consistent track-record infers that the actual track-record of the incumbent is consistent with probability:

$$p_C = Pr(\tau = C|\tilde{\tau} = C) = \frac{gPr(\tau = C)}{gPr(\tau = C) + (1-g)Pr(\tau = F)}$$

while an analogous expression, denoted by p_F, represents the probability that the true track record is flip flopping given an observed flip-flopping record. Given her guess of the incumbent's strategy, the voter updates her beliefs and assigns a reputation to each true track-record. Since she does not observe the true track record, the actual reputation that she assigns to the politician is simply the weighted average of the reputations following each track-record, weighted by the probability that the observed track-record is of each type conditional on the observed media report. As a result, if $g = 1$ there is no noise (full transparency) in the reporting media and we recover the baseline model, whereas if $g = \frac{1}{2}$ there is no transparency and the reputation associated to each track-record is simply λ.

Since the lower g, the less voters can learn about the incumbent independently of his behaviour, it is intuitive that truthful play can be restored for g low enough. It is also clear that if this g, which I'll denote by g^*, is larger than $\frac{1}{2}$, then having $g < g^*$ is never optimal, because once truthful play has been restored, lowering g only makes learning worse. What is not straightforward is the fact that increasing g starting from g^* is never

[14]Prat (2005) again as the main example.

[15]Notice that I do not need to fully specify whether the noise comes from the report of the second or the first period action, or a mix of both: all I need is that noise is symmetric across the actions played by politicians.

optimal. The following proposition shows that the optimal reporting accuracy is always the maximum g sustaining truthful play.

Proposition 3. *The optimal media reporting accuracy is $g^* \leq 1$. The value of g^* is the largest possible such that incumbents play truthfully. Therefore, $g^* < 1$ whenever a truthful equilibrium is not sustainable in the baseline model.*

Proof. See Appendix. □

5.3 Delegating the First Action

An alternative way of thinking about transparency of the politician's action is to consider the possibility of delegating the first action to an independent agent. In many contexts in which multiple decisions have to be taken on an issue, as a matter of fact, the politician or top level official deos not take care of all the steps in the process, but only acts in some stages. One can think for example about a committee doing preliminary work before a bill is voted, or about a bureaucrat drafting a reform, or even at a local government experimenting on a policy before the central government adopts it. I will therefore consider a situation in which the first action is observable and taken by an independent agent, whereas the second action, also observable, is taken by the incumbent. This variation of the game also allows me to shed light on the relationship between the game of repeated policy decisions I described above and a one-action pandering game à la Prat (2005). Suppose that the utility citizens receive from the first action matching the state is now $\alpha > 0$ (which could be larger or smaller than one) instead of 1 as in the benchmark model. Moreover, for simplicity I assume that the first decision maker is incompetent with probability 1. [16] This means that the information an incompetent incumbent has before taking the second action is exactly the same as that he has in the benchmark model: this also means that the cost of ignoring her signal is the same. The competent type, on the other hand, has less initial information than in the benchmark model, since the first signal is less informative than the one a competent agent would have. As a result, the signal of the competent incumbent flips more often with respect to

[16] If the action were delegated to an agent with signal accuracy lower than q, the incentive to distort the action in the second period would be even smaller, but the cost of faking consistency would also decrease, so the effect for high levels of ϕ would be ambiguous. If instead the action were delegated to a more competent agent, as long as the signal of the incompetent incumbent is strong enough to dictate the optimal action in the second period, the case for delegation would be even more compelling. If the action was delegated to a very competent agent (and $\gamma > q$, then the incompetent second period incumbent should just repeat the same action independently of her signal, which would allow the voter to perfectly identify the competent incumbent whenever a flip-flop occurs. In this context, delegation can create opportunistic flip-flops.

the first agent than with respect to his own first signal. Therefore,flip-flopping [17] has a better reputation than in the benchmark model. If electoral concerns are high, a partially truthful equilibrium in both games has the same reputational wedge between consistency and flip-flopping, equal to $\frac{2\rho-1}{\phi}$. However, achieving the same wedge requires a smaller distortion in the game with delegation. This means that delegating leads to a gain in terms of second action welfare, but a loss in terms of first action welfare, whereas nothing changes in terms of electoral selection. It follows that if α is small enough, delegating to an incompetent agent is beneficial.

Proposition 4. *Suppose that a_1 was delegated to an agent drawn from a pool of incompetents. Suppose further that $\phi > \phi_D$, so that truthful equilibria are not sustainable in either arrangement. Then, if the welfare weight of the first action α is small enough, welfare in the game with the delegated first period action is strictly larger than in the game with repeated actions from the politician.*

This result is interesting not only because it provides a rationale for delegation in contexts of repeated decision, but also because it shows that if we consider a one shot pandering model as a reduced form for a repeated decision model, then we are likely to underestimate the distortive incentives of elections. [18]

Moreover, a corollary of the result is that delegation is more likely to be beneficial when electoral concerns are high rather than when they are low, since with no distortions it is always better not to delegate.

5.3.1 Delegating the First vs the Second Action

Suppose that instead of the first one, the second action can be delegated to an independent agent. In this case, the politician always follows his signal in the first period. As a result, the scope for welfare improvement is even larger than in the example with delegation of the first action. In other words, in this model whenever delegating the first action to an agent is welfare improving, delegating the second action to the agent and letting the politician take the first achieves an even better result [19].

[17] In this case, flip-flopping just means contradicting the action of the first agent

[18] Using a belief on the type taking the first action putting some weight on each of the two types, for example, always leads to underestimating the distortion of the second action.

[19] An interesting extension that we leave for future research concerns the possibility that a politician can choose the ability of the agent to which the action is delegated.

5.4 Transparency of Consequences: Commentator Media

So far I assumed that voters have no feedback on the state of the world before elections. This serves the purpose of creating an environment in which all learning about the incumbent is done through his track-record. In many situations, however, voters have some information about the right policy choice. A fundamental role in this respect is again played by the media.

In this section, therefore, I augment the baseline model with an additional player, which I call commentator media to follow up on the contribution by Ashworth and Shotts (2010). To keep the analysis as simple as possible I assume that the media is only active at the end of $t = 2$, i.e. just before elections: this also reflects the fact that the coverage of politics in the vicinity of elections is particularly salient. I assume that the media is endowed with a signalling technology of accuracy $q_M > \frac{1}{2}$ (conditionally independent of the signal that incumbents receive) and I abstract from strategic consideration on the part of the media assuming that before election, the media truthfully reveals the realization of its signal s_M.

What is going to happen is simply that voters now have another piece of information to use when updating their beliefs on the incumbent type: reputation will not only depend on whether the incumbent flip-flopped or played consistently, but also on whether the media report endorses his second period action. In such a setup it seems natural that having an informative signal on the state of the world will act as disciplining device for the incumbent. The main result of this section, however, is to show that things need not work this way.

Proposition 5. *Increasing the accuracy of the commentator media can increase flip-flopping avoidance.*

Proof. See Appendix. □

It is interesting to describe in some detail the circumstances in which an increase in media accuracy increases the distortion to accountability. What is needed is a highly persistent environment (high γ) and high level of competence of politicians (both high q and high λ). In this context, increasing q_M from the lower bound of $\frac{1}{2}$ is initially marginally beneficial, but as q_M increases the effect reverses and further increasing q_M (up to a point) increases the distortion. When q_M starts becoming very precise, however, there is a large benefit in further increases of q_M. In other words, in these environments a marginally informative media is slightly beneficial, an informative but not very precise media is detrimental but a very precise media is again beneficial.

The intuition behind this result is what we might define as politicians gambling on the endorsement by the media. The force driving the result is the fact that when λ, γ and q are large, then an increase in media accuracy q_M sharply increases the payoff value of a successful gamble, which is the difference in reputation between being consistent and endorsed by the media and being a flip-flopper and opposed by the media, denoted by $\mu_{C,E} - \mu_{F,O}$. On the other hand, the the payoff from a failed gamble, which is the difference in reputation between a consistent politician opposed by the media and a flip-flopper endorsed by the media, or $\mu_{C,O} - \mu_{F,E}$, decreases more slowly (but it suddenly drops when q_M is sufficiently high, meaning that when the media is very well informed, a non-endorsement is very costly in terms of reputation). Moreover, as long as the media is not too informative, the probability S that the media signal matches the politician's flip-flopping signal remains close to $\frac{1}{2}$, making it likely for a politician avoiding a flip-flop to gamble successfully. This result relies heavily on λ being high, which is what gives the necessary curvature to the payoff from gambling on avoiding the flip-flop: in other words, when most politicians are competent, increasing the informativeness of the media increases the value of an endorsement significantly.

A related result is that when persistence γ, the share of competent politicians λ and electoral concerns ϕ are all very high, it might be the case that incompetent politicians distort their behaviour even under a perfect media signal $q_M = 1$. In this case, flip-flopping is such a strong signal of incompetence that politicians are willing to take a gamble in which they lose the election for sure unless they receive a media endorsement.

The existence of potentially non-monotone responses of policy distortions to media accuracy can have interesting implications for issues such as the public subsidization of media outlets. In persistent environment with a prevalence of competent politicians, subsidizing media is only beneficial if very precise commentary is achieved. This could for example suggest that concentrating resources into one or few high-quality outlets is better than subsidizing many average ones. In particular, once the media is sufficiently informative, returns to small increases in informativeness can be very large: in other words, in these situations *the devil is in the details*.

Notice that when an increased accuracy of the commentator media increases flip-flopping avoidance, accountability welfare decreases but selection welfare might still improve, so I are not able to conclude that a more accurate commentator media is detrimental to welfare *tout court*.

5.5 Transparency of Consequences: Poll-Matching

In this section I analyze the game in which a public signal concerning the state of the world in the second period is released *before* the politician takes the second action. In this modified game, the incumbent has one additional piece of information before choosing the second action a_2. In particular, I assume that the signal m (which I will call the poll) has accuracy $Pr(m = \omega_2) = p = q$, i.e. the same accuracy as the signal of the incompetent incumbent, and that it is conditionally independent with respect to the signal of the incumbent. Moreover, unlike in the benchmark model, the accuracy of the signal of the competent incumbent is now $q_H = h < \bar{q}$. In this setup, the optimal action in the second period does not only depend on the private signal of the incumbent, but also on the poll. In particular, the truthful equilibrium benchmark does not involve the politician always following his private signal: as a matter of fact, if the poll confirms the first action of the politician (that is to say $m = s_1 = a_1$) the optimal action for both types is to match the poll independently of the private signal [20]. If the poll instead calls for a flip-flop, the optimal action depends on the private signal s_2. In this setup, flip-flopping signals the fact that the private signal matches the state, i.e. $s_2 = m$; since in this case the *smart action* is indeed to match the signal, flip-flopping to match the poll delivers a better reputation than not changing one's mind. As a result, politicians have an incentive to flip-flop whenever the poll calls for it, and in equilibrium incompetent politicians mix by flip-flopping to match the poll even when their signal calls for sticking to the previous action. I summarize this result in the following proposition.

Proposition 6. *If $h \in (p, \bar{h}]$ and $\phi > \bar{\phi}^m$, the partially truthful equilibrium of the game is such that $a_1 = s_1$ for both types and:*

$$a_2 = \left\{ \begin{array}{ll} m, & \text{if } m = a_1 \\ s_2, & \text{if } m \neq a_1 \end{array} \right\}$$

for the competent type, whereas when $m \neq s_1$ and $s_2 \neq m$, the incompetent politician plays $a_2 = m$ with probability $1 - \sigma_m^*$.

Proof. Let's start by analyzing the undistorted equilibrium (i.e. the analogue of the truthful equilibrium in this setup). In the second period, there are 4 possible situations: first of all, s_2 can either match or differ from the poll m. In the former case, the optimal choice is always to play $a_2 = s_2 = m$ independently of s_1 and the type. If however

[20]For this reason we need q_H to be sufficiently low, since if q_H is too large then the high type would choose $a_2 = s_2$ independently of m.

the poll and the private signal are conflicting, the first signal becomes pivotal for the decision: for the low type, the private signal and the poll cancel each other out and given the persistence of the state the first signal becomes pivotal. For the high type, instead, the assumption $p < q_H < \bar{q}$ guarantees that despite the fact that the private signal is stronger than the public signal, the prior prevails.

Let's now consider reputations. The voter now has an additional signal m, therefore the reputation depends not only on flip-flopping versus consistency, but also on whether the action matches the poll or not. Moreover, notice that since no politicial would flip-flop to contradict the poll, $\lambda_{F,K}$ is not defined by Bayes rule but is derived out-of equilibrium. I will assume that out-of-equilibrium beliefs are pessimistic enough to discourage a flip-flopping track-record when it doesn't match the poll.

Suppose that incumbents play using truthful strategies. If the poll signal matches the first action, the incumbent always plays $a_2 = m$ and reputation can be written as:

$$\lambda_{C,M} = \frac{\lambda C_H}{\lambda C_H + (1-\lambda) C_L}$$

where (for accuracy q; for the high type one simply needs to replace it with h)

$$C_q = \gamma[q^2 p + (1-q)^2(1-p) + q(1-q)] + (1-\gamma)[q^2(1-p) + (1-q)^2 p + q(1-q)].$$

If on the other hand the signal does not match the first action, the incumbent should optimally follow his private signal, and reputations can be written as:

$$\lambda_{F,M} = \frac{\lambda F_H}{\lambda F_H + (1-\lambda) F_L}$$

in the case of flip-flopping and matching the poll with the second action, and

$$\lambda_{C,K} = \frac{\lambda K_H}{\lambda K_H + (1-\lambda) K_L}$$

in the case of consistent play and not matching the poll with the second action.

where the following are defined (again for accuracy q, for the high type one simply needs to replace it with h):

$$F_q = \gamma[q(1-q)] + (1-\gamma)[q^2 p + (1-q)^2(1-p)]$$

and

$$K_q = \gamma[q^2(1-p) + (1-q)^2 p] + (1-\gamma)[q(1-q)].$$

Let's focus on the case in which the poll does not match the first action. If the incumbent flip-flops, that signals incompetence but at the same time having a signal that matches the poll is a sign of competence. If p is large enough, therefore, the second effect outweighs the first, and flip-flopping to match the poll gives a higher reputation than not flip-flopping but not matching the poll. In particular, it can be shown that some values of p, $\lambda_{F,M} > \lambda_{C,K}$. This requires:
$\frac{F_q}{F_h} < \frac{K_q}{K_h}$

As a matter of fact, the former equation can be rewritten as:

$$\frac{\gamma p(1-p) + (1-\gamma)(p^3 + (1-p)^3)}{\gamma(p^2(1-p) + (1-p)^2 p) + (1-\gamma)p(1-p)} < \frac{\gamma h(1-h) + (1-\gamma)(h^2 p + (1-h)^2(1-p))}{\gamma(h^2(1-p) + (1-h)^2 p) + (1-\gamma)h(1-h)}$$

and it can be checked with a solver that there exist $\underline{p} < g$ and $\bar{p} = h$ such that it is verified. for $p \in [\underline{p}, \bar{p}]$. Notice how this requires further that $g < h$, because if g is too high the bad reputational effect of flip-flopping is too high to be counterbalanced by matching the poll.

Therefore, when $s_2 = m = s_1 = a_1$ or when $s_2 \neq m = s_1 = a_1$, matching the poll is optimal both in terms of matching the state and in terms of reputation, since $\lambda_{C,M} > \lambda_{F,K}$. [21] When on the other hand $m \neq s_1 = a_1$, if $s_2 = m$ the choice of $a_2 = m$ is straightforward since it maximizes the probability of matching the state and $\lambda_{F,M} > \lambda_{C,K}$; however, if $m \neq s_1$ and $s_2 \neq m$, then following the private signal leads to the optimal state matching decision but $\lambda_{C,K}$, whereas posturing to match the poll is costly in terms of policy performance but gives $\lambda_{F,M} > \lambda_{C,K}$. As a result, the politician has a trade-off and the size of the electoral concerns ϕ determines whether an undistorted equilibrium is feasible or not. In a partially truthful equilibrium, the probability that the incompetent incumbent follows his signal when $m \neq a_1 = s_2$, denoted by σ_m, solves the following equation, in a very similar way to what happened in the benchmark model:

$$\frac{\lambda F_h}{\lambda F_h + (1-\lambda)(F_l + K_l(1-\sigma_m^*))} - \frac{\lambda K_h}{\lambda K_h + (1-\lambda)K_l\sigma_m^*} = \frac{2\rho_m - 1}{\phi}.$$

□

Compared to the benchmark model, in this partially truthful equilibrium a flip-flopping track-record can be the result of opportunistic posturing carried out by an incompetent politician to match the public poll. Unlike in the benchmark model, however,

[21] Notice that the reputation when the poll matches the first action is larger than when it does not: in other words, incumbents still have the incentives to follow their signal in the first period.

flip-flopping is now beneficial for a politician's reputation, given that it only happens to match an informative poll.

6 Predictability, Opportunism and Flip-flopping

The central aim of this paper is to provide a rationale of why voters dislike flip-flopping politicians. The explanation I offer in this paper has to do with signalling competence: since the information available to badly informed politicians is more likely to change, flip-flopping is associated with incompetence.

Whereas information is certainly a component to the reasons why flip-floppers are disliked, other explanations that are often brought forward in the political discourse have to do with predictability and opportunism. In this section I am going to show that the structure of my model can also be used to shed some light on these alternative explanations of flip-flopping.

Let's begin with predictability. One of the reasons voters dislike flip-floppers, as a matter of fact, is because flip-floppers are seen to be less predictable candidates, and predictability can be desirable when for example politicians lack commitment. In my model, when voters see a flip-flop, they not only update negatively on the politician's competence, but they also assign a higher probability to a further flip-flop happening in the future period. In other words, the model is consistent with the idea that flip-flopping politicians are unpredictable (although predictability has no value per se for the voter in my model).

Proposition 7. *Politicians who flip-flop before elections are also more likely to flip-flop again if re-elected.*

Proof. See Appendix. □

What about opportunism? My model is a model of opportunistic avoidance of flip-flops rather than one of opportunistic flip-flops, and there is no readily available relabelling of the model to fully catch that feature. However, if the signal politicians receive in my model is interpreted as a signal concerning the opinion of the majority, then we can reinterpret the model as saying that flip-floppers have a bad reputation because they are more likely to be unskilled at understanding the opinion of majority of voters. In a world in which all politicians are opportunists, in other words, flip-floppers are more likely to be bad opportunists.

In future research I would like to address the question of opportunistic flip-flops more in detail.

7 Conclusion

This paper describes circumstances in which career concerned politicians have the incentive to inefficiently stick to their previous policy positions in order to avoid the reputational stigma of flip-flopping. This happens because voters are aware that policy shifts are more likely to be performed by incompetent leaders. The incentive to avoid efficient policy shifts damages voters' welfare both in terms of policy effectiveness and selection of competent candidates through elections.

In other words, this paper rationalizes the conventional wisdom that flip-flopping is bad for the reputation of a politician; however, my results also suggests that the level of flip-flopping delivered by electoral competition might be excessively low and that democracy could benefit from politicians being more willing to change their mind.

An additional interesting implication of the mechanism described in the paper is that changes in the fundamentals driving policy choices are likely to bring to leadership change, but politicians substituting an incumbent after a change of state has occurred are on average less competent than politicians entering power in stable times.

To sum up, my analysis provides a new example of how electoral competition can decrease welfare: elections enable voters to retain the politicians they believe to be competent, but at the same time they give potentially distortive incentives to politicians. In this context, therefore, it might be optimal not to subject politicians to electoral incentives, by either setting up a single term limit rule or handing decision power to a judiciary.

Another implication of my model concerns the observability of politicians' actions: making some of the actions taken by the politician unobservable (or observable with noise) can eliminate the incentives for inefficient policy choices. One interpretation of this result can be that making some parts of a policy-making process secret might be beneficial (one can think of secret voting on parliaments, closed-doors committee meetings, or non-disclosure of the early stages of a policy process). Another interpretation has to do with the media environment and the fact technological developments such as social media have made it very easy for voters to know what a politician previously did or stated: according to my model, this might lead to distortions and decrease the responsiveness of policies to information.

A similar trade-off between noise and distortion arises when the first action is observable but delegated to an independent agent: even if the agent is less competent than the incumbent, delegation might be beneficial by reducing the incentives to distort the second action.

Finally, I also show that when media act as commentators of the quality of a policy choice, an increase in the accuracy of the media signal can in some circumstances incentivize rather than deter the distortive behaviour of politicians. In other words, giving voters more information to evaluate politicians' track records can backfire.

References

Aghion, P., and M. O. Jackson (2016): "Inducing Leaders to Take Risky Decisions: Dismissal, Tenure, and Term Limits," *American Economic Journal: Microeconomics*, 8(3), 1–38.

Agranov, M. (2016): "Flip-Flopping, Primary Visibility, and the Selection of Candidates," *American Economic Journal: Microeconomics*, 8(2), 61–85.

Ashworth, S., and K. W. Shotts (2010): "Does informative media commentary reduce politicians' incentives to pander?," *Journal of Public Economics*, 94(11-12), 838–847.

Calvino, I. (1969): *Prima che tu dica pronto*. Mondadori.

campaign for freedom of information (1997): "Freedom of Information: Key Issues," .

Canes-Wrone, B., M. C. Herron, and K. W. Shotts (2001): "Leadership and Pandering: A Theory of Executive Policymaking," *American Journal of Political Science*, 45(3), 532–550.

Dewan, T., and R. Hortala-Vallve (2014): "Electoral Competition, Control and Learning," *Working Paper*, pp. 1–31.

Doherty, D., C. M. Dowling, and M. G. Miller (2015): "When is Changing Policy Positions Costly for Politicians? Experimental Evidence," *Political Behavior*, 38(2), 455–484.

Fearon, J. D. (1999): "Electoral Accountability and the Control of Politicians: Selecting Good Types versus Sanctioning Poor Performance," in *Democracy, Accountability and Representation*, ed. by A. Przeworski, and S. C. Stokes, chap. 2, pp. 55–97. Cambridge University Press.

Frisell, L. (2009): "A theory of self-fulfilling political expectations," *Journal of Public Economics*, 93(5–6), 715 – 720.

Fu, Q., and M. Li (2014): "Reputation-concerned policy makers and institutional status quo bias," *Journal of Public Economics*, 110, 15–25.

Gentzkow, M., and J. M. Shapiro (2006): "Media Bias and Reputation," *Journal of Political Economy*, 114(2), 280–316.

Hummel, P. (2010): "Flip-flopping from primaries to general elections," *Journal of Public Economics*, 94(11–12), 1020 – 1027.

Levendusky, M. S., and M. Horowitz (2012): "When Backing Down Is the Right Decision: Partisanship, New Information and Audience Costs," *Journal of Politics*, 74, 323–338.

Levy, G. (2004): "Anti-herding and strategic consultation," *European Economic Review*, 48(3), 503–525.

——— (2007): "Decision Making in Committees: Transparency, Reputation, and Voting Rules," *The American Economic Review*, 97(1), 150–168.

Li, W. (2007): "Changing One's Mind When the Facts Change: Incentives of Experts and the Design of Reporting Protocols," *The Review of Economic Studies*, 74(4), 1175–1194.

Majumdar, S., and S. W. Mukand (2004): "Policy Gambles," *American Economic Review*, 94(4), 1207–1222.

Maskin, E. S., and J. Tirole (2004): "The Politician and the Judge: Accountability in Government," *American Economic Review*, 94(4), 1034–1054.

Patacconi, A., and N. Vikander (2015): "A model of public opinion management," *Journal of Public Economics*, 128, 73 – 83.

Prat, A. (2005): "The Wrong Kind of Transparency," *The American Economic Review*, 95(3), 862–877.

Prendergast, C., and L. Stole (1996): "Impetuous Youngsters and Jaded Old-Timers: Acquiring a Reputation for Learning," *The Journal of Political Economy*, 104(6), 1105–1134.

Tavits, M. (2007): "Principle vs. pragmatism: Policy shifts and political competition," *American Journal of Political Science*, 51(1), 151–165.

Tomz, B. M., and R. V. Houweling (2012): "Political Repositioning," *Working Paper*, pp. 1–34.

Tomz, M., and R. P. Van Houweling (2012): "Candidate Repositioning," pp. 1–43.

A Proofs

Proposition 1

Proof. Let's consider the choice of action at $t = 2$, politicians have already taken action a_1 and they know that if re-elected, which happens with probability $\frac{1}{2} + \frac{\mu(a_1,a_2)-\lambda_O}{2}$, they will get 2ϕ. On top of that, politicians get utility of 1 whenever they match the state of the world, the probability of which depends on the posterior belief, denoted by $Pr(\omega_2 = s_2|s_1, q_\theta) = \rho_2(s_2, s_1, q_\theta)$. In order to slightly simplify notation, denote the probability of winning given reputation μ by $r(\mu)$. So $r(\mu) = \frac{1}{2} + \frac{\mu-\lambda_O}{2}$ as $r(\mu)$. Given that there are only two actions available, the politicians will calculate the reputation associated to each action and follow his signal if and only if:

$$\rho_2(s_2, s_1, q_\theta) + r(\mu(a_2 = s_2, s_1))2\phi \geq [1 - \rho_2(s_2, s_1, q_\theta)] + r(\mu(a_2 \neq s_2, s_1))2\phi.$$

It follows that in order to have $\sigma(s_2) = 1$, the above condition needs to hold for both types and both signal realizations given any of the two possible choices a_1, which results in a set of 8 inequalities.

However, it is immediate to notice that whenever $a_2 = s_2$ is the most reputable action, i.e. $\mu(a_2 = a_1, a_1) \geq \mu(a_2 \neq a_1, a_1)$, following the signal is unquestionably optimal. Since consistency has a better reputation than flip-flopping, it follows that whenever $s_2 = a_1$, $a_2 = s_2$, for each a_1 and each type. This means that we are left with 4 conditions. Moreover, since $\mu^T(0, 0) = \mu^T(1, 1)$ and $\mu^T(1, 0) = \mu^T(0, 1)$ and the same holds for $\rho_2(1, 0) = \rho_2(0, 1)$, we are left with only two conditions, one for each type:

$$\rho_2(s_2 \neq s_1, q_\theta) + r(\mu^T_F)2\phi \geq [1 - \rho_2(s_2 \neq s_1, q_\theta)] + r(\mu^T_C)2\phi,$$

which can be rearranged to:

$$\frac{2\rho_2(s_2 \neq s_1, q_\theta) - 1}{\phi} \geq \mu^T_C - \mu^T_F.$$

Thanks to the single-crossing property, the binding constraint for a truthful equilibrium is the condition concerning the incompetent politician: whenever the incompetent politician follows his signal, or is at least indifferent, the competent politician does, too. Conversely, if the competent politician is indifferent or he doesn't follow his signal, the incompetent will also not follow it.

As a result, we know that if $\frac{2\rho_2(s_2\neq s_1,q)-1}{\phi} \geq \mu^T_C - \mu^T_F$, politicians will follow their signal at $t = 2$. Rearranging we get the condition on ϕ, $\phi \leq \frac{2\rho_2(s_2\neq s_1,q)-1}{\mu^T_C-\mu^T_F}$.

Let's now see what happens at $t = 1$. If politicians know that they are going to follow their signal at $t = 2$, then the dominant strategy at $t = 1$ is to follow their signal. Since all that matters is being consistent versus flip-flopping, then given the persistence of the state, it is more likely to end up in the favourable situation of playing consistently by following one's signal in the first period.

□

Trembling-hand perfection refinement eliminates pooling equilibria

Proof. In a pooling equilibrium, both politicians play the same track-record with probability one independent of the signals received. In other words, in these equilibria an action in every period remains off-equilibrium, hence Bayes rule cannot restrict beliefs on these actions in any way. It follows that if the reputation attached to any off equilibrium track-record is sufficiently bad, then when electoral concerns are high enough politicians will not have any incentive to deviate from the pooling track-record. Sufficient conditions for any pooling equilibrium to be sustainable are that:

$$2\rho_1 - 1 + Pr(s_2 = s_1|s_1)(2\bar{\rho}_2 - 1) < (r(\mu = \lambda) - r(\mu = 0))2\phi$$

and

$$2\bar{\rho}_2 - 1 < r(\mu = \lambda) - r(\mu = 0))2\phi$$

Let's now introduce the trembling-hand perfection requirement. Assume that with probability $\epsilon > 0$ close to zero, a politician willing to play action a will instead play action a'. Assume that ϵ is the same for both types. Take a pooling equilibrium. In any period, with probability ϵ the voter observes an action different from that on the pooling track-record. Since both politicians have the same strategy and ϵ is the same for both types, then the reputation the voter must attach to actions outside the pooling track-record has to be λ, the same as the reputation of the pooling track-record. However, this cannot happen in equilibrium, because if the reputation of any track-record is the same, then incumbents have an incentive to always follow their signal. In other words, no pooling equilibrium can survive the trembling hand perfection requirement. □

Proof of Theorem 1

Proof. I split the proof in several parts. First of all I characterize the symmetric partially truthful equilibrium.

Claim 1: Existence of Partially Truthful Equilibrium

In this equilibrium, $\mu(0,0) = \mu(1,1) \equiv \mu_C$ and $\mu(0,1) = \mu(1,0) \equiv \mu_F$ by symmetry and $\mu_C > \mu_F$. Moreover, assume for now that incumbents always follows their signal at $t = 1$. Consider now period $t = 2$. After following their signal in period 1, in period 2 the incumbent has to decide whether to follow his signal or not. When $s_2 = a_1$, the incumbent always follows his signal, because:

$$\rho_2(s_1, s_2 = s_1, \theta) + r(\mu_C)2\phi > (1 - \rho_2(s_1, s_2 = s_1, \theta) + r(\mu_F)2\phi$$

since $\mu_C > \mu_F$ insures that $r(\mu_C) > r(\mu_F)$ and $\rho_2 > 1 - \rho_2$ by the decision relevance of signals. However, when $s_2 \neq a_1$, the incumbent has a tradeoff. From Proposition 1, a truthful equilibrium requires that:

$$\rho_2(s_1, s_2 \neq s_1, L) + r(\mu_F)2\phi \geq (1 - \rho_2(s_1, s_2 \neq s_1, L)) + r(\mu_C)2\phi,$$

from which the upper bound $\bar{\phi}$ was derived. Moreover, we know from Fact 2 (single-crossing property), that the following holds:

$$\begin{aligned}
&\rho_2(s_1, s_2 \neq s_1, L) + r(\mu'_F) \geq 1 - \rho_2(s_1, s_2 \neq s_1, L) + r(\mu'_C) \Rightarrow \\
&\rho_2(s_1, s_2 \neq s_1, H) + r(\mu'_F) > 1 - \rho_2(s_1, s_2 \neq s_1, H) + r(\mu'_C)
\end{aligned}$$

and

$$\begin{aligned}
&\rho_2(s_1, s_2 \neq s_1, H) + r(\mu'_F) \leq 1 - \rho_2(s_1, s_2 \neq s_1, H) + r(\mu'_C) \Rightarrow \\
&\rho_2(s_1, s_2 \neq s_1, L) + r(\mu'_F) > 1 - \rho_2(s_1, s_2 \neq s_1, L) + r(\mu'_C)
\end{aligned}$$

This means that we are left with three possibilities: both politicians play $a_2 \neq s_2$ when that involves flip-flopping, or the high type mixes between $a_2 = s_2$ and $a_2 = s_2$, or the high type always plays $a_2 = s_2$ and the low type mixes. I will now prove that the first two cannot be part of an equilibrium. Assume that both politicians play $a_2 \neq s_2$ when $s_2 \neq a_1$. Then, nobody would flip-flop and therefore, in a candidate equilibrium $\mu_C = \mu_F$; in this case, however, the optimal strategy for the incumbent is to follow his signal. Consider now the other case, i.e. that when $s_2 \neq a_1$, the high type mixes between $a_2 = s_2$ and $a_2 \neq s_2$ while the low type always plays $a_2 \neq s_2$. If this were an equilibrium, then flip-flopping would reveal the high type, and therefore $\mu_F = 1 > \mu_C$. In such a situation, however, the incumbent would always flip-flop. It follows that the only possibility when $s_2 \neq a_1$ is that the high type follows his signal whereas the low type mixes between the two actions. The low type mixes when the following holds:

$$\frac{2\rho_2(s_1, s_2 \neq s_1, L) - 1}{\phi} = \underbrace{\frac{\lambda\gamma}{\lambda\gamma + (1-\lambda)(1-A\sigma^*)}}_{\mu_C} - \underbrace{\frac{\lambda(1-\gamma)}{\lambda(1-\gamma) + (1-\lambda)A\sigma^*}}_{\mu_F} \tag{2}$$

In order to show existence and uniqueness of such an equilibrium, consider $\sigma^* \in [0, 1]$. It holds from Proposition 1 that at $\sigma^* = 1$, $\frac{2\rho_2(s_1,s_2\neq s_1,L)-1}{\phi} < \mu_c - \mu_F$. At the same time, at $\sigma^* = 0$ it has to be that $\frac{2\rho_2(s_1,s_2\neq s_1,L)-1}{\phi} > \mu_c - \mu_F$, since in that case $\mu_F = 1$. By continuity of $\mu_C - \mu_F$, an equilibrium with $\sigma^* \in (0, 1)$ exists. Moreover, notice that μ_C is strictly increasing in σ^* while μ_F is strictly decreasing in σ^*, and therefore the σ^* such that (2) is unique. Finally, consider the action choice at $t = 1$. It holds that $\mu_C > \mu_F$ and the incumbent knows he is going to follow his signal when $s_2 = s_1$. Moreover, the incumbent knows that, because of the persistence of the state of the world, given s_1 it is more likely for him to receive $s_2 = s_1$. However, the utility of receiving $s_2 = s_1$ is higher when $a_1 = s_1$ than when $a_1 \neq s_1$. As a result, following the signal at $t = 1$ is optimal both in terms of instantaneous payoff and in terms of future payoff. Mathematically, denote by $\pi(\theta) = Pr(s_2 = s_1|s_1)$. Notice that since $\gamma > \frac{1}{2}$, $\pi > \frac{1}{2}$ and notice that $\pi(\theta)\rho_2(s_1, s_2 = s_1, \theta) + (1 - \pi(\theta))\rho_2(s_1, s_2 \neq s_1, \theta) = \rho_1(s_1, \theta)\gamma + (1 - \rho_1(s_1, \theta))(1 - \gamma)$. Denote also by $\bar{\rho}_2 = \rho_2(s_1, s_2 = s_1, \theta)$ and by $\underline{\rho}_2 = \rho_2(s_1, s_2 \neq s_1, \theta)$. With this notation, following one's signal at t_1 requires the following condition to hold:

$$\rho_1 + \pi(\bar{\rho}_2 + r_C 2\phi) + (1-\pi)(\underline{\rho}_2 + r_F 2\phi) \geq (1-\rho_1) + \pi(\bar{\rho}_2 + r_F 2\phi) + (1-\pi)(\underline{\rho}_2 + r_C 2\phi)$$

which can be rearranged to

$$(2\rho_1 - 1) \geq -(2\pi - 1)(r_C - r_F)2\phi$$

which always holds. Hence both types follow their signal at $t = 1$.

To sum up, I have showed the existence of a unique symmetric partially truthful equilibrium (it would be truthful if $\phi \leq \bar{\phi}$) with and $\mu_C > \mu_F$. Both politicians follow their signal at $t = 1$. At $t = 2$, the high type always follows his signal, whereas the low type follows his signal with probability 1 when if $s_2 = s_1$ and mixes playing $a_2 = s_2$ with probability $\sigma^* \in (0,1)$ when $s_2 \neq s_1$.

Claim 2: Uniqueness of Symmetric Non-Pooling Equilibrium

In Claim 1 I have shown that if $\mu_C > \mu_F$, then the equilibrium is the partially truthful one I characterized (or truthful, depdending on ϕ). I will now prove that remaining in the realm of symmetric equilibria, there exists no equilibrium where $\mu_F > \mu_C$. By claim 1, if $a_1 = s_1$ for both politicians the unique equilibrium is the (partially) truthful one. Let's then assume that $(2\rho_1 - 1) < (2\pi - 1)(r_F - r_C)2\phi$ for some incumbent type. Assume that the incumbent played an action $a_1 \neq s_1$. Since $\mu_F > \mu_C$, at $t = 2$ the optimal choice is to follow the signal when $s_2 = s_1$ whereas a tradeoff arises when $s_2 \neq s_1$. However, notice that $\frac{2\rho_1 - 1}{2\pi - 1} > 2\rho_2 - 1$, because $\pi = \gamma q + (1-\gamma)(1-q) < q$ given that $q > \frac{1}{2}$. Therefore, whenever the incumbent plays $a_1 \neq s_1$, at $t = 2$ he would always play $a_2 \neq a_1$. However, this cannot be part of an equilibrium. If the incumbent knows he is always going to flip flop, as a matter of fact, then the optimal strategy is to follow the first signal and then deviate if $s_2 = s_1$. The reason is that the such a strategy gives the incumbent the same reputation but it is less expensive in terms of policy costs, since $\pi(2\bar{\rho}_2 - 1) < (2\rho_1 - 1) + (1-\pi)(2\underline{\rho}_2 - 1)$. This can in fact be rearranged to yield:

$$[2(\rho_1\gamma + (1-\rho_1)(1-\gamma)) - 1] - (1-\pi)(2\underline{\rho}_2 - 1) < 2\rho_1 - 1 + (1-\pi)(2\underline{\rho}_2 - 1).$$

Claim 3: Non-Existence of Asymmetric Equilibria

I prove this result in several substeps.
Step 1: If $a_1 = s_1$, the equilibrium in the subgame starting at $t = 2$ is either truthful or partially truthful.

Proof. This was proved in Claim 1. □

Step 2: If for some type $a_2 = s_2$ with strictly positive probability, then for that type $a_1 = s_1$.

Proof. Denote by α and β the gap between the consistent and flip-flopping reputation after action $a_1 = 1$ and $a_1 = 0$ respectively. Moreover, denote by $c_1 = \frac{2\rho_1 - 1}{\phi}$. Denote by

$\pi = Prob(s_2 = s_1|s_1)$. If some type contradicts the signal at $t = 1$ with some probability after both signal realizations, and then plays $a_2 = s_2$ with some positive probability at $t = 2$, the following has to hold:

$$\mu(1,0) - \mu(0,1) = c_1 + \pi[\mu(0,0) - \mu(0,1)] - (1-\pi)[\mu(1,1) - \mu(1,0)]$$
$$\mu(0,1) - \mu(1,0) = c_1 + \pi[\mu(1,1) - \mu(1,0)] - (1-\pi)[\mu(0,0) - \mu(0,1)]$$
$$\mu(0,0) - \mu(0,1) = \beta$$
$$\mu(1,1) - \mu(1,0) = \alpha$$

Substituting the last two equations into the first two, and then the second into the first, we obtain the following condition:

$$c_1 = -(2\pi - 1)\frac{\alpha + \beta}{2}.$$

Now, whenever negative, $\frac{\alpha+\beta}{2} \in [-1, 0]$. However, $\frac{c_1}{2\pi-1} > 1$ since $\pi < \rho_1$. So this equality can never hold. This means that it is never the case that an incumbent contradicts the signal with some probability after both realizations of s_1.

Consider the case of a potential equilibrium in which in the first period, the incompetent politician mixes after receiving one of the signals. Let's consider, without loss of generality, a candidate equilibrium in which $Pr(a_1 = 1|s_1 = 1) = \sigma_1 < 1$. When the voter observes $a_1 = 1$, however, she knows that it is a genuine realization of $s_1 = 1$. As a result, the game play after $a_1 = 1$ is the same as that in the partially truthful equilibrium. Hence, $\mu(1,1) > \mu(1,0)$ and $\mu(1,1) - \mu(1,0) \leq c_2$. There are now therefore two cases: $\mu(0,0) > \mu(1,0)$ and $\mu(0,0) < \mu(1,0)$. If $\mu(0,0) > \mu(0,1)$, then since we are considering equilibria in which there is at least partialy truthful play at $t = 2$, it has to be that $\mu(0,0) - \mu(0,1) \leq c_2$. Since the average reputation after $a_1 = 1$ is larger than after $a_1 = 0$, because just incompetents are deviating, then it has to be that $\mu(1,1) \geq \mu(0,1)$. However, in order for the incumbent to have the incentive to deviate from $s_1 = 1$ it has to be the case that $c_1 + (\pi\mu(1,1) + (1-\pi)\mu(1,0)) \leq \pi\mu(0,1) + (1-\pi)\mu(0,0)$. However, $\pi\mu(0,1) + (1-\pi)\mu(0,0) - (\pi\mu(1,1) + (1-\pi)\mu(1,0))$ can be at most $2(1-\pi)c_2$ without the condition that $\mu(1,1) \geq \mu(0,1)$ being violated. Since $c_1 > 2(1-\pi)c_2$, then this distance is not enough to give the politician the incentive to deviate after $s_1 = 1$.

If $\mu(0,0) > \mu(0,1)$, and $\mu(0,0) - \mu(0,1) \leq c_2$, consider that since the average reputation after $a_1 = 1$ is larger than after $a_1 = 0$, because just incompetents are deviating, then it has to be that $\mu(1,1) \geq \mu(0,1)$. However, in order for the incumbent to have the incentive to deviate from $s_1 = 1$ it has to be the case that $c_1 + (\pi\mu(1,1) + (1-\pi)\mu(1,0)) \leq$

$\pi\mu(0,1)+(1-\pi)\mu(0,0)$. However, $\pi\mu(0,1)+(1-\pi)\mu(0,0)-(\pi\mu(1,1)+(1-\pi)\mu(1,0))$ can be at most $2(1-\pi)c_2$ without the condition that $\mu(1,1)\geq\mu(0,1)$ being violated. Since $c_1>2(1-\pi)c_2$, then this distance is not enough to give the politician the incentive to deviate after $s_1=1$.

Let's now consider the case in which $\mu(0,1)>\mu(0,0)$. If $\mu(0,1)-\mu(0,0)\leq c_2$, then we can immediately see that $\pi\mu(0,1)+(1-\pi)\mu(0,0)-(\pi\mu(1,1)+(1-\pi)\mu(1,0))\leq c_2$, but since $c_2<c_1$, again the incumbent cannot have the incentive to deviate without $\mu(0,0)>\mu(1,1)$; however, if $\mu(0,0)>\mu(1,1)$ the reputation after action $a_1=0$ is strictly larger than after $a_1=1$, which is a contradiction.

□

Step 3: there is no equilibrium in which some type does not follow the signal at $t=1$ with probability 1 and then contradicts one realization of the signal at $t=2$.

Proof. Impossible for a high type to contradict the signal at $t=2$, unless that is a pooling equilibrium (which we remove). Therefore, assume a low type contradicted s_1. Without loss of generality, assume that a low type who received signal $s_1=0$ played instead $a_1=1$. If $\mu(1,1)>\mu(1,0)$, then there are two possibilities: if action $a_1=1$ is only played by high types and low types who received a signal $s_1=0$, then contradicting signal $s_2=0$ would lead to $\tau=(1,0)$ to be revealing of the high type, which cannot happen in equilibrium. If $a_1=1$ were also to be played by low types receiving signal $s_1=1$, then if the low types who received $s_1=0$ contradict their signal $s_2=0$, also the low types receiving $s_2=0$ after having received $s_1=1$ and played $a_1=1$ would do that. And therefore $\tau=(1,0)$ would be revealing of the high type, which cannot hapen in equilibrium. So contradicting the signal s_10 at $t=1$ requires $\mu(1,0)>\mu(1,1)$. Since the high type plays truthfully at $t=2$, however, conditional on $a_1=1$ he plays $a_2=1$ more often than $a_2=0$. Therefore, $\mu(1,0)$ cannot be larger than $\mu(1,1)$.

□

□

Proof of Lemma 1

Proof. The average reputation is constant in equilibrium and equal to λ, hence $Pr(\tau=C)\mu_C+Pr(\tau=F)\mu_F=\lambda$. This allows us to rewrite $Pr(\tau=C)=\frac{\lambda-\mu_F}{\mu_C-\mu_F}$. From the model we know that selection improves if and only if the second moment m_2 of the distribution of reputation, $m_2=Pr(\tau=C)\mu_C^2+(1-Pr(\tau=C))\mu_F^2$, increases. Moreover, in equilibrium it holds that $\mu_C \geq \mu_F$ and $Pr(\tau = C) > Pr(\tau = F)$. Taking the expression for the second moment and substituting in the restriction due to the constant mean yields:

$$m_2 = \frac{\lambda - \mu_F}{\mu_C - \mu_F}\mu_C^2 + [1 - \frac{\lambda - \mu_F}{\mu_C - \mu_F}]\mu_F^2$$

Let's now write the expression of a contour line, in the (μ_C, μ_F) plane, along which the value of the second moment is constant; using implicit differentiation one gets:

$$\frac{\partial \mu_F}{\partial \mu_C} = \frac{\lambda - \mu_F}{\mu_C - \lambda}$$

It can be seen that this contour line is increasing in μ_C and concave. This means that the second moment increases by increasing μ_C and decreasing μ_F. Moreover, notice that since in equilibrium $\mu_C > \mu_F$, $Pr(\mu_C) > Pr(\mu_F)$ and $Pr(\mu_C)\mu_C + Pr(\mu_F)\mu_F = \lambda$, then $\mu_C - \lambda < \lambda - \mu_F$. So $\frac{\partial \mu_F}{\partial \mu_C} \geq 1$ in equilibrium. As a result, therefore, whenever $\mu_C - \mu_F$ increases and μ_C increases, the second moment m_2 has to increase and with it also selection welfare increases. Thus can be immediately be seen graphically, since the contour line of the second moment is increasing and concave and the line along which $\mu_C - \mu_F$ is constant is increasing with a slope of 1. In other words, knowing the size of $\mu_C - \mu_F$ is sufficient to compare the selection welfare across different equilibria. $\square$

Proof of Proposition 2

Proof. The welfare expression (1) derived in the text reads:

$$W = [\lambda + (1-\lambda)q] + [\lambda + (1-\lambda)\tilde{q}] + b\left[\frac{1}{4} + \frac{\lambda_O^2}{4} - \frac{\lambda\lambda_O}{2} + \frac{\lambda + \lambda_O}{2} + \frac{\mathbb{E}\mu^2}{4}\right]$$

With a single term limit, politicians can do no better than following their signal, since there is no re-election possibility. This generates a gain of

$$(1-\lambda)A(1-\sigma^*)(2\rho_2 - 1)$$

At the same time, however, a single term limit corresponds to a commitment to choosing the challenger no matter what the belief about the incumbent is. Therefore, society gets the benefit b with probability λ_O, plus $\mathbb{E}v = 0$. The utility from the selection of the right type of politician is therefore $\lambda_O b$ rather than $b\left[\frac{1}{4} + \frac{\lambda_O^2}{4} - \frac{\lambda\lambda_O}{2} + \frac{\lambda+\lambda_O}{2} + \frac{\mathbb{E}\mu^2}{4}\right]$ This generates a loss in terms of selection welfare expressed by:

$$b\left(\frac{1}{4}+\frac{\lambda_O^2}{4}-\frac{\lambda\lambda_O}{2}+\frac{\lambda-\lambda_O}{2}+\frac{\mathbb{E}\mu^2}{4}\right)$$

It follows that having a single term limit is beneficial if the following inequality holds:

$$(1-\lambda)A(1-\sigma^*)(2\rho_2-1)\geq b\left(\frac{1}{4}+\frac{\lambda_O^2}{4}-\frac{\lambda\lambda_O}{2}+\frac{\lambda-\lambda_O}{2}+\frac{\mathbb{E}\mu^2}{4}\right)$$

□

Lemma 3: p_C+p_F increasing in σ

Proof. $p_C=\frac{gPr(\tau=C)}{gPr(\tau=C)+(1-g)Pr(\tau=F)}$ and $p_F=\frac{gPr(\tau=F)}{gPr(\tau=F)+(1-g)Pr(\tau=C)}$. Now, remember that $\mu_C=\frac{\lambda\gamma}{\lambda\gamma+(1-\lambda)(1-A\sigma)}=\frac{\lambda\gamma}{Pr(\tau=C)}$ and similarly, $\mu_F=\frac{\lambda(1-\gamma)}{Pr(\tau=F)}$. Substituting these into the expression for p_C and p_F one gets the following expressions:

$$p_C=\frac{g\gamma\mu_F}{g\gamma\mu_F+(1-g)(1-\gamma)\mu_C}$$

and

$$p_F=\frac{g(1-\gamma)\mu_C}{g(1-\gamma)\mu_C+(1-g)(1-\gamma)\mu_F}.$$

Let's now denote by D_C and D_F the denominators of μ_C and μ_F respectively and by Den_C and Den_F the denominator of p_C and p_F respectively. First of all, let's establish that $Den_F<Den_C$. This clearly holds since $g(1-\gamma)\mu_C+(1-g)\gamma\mu_F<g\gamma\mu_F+(1-g)(1-\gamma)\mu_C$ can be rearranged to yield $D_C>D_F$, which is always satisfied. Let's now look at the numerator of the derivatives of p_C and p_F with respect to σ. The former is:

$$-\frac{g\gamma(1-\gamma)\lambda(1-\lambda)A}{D_F^2}Den_C-\lambda(1-\lambda)A\gamma(1-\gamma)A[\frac{(1-g)}{D_C^2}-\frac{g}{D_F^2}]g\gamma\mu_F$$

and the latter is:

$$\frac{g\gamma(1-\gamma)\lambda(1-\lambda)A}{D_C^2}Den_F-\lambda(1-\lambda)A\gamma(1-\gamma)A[\frac{g}{D_C^2}-\frac{(1-g)}{D_F^2}]g(1-\gamma)\mu_C$$

Substituting for Den_F and Den_C and rearranging one can see that the numerators only differ in the sign: p_C decreases as σ increases, whereas p_F increases. However, since $Den_F<Den_C$, then the positive effect from p_F dominates. □

Proof of Proposition 3

Proof. In terms of first period behaviour, nothing changes with respect to the baseline model: it can therefore be proved that all incumbents follow their signal at $t = 1$ by using the same arguments provided for the baseline model (refer to the proof of Theorem 1). Let's now move to the decision to be taken at $t = 2$. An incumbent receiving a signal that suggests him to flip-flop knows that if he does flip-flop, with probability g voters will observe a media report indicating he has been consistent, and viceversa if he goes against his signal and avoids the flip-flop there is a probability $1 - g$ that his track record will come across as flip-flopping. It immediate to see, therefore, that given any two reputation levels, the incentive to strategically avoid the flip-flop is smaller when the reporting media is noisy. However, the voter is aware of the noisy signal received from the media and therefore she also adjusts her beliefs: whenever observing a flip-flopping track-record, for example, the voter knows that with some probability, the politician actually played consistently and viceversa. In particular, let's fix the strategy played by incumbents, and in particular let's assume that incompetent incumbents follow a flip-flopping signal with probability σ. If the voter were able to observe track-records perfectly, beliefs would be $\mu_C(\sigma)$ and $\mu_F(\sigma)$. Given the noise, the belief the incumbent is competent when observing a consistent track-record is $\tilde{\mu}_C = p_C\mu_C(\sigma) + (1 - p_C)\mu_F(\sigma)$, where $p_C = \frac{gPr(\tau=C)}{gPr(\tau=C)+(1-g)Pr(\tau=F)}$. Analogous expressions, denoted by $\tilde{\mu}_F$ and p_F, hold for the case in which the voter observes a flip-flopping track-record and are derived by simply swapping C with F. Denote now for simplicity by $\rho \equiv \rho_2(s_2 \neq s_1, q)$. It follows that the incumbent prefers to follow a flip-flopping signal whenever the following inequality holds:

$$\rho + gr(\tilde{\mu}_C)2\phi + (1 - g)r(\tilde{\mu}_F)2\phi \geq 1 - \rho + gr(\tilde{\mu}_C)2\phi + (1 - g)r(\tilde{\mu}_F)2\phi$$

Using the definitions of $\tilde{\mu}_C$ and $\tilde{\mu}_F$, this expression can be rearranged to yield the following:

$$\frac{2\rho - 1}{\phi} \geq (2g - 1)(\tilde{\mu}_C - \tilde{\mu}_F)$$

and further to get:

$$\frac{2\rho - 1}{\phi} \geq (2g - 1)(p_C + p_F - 1)(\mu_C - \mu_F)$$

Notice that if $g = 1$, i.e. no noise, then one gets back to the expression from the baseline model, given also that $p_C(g = 1) = 1 = p_F(g = 1)$. Let's now consider the case in which a truthful equilibrium is not sustainable when $g = 1$, since $\frac{2\rho-1}{\phi} < (\mu_C^T - \mu_F^T)$. When noise is introduces, the right hand side becomes $(2g - 1)(p_C + p_F - 1)(\mu_C^T - \mu_F^T)$. It is immediate to check that given μ_C^T and μ_F^T, the right-hand side strictly decreases as g

becomes smaller than one. This means that there is a level of noise $g^* \in \left(\frac{1}{2}, 1\right)$ such that

$$\frac{2\rho - 1}{\phi} = (2g^* - 1)(p_C(g^*, \mu^T) + p_F(g^*, \mu^T) - 1)(\mu_C^T - \mu_F^T).$$

Moreover, remember that $\mu_C(\sigma) - \mu_F(\sigma)$ is strictly increasing in σ and by Lemma 3 it follows that $p_C(\hat{g}, \mu^T) + p_F(\hat{g}, \mu^T) - 1$ is also strictly increasing in σ. As a result, then, decreasing g always increases the equilibrium level of σ. This means that as long as $\sigma < 1$, decreasing g will alleviate the accountability distortion caused by incumbents avoiding flip-flops. Since the equilibrium level of $\tilde{\mu}_C - \tilde{\mu}_F$ increases as g decreases (as long as $\sigma < 1$), because in equilibrium $\frac{2\rho-1}{\phi} \geq (2g - 1)(\tilde{\mu}_C - \tilde{\mu}_F)$, then the selection of politicians also improves thanks to Lemma 1. As a result, decreasing g improves welfare as long as it crowds out the lies of politicians. Decreasing g below g^*, however, $\mu_C - \mu_F$ cannot increase further and hence incompetent incumbents start having a strict preference towards following their signal when it prescribes a flip-flop. Therefore, decreasing g further will hurt learning and have no benefit on accountability. It follows that g^* is the optimal level of news informativeness. □

Proof of Proposition 4

Proof. I will first show that for each σ, the reputation spread without delegation, denoted by $\lambda_C - \lambda_F$, is larger than the reputation spread with delegation, denoted by $\lambda_C^D - \lambda_F^D$, that is to say

$$\lambda_C - \lambda_F > \lambda_C^D - \lambda_F^D$$

In order to show this, write the reputation spread $\lambda_C - \lambda_F$ as follows:

$$\frac{\lambda\gamma}{\lambda\gamma + (1-\lambda)(1-A\sigma)} - \frac{\lambda(1-\gamma)}{\lambda(1-\gamma) + (1-\lambda)A\sigma},$$

where $A = (1-\gamma)(q^2 + (1-q)^2) + 2\gamma q(1-q)$. whereas $\lambda_C^D - \lambda_F^D$ can be written as:

$$\frac{\lambda\pi}{\lambda\pi + (1-\lambda)(1-A_d\sigma)} - \frac{\lambda(1-\pi)}{\lambda(1-\pi) + (1-\lambda)A_D\sigma},$$

where $\pi = \gamma q + (1-\gamma)(1-q)$ and $A_D = (1-\pi)q + \pi(1-q)$. Let's start from comparing λ_F and λ_F^D. The latter is larger whenever $\frac{A_D}{1-\pi}\sigma < \frac{A}{1-\gamma}\sigma$. First of all, substituting immediately shows that $A_D = A$. This is due to the fact that the number of flip-flopping sequences for the incompetent politician remains the same under delegation, given that the action is delegated to another incompetent politician. Since $\pi < \gamma$, then it immediately follows that $\lambda_F^D > \lambda_C$. [22] Therefore, $\lambda_F^D > \lambda_C$ for each value of σ. In order to see that $\lambda_C > \lambda_C^D$, notice that $\frac{1}{\pi}(1 - A_D\sigma) > \frac{1}{\gamma}(1 - A\sigma)$. Therefore, we have that fixing a level of σ, the reputational spread between consistency and flip-flopping is always larger when the first action is not delegated. Since $\lambda_C - \lambda_F$ and $\lambda_C^D - \lambda_F^D$ are monotonically increasing in σ, this also means that, denoting by σ^D the σ relative to the delegation game,

$$\lambda_C - \lambda_F = \lambda_C^D - \lambda_F^D \Rightarrow \sigma_D > \sigma.$$

In other words, the same reputational spread is achieved with more distortion in the game without delegation. Let's now denote by $\lambda_C^{D,T}$ and $\lambda_F^{D,T}$ the reputations obtained in the delegation game under $\sigma_D = 1$ (i.e. truthful play): if $\phi > \frac{2\rho-1}{\lambda_C^{D,T} - \lambda_F^{D,T}} \equiv \phi_D$, then a truthful equilibrium is not sustainable in the delegation game. Since $\lambda_C^T - \lambda_F^T > \lambda_C^{D,T} - \lambda_F^{D,T}$, a truthful equilibrium is also not sustainable in the game without delegation. It follows that in both games, $\Delta^* = \Delta_D^* = \frac{2\rho-1}{\phi}$. Therefore, $\sigma^* < \sigma_D^*$. In other words, the amount of learning that can be sustained in equilibrium is the same in both games. Therefore, the expected competence of the elected politician is the same in both games. However, delegating the first action decreases the amount of distortion, since $\sigma_D^* > \sigma$: the benefit of delegation can be written as $(\sigma_D^* - \sigma^*)(2\rho - 1)(2q - 1)(1 - \gamma)$. The cost of delegation, on the other hand, is due to the fact that the first action is taken with probability 1 by an incompetent agent, that is $\alpha\lambda(1 - q)$. However, if the welfare weight of the first action is small enough, i.e.

$$\alpha < \underline{\alpha} \equiv \frac{(\sigma_D^* - \sigma^*)(2\rho - 1)(2q - 1)(1 - \gamma)}{\lambda(1 - q)}$$

, then delegating the first action to an incompetent agent is optimal. Notice that $\underline{\alpha}$ can be larger than 1, meaning that it can be the case that even if the first action is relatively more important than the second, delegation is optimal. □

Proof of Proposition 5

Proof. Let's start the analysis by writing the modified reputations. Given that there are now two signals (the track record and the media signal) we now have four different reputations. I denote by $\mu_{C,E}$ and $\mu_{F,O}$ the reputation from consistent play given that the media endorses (E) or opposes (O) the politician's decision. Given this notation, the reputation expressions can be written in the following way:

[22] If the action was delegated to a more competent agent, the number of flip-flops of the incompetent incumbent would decrease, improving even further the reputation associated to a flip-flop.

$$\mu_{C,E} = \frac{\lambda\gamma q_M}{\lambda\gamma q_M + (1-\lambda)[(p_2 q q_M + (1-p_2)(1-q)(1-q_M)) + ((1-p_2)q(1-q_M) + p_2(1-q)q_M)(1-\sigma)]}$$
$$\mu_{F,E} = \frac{\lambda(1-\gamma)q_M}{\lambda(1-\gamma)q_M + (1-\lambda)((1-p_2)q(1-q_M) + p_2(1-q)q_M)\sigma}$$
$$\mu_{C,O} = \frac{\lambda\gamma(1-q_M)}{\lambda\gamma(1-q_M) + (1-\lambda)[((1-p_2)q(1-q_M) + p_2(1-q)q_M) + ((1-p_2)q q_M + p_2(1-q)(1-q_M))(1-\sigma)]}$$
$$\mu_{F,O} = \frac{\lambda(1-\gamma)(1-q_M)}{\lambda(1-\gamma)(1-q_M) + (1-\lambda)((1-p_2)q(1-q_M) + p_2(1-q)q_M)\sigma}$$

Notice that if $q_M = \frac{1}{2}$, then we get back to the expressions used in the baseline model. Moreover, notice that compared to the reputations from the baseline model, $\mu_{C,E} > \mu_C > \mu_{C,O}$ and the analogous inequality holds for the flip-flopping reputations. Now, denote by $S = Pr(s_M = F | s_P = F) = \rho_2 q_M + (1-\rho_2)(1-q_M)$, i.e. the probability that, given the incumbent's signal is flip-flopping, the media signal also endorses a flip-flop. It turns out that when they receive a flip-flopping signal, incompetent incumbents follow their signal if the following inequality holds:

$$\frac{2\rho_2 - 1}{\phi + q} = [S\mu_{C,O} + (1-S)\mu_{C,E}] - [S\mu_{F,E} + (1-S)\mu_{F,O}]$$

The expression can be rearranged to yield:

$$\frac{2\rho_2 - 1}{\phi + q} = S[\mu_{C,O} - \mu_{F,E}] + (1-S)[\mu_{C,E} - \mu_{F,O}]$$

Using implicit differentiation one can see that σ^* can decrease when q_M increases. In such a situation, accountability welfare decreases. In fact, it is sufficient to look at is the derivative of the right-hand side of the expression above with respect to q_M. In particular, if the right hand side increases in q_M, then σ^* needs to decrease for equilibrium to be restored. Differentiating the right-hand side we get the following condition for an increase in q_M to decrease σ^*:

$$(2\rho_2 - 1)[(\mu_{C,O} - \mu_{F,E}) - (\mu_{C,E} - \mu_{F,O})] + (1-S)\left[\frac{\partial \mu_{C,E}}{\partial q_M} - \frac{\partial \mu_{F,O}}{\partial q_M}\right] - S\left[\frac{\partial \mu_{F,E}}{\partial q_M} - \frac{\partial \mu_{C,O}}{\partial q_M}\right] > 0$$

It turns out that the above inequality holds when γ and q are both high (with γ potentially higher than q), λ is sufficiently high and q_M is not too high. In this situation, increasing the informativeness of the commentator signal has the effect of increasing the distortion to accountability.

□

Proof of Proposition 5

Proof. The proof just consists of simple algebra. Denote by $p_3 = \gamma\rho_2 + (1-\gamma)(1-\rho_2)$ and $\bar{p}_3 = \gamma\bar{\rho}_2 + (1-\gamma)(1-\bar{\rho}_2)$. We have:

$$
\begin{aligned}
&Pr(a_3 = a_2 | a_2 = a_1) = \\
&= \mu_C\gamma + (1-\mu_C)[Pr(s_2 = s_1 | a_2 = a_1, L)(q\bar{p}_3 + (1-q)(1-\bar{p}_3)) + \\
&(+1 - Pr(s_2 = s_1 | a_2 = a_1, L))((1-q)p_3 + q(1-p_3))]
\end{aligned}
$$

and

$$
Pr(a_3 = a_2 | a_2 = a_1) = \mu_F\gamma + (1-\mu_F)(qp_3 + (1-q)(1-p_3))
$$

Moreover, $Pr(s_2 = s_1 | a_2 = a_1, L) = \frac{1-A}{1-A\sigma^*}$. All I want to show is that $Pr(a_3 = a_2 | a_2 = a_1) \geq Pr(a_3 = a_2 | a_2 = a_1)$. Using the expressions for the probabilities and rearranging yields:

$$
\frac{1-A}{1-A\sigma^*} \geq \frac{2\bar{p}_3 - 1}{\bar{p}_3 + p_3 - 1}
$$

Since only the left-hand side depends on σ^*, and it is strictly increasing in σ^*, then evaluating the inequality at the lower bound on σ^*, i.e. $\underline{\sigma} = \frac{1-\gamma}{A}$, works as a sufficient condition. Substituting in for $\bar{p}_3$ and p_3 and $\underline{\sigma}$ yields:

$$
\frac{2(1-\gamma) - 1 + 2(2\gamma-1)\rho_2}{2(1-\gamma) - 1 + (2\gamma-1)(\rho_2 + \bar{\rho}_2)} \geq \frac{1-\gamma}{\gamma} 2q(1-q) + (q^2 + (1-q)^2)
$$

and it can be checked numerically that this condition is always satisfied for the parameter values of interest for the model. □