Flip-flopping and Electoral Concerns*

Giovanni Andreottola ⋆

July 2020

Abstract

Policy-making in changing environments requires policy revisions. At the same time, however, the action record of a politician contains information about his competence that can create a reputation premium for consistent decision-making. As a result, politicians with high reputation concerns are deterred from updating their policies in light of new information, decreasing both the quality of the decisions taken and the ability of voters to select competent politicians. My model thus provides a rationale for the stigma often associated with policy shifts (flip-flopping) and increases our understanding of the factors shaping the degree to which inconsistent policy-making is punished by voters. In addition to this, the model has implications concerning several important issues, such as the behavior of term-limited politicians, the timing of policy shifts and the types of tasks to assign to a reputation concerned agent.

Keywords: flip-flopping; elections; political agency; accountability; reputation; persistence; policy change

*I would like to thank the Editor Tom S. Clark and three anonymous referees for excellent comments and suggestions. I am grateful to my advisors Andrea Mattozzi and David Levine for their great supervision during my PhD at the EUI, and to Larry Samuelson for hosting me as a postdoc at the Cowles Foundation in Yale between 2018 and 2019. The paper also benefitted from useful discussions with Piero Gottardi, Andrea Galeotti, Ronnie Razin, Alessandro Riboni and several others. I would also like to thank audiences at the Leitner Political Economy Seminar in Yale, the EUI Microeconomics Working Group, the EUI Workshop on Media Effects, the University of Bologna, the University of Naples, the Max Planck Institute in Munich, the NICEP Workshop in Nottingham, the IOEA in Cargese and the SIOE conference in Montréal for their questions and comments.

⋆CSEF and Federico II University of Naples.
Email: Giovanni.Andreottola@eui.eu Website: `www.giovanniandreottola.com`

Policy-making in a changing environment requires revising and updating decisions over time, in response to the arrival of new information. With respect to many key policy areas, such as foreign policy or financial regulation, voters observe the stream of decisions taken over time by the incumbent, but they can neither directly observe the information which politicians receive (for example because of confidentiality, or because of a high degree of technical complexity), nor they are able (at least within the relevant timeframe for elections or media commentary) to judge whether the decisions taken were right or wrong.

In such environments, does the record of actions chosen by a representative contain information about his competence? If so, what incentives does this provide to a political decision maker? Starting from these two key questions, I build a model of political agency in which an incumbent politician needs to take two decisions in sequence, having some private information at hand. My main result is to show that if the state of the world is persistent – that is if the right approach to the policy issues addressed in the two periods of office is likely to be the same – the incumbent politician has an incentive to disregard new information in order to appear competent. When reputation concerns are sufficiently strong, these incentives inhibit policy responsiveness, resulting in an inefficient degree of policy persistence. On top of that, when policy choices only imperfectly reflect the incumbent's information the ability of voters to learn about the incumbent's competence from his track-record is diminished, creating a further cost for society.

By uncovering the incentives deriving from the repeated nature of political decision-making, a contribution of the model is to rationalize the widespread stigma politicians face for policy reversals, which are also disparagingly known as flip-flops.

The effect of electoral concerns on policy reversals has several implications. For example, politicians for whom reputation concerns are less relevant, be it because their re-election prospects are safe, or because they face a binding term-limit, should be more willing to reverse their previous decisions than politicians facing more competitive elections. What is more, my model implies that this turning back from previous decisions or electoral promises need not be the sign of a lack of accountability, but to the contrary it might be a beneficial consequence of not being subject to election concerns[1]. The fact that politicians facing strong reputation concerns are less likely to implement policy reversals also bears the consequence that it might be necessary to look back at behavior taking place at *unsuspecting times* in order to draw information from a politician's consistency. This also means that a politician might be negatively evaluated for flip-flops that occurred far away in the past.

In addition to the conventional wisdom and anecdotal evidence of the negative evaluation of policy changes, there is also fairly abundant empirical and experimental evidence that voters value consistency as a trait in their political representatives. In many instances, inconsistency is negatively evaluated independently of the direction of the policy change and the partisanship of the flip-flopping politician[2]. Among the several reasons reported by survey respondents to motivate the negative evaluation of inconsistent decision-makers, a recurring one is that inconsistent decision-makers are more likely to be incompetent[3]. Along with competence being directly mentioned as a reason for the

[1]List and Sturm (2006) show that term-limited governors are more likely to reverse previous policies that they have previously enacted.

[2]See for example Allgeier, Byrne, Brooks, and Revnes (1979), Carlson and Dolan (1985), Tomz and Van Houweling (2012), Doherty, Dowling, and Miller (2015), Robison (2017). Please refer to the literature review section for additional details on empirical findings on flip-flopping.

[3]The common charges of indecisiveness or weakness can also be related to competence. Incompetence as a reason for the negative evaluation of flip-floppers is reported by Tomz and Van Houweling (2012), Levendusky and Horowitz (2012b) among others.

stigma associated with inconsistency, there is also indirect evidence of the link between consistency and competence. For example, the time between a decision and its reversal decreases the negative evaluation of a flip-flop[4]; similarly, policy reversals on complex issues, where new information is more likely to sway the optimal decision, are evaluated less badly than flip-flops on simpler issues[5]. It has also been shown that the presence of supporting information on necessity of a policy change improves the evaluation of an inconsistent policy-maker[6].

In this respect, the contribution of my work is to provide a simple and flexible formalization that captures many of the "stylized facts" highlighted by empirical studies. In particular, my model emphasizes that a fundamental role to understand the consequences of policy changes is played by the persistence of the state of the world, that is the likelihood of the correct policy approach to remain the same across periods. Persistence can be interpreted as being a function of several key aspects of a policy-making environment, such as i) time between decisions (the shorter the time between decisions, the higher persistence); ii) the type, complexity and similarity of the issues decided upon; iii) the degree of public information on whether a change in the environment occurred. With respect to these elements of the decision-making environment, an implication of my model is that because of electoral incentives, the quality of decision-making and the informativeness of a politician's track-record are non-monotonic in the level of persistence. On top of the positive implications discussed above, this can also have normative, institutional design

[4]A famous quote by Congressman Jack Flynt goes: *"I can vote conservative or I can vote liberal, and my constituents don't care. But if vote one way at two o'clock and another way at two-thirty, they'll think they sent an idiot to Congress"*. See also Doherty, Dowling, and Miller (2015) for an analysis of how time affects the perception of a flip-flop.

[5]See again Doherty, Dowling, and Miller (2015) and Tavits (2007) for a related point on principle versus pragmatic issues.

[6]See Levendusky and Horowitz (2012b) and Robison (2017).

implications for the timing of decision-making, the types of tasks to assign to reputation concerned decision-makers and the degree of transparency over the environment in which a decision-maker acts.

Finally, an interesting consideration concerns the role of consistency in the evaluation of politicians: what my model predicts is that a politician's action record provides voters with indirect information on the competence of a decision-maker. This suggests that consistency - and thus flip-flopping - are likely to play a larger role in circumstances in which there is little direct information on the quality of the decisions taken by a politician: this can happen for reasons of confidentiality, lack of transparency, but also because of credibility issues. The action record of a politician, as a matter of fact, is generally a rather objective or easily fact-checkable source of information on the political representative.

Literature Review

Literature on Consistency and Flip-flopping

There is a long-standing literature investigating the effects of consistency in shaping an agent's evaluation. Allgeier, Byrne, Brooks, and Revnes (1979) is the seminal study introducing the so called *waffle pheonomenon*: attitude changes are perceived negatively, even if the change increases the similarity with the evaluator. Carlson and Dolan (1985) reach analogous results in a similar experiment specifically involving political candidates. Despite not finding direct evidence of the waffle phenomenon, Hoffman and Carver (1984) find that agreement with the decision-maker following an inconsistent track-record is not rewarded by voters[7]. Recent research has gotten into increasing detail concerning the conditions leading flip-flopping politicians to be punished by voters. Doherty, Dowling, and Miller (2015) consider the following factors: the degree of agreement on the current position; the amount of time passed and the type of issue. Their results are in line with the predictions of my model, in particular concerning the effect of persistence: consistency matters and is only attenuated by the effect of agreement; the time passed between a decision and its reversal makes voters more forgiving; finally, flip-flops on complex issues are evaluated less badly than flip-flops on simpler ones[8]. Robison (2017) emphasizes that consistency is more relevant in the absence of direct information on the quality of the action taken by the politician, which mirrors the point my paper makes about consistency as indirect information about the politician's competence. Another important reference is

[7]Similarly, Croco (2016) finds that alignment with the decision-maker tends to prevail over the judgment of consistency, but she finds an effect of consistency among voters who are neutral on the issue at stake.

[8]Relatedly, Tavits (2007) uses election data to show that flip-flops on pragmatic issues are seen less badly than flip-flops on issues of principle.

Tomz and Van Houweling (2012), who find that policy changes are costly not only because of the ideological ambiguity they cause (a channel I do not consider in my model), but also because of the negative inference voters draw about a candidate's character, of which competence constitutes an important component[9]. Sigelman and Sigelman (1986) and Sorek, Haglin, and Geva (2018) also highlight the close link between consistency and the perceived competence of a politician, the latter showing that the more a leader is seen as competent, the better is his reputation following a flip-flop, which reflects the results of my model.

With respect to this branch of the literature, my model aims to provide a simple formalization capable of matching several of the stylized facts emerging from empirical studies.

There are also some theoretical papers broadly related to the idea of flip-flopping: Hummel (2010) builds a spatial model in which candidates suffer a negative valence shock when changing platform between the primary and the general election. Whereas his model uses an exogenously given cost of policy change to derive predictions on the platforms chosen by candidates in a two-stage election, my model explores the reasons why policy changes cause a reputation loss, with implications for policy persistence and its responsiveness to new information. In Bernhardt and Ingerman (1985), the penalization for policy change derives from the perceived ideological riskiness of the platform proposed and is used to explain policy divergence with respect to the median voter's preferences. Similarly, in Agranov (2016) candidates undertake costly electoral posturing to appear liberal (in the primary election) or moderate (in the general). These models have therefore

[9]Related studies by the same authors, with similar findings, include Tomz and Van Houweling (2012) and Tomz and Van Houweling (2014).

a very different focus and predictions unrelated to those of my work.

Experts and Reputation Concerns

The electoral concerns model that I consider builds on the seminal contributions of Canes-Wrone, Herron, and Shotts (2001), Maskin and Tirole (2004) and Prat (2005). In this respect, my contribution is to study a two-period political-agency framework and its implications, such as the key one concerning flip-flopping[10]. My work is also related to models of repeated expert behavior, including Prendergast and Stole (1996), Majumdar and Mukand (2004), Li (2007) and Aghion and Jackson (2016). Along with several differences in the model structure, my contribution revolves around the focus on a changing but persistent environment.

Audience Costs

Despite being more general and not specifically targeted to international relations, my model is also related to the idea of *audience costs*, i.e. that leaders who turn back on previously made threats suffer a reputation loss (the seminal contribution is Fearon (1994)). Tomz (2007) uses survey experiments to provide direct evidence of the existence of such costs, also showing that audience costs are deeply intertwined with considerations about reputation and competence. Schultz (2001) also supports the idea that leaders who back down are perceived as less competent. Along similar lines, a key result in Levendusky and Horowitz (2012a) is to show that a leader changing his or her mind is seen as less competent than one who stays coherent, with partisanship playing hardly any role[11].

[10] Other political-agency models I draw on are Ashworth and Shotts (2010) and Fu and Li (2014).

[11] A theoretical model on conflict sharing some insights with my analysis is Smith (1998). In his model, however, competence is related to the probability of a victorious outcome in a conflict and not to the

The Model

There are two periods $t \in \{1, 2\}$; an incumbent politician takes a binary action $a_t \in \{0, 1\}$ in each period on behalf of a representative voter. The payoff to the representative voter is 1 if the action a_t matches the state of the world $\omega_t \in \{0, 1\}$, and zero otherwise. The commonly known prior probability that ω_1 takes value 0 or 1 is equal to $\frac{1}{2}$. This simplifies the analysis but it is by no means crucial for the results. The state of the world is persistent, so that $Pr(\omega_2 = \omega_1|\omega_1) = \gamma > \frac{1}{2}$. The persistence parameter γ, which as we will see is of great importance in the model, can be thought of as representing how informative past signals (and thus past policy choices) are for future policy choices. In this respect, it can also be thought of as measuring the rate of depreciation of the information guiding policies (γ close to 1/2 represents a high rate of depreciation).

Incumbents can be of two types $\theta \in \{H, L\}$, i.e. competent and incompetent. The incumbent's type is fixed across periods. The probability that an incumbent is competent is $\lambda \in (0, 1)$, which is common knowledge. Incumbents do not observe their type[12]. In each period, both types receive a realization $s_t \in \{0, 1\}$ of a binary signal about the state of the world: the accuracy of the signal, $Pr(s_t = \omega_t|\omega_t) = q_\theta$, depends on the politician's type: $\frac{1}{2} \leq q_L < q_H \leq 1$. The signal is private information of the incumbent and s_1 and s_2 are independent conditional on the realization of the state of the world ω_t.

The incumbent's private signals are the only information available about the state of the world[13]. Therefore, for all $\gamma \leq 1$ and $q_H > q_L \geq 1/2$, the signal received by

accuracy of the information on whether to escalate the conflict. As a result, reputation concerns make the behavior of the politician more hawkish rather than more persistent as in my model; on top of that, in his model threats are never reversed in equilibrium.

[12]This assumption simplifies the analysis but the results remain qualitatively unchanged under the alternative assumption that incumbents know their type: in fact, previous drafts of the paper were based on the assumption of privately observed competence.

[13]In the Online Appendix I consider the scenario in which, in addition to the incumbent's private

any politician is decision-relevant, meaning that $a_t = s_t$ maximizes the probability of matching the action to the state.

The reputation of the incumbent at the end of $t = 2$ consists of the belief that the representative voter holds of him being competent. I denote this by $r(a_1, a_2) = Pr(\theta = H|a_1, a_2)$. In addition to utility from the policy actions chosen by the incumbent, the representative voter benefits from the ability to select competent politicians, for example through elections: for maximum generality, I model this in the form of a strictly increasing and strictly convex function $\nu(\cdot)$ of the incumbent's reputation[14]. To sum up, the representative voter's utility is:

$$U_v = \sum_{t=1}^{2} \mathbb{1}_{a_t=\omega_t} + \delta\nu(r(a_1, a_2)) \tag{1}$$

where δ measures the welfare weight of selection relative to that of the quality of the policies chosen by the incumbent.

The utility of the incumbent depends both on the probability of matching the state of the world and on the reputation earned at the end of the two periods[15]. The incumbent's objective function is hence as follows:

$$U_p = \sum_{t=1}^{2} \mathbb{1}_{a_t=\omega_t} + \phi r(a_1, a_2) \tag{2}$$

signals, the state of the world ω_1 is revealed to the public after the first action a_1 has been taken. This can be thought of as feedback on the policy implemented in the first period, or as the result of a journalistic investigation.

[14]If the incumbent is subject to re-election in the presence of a uniformly distributed valence shock, $\nu(\cdot)$ would take the quadratic form. For more details, please refer to the Online Appendix.

[15]Just like the utility of the representative voter, this can be shown to be the reduced form of a game ending with an election with a uniformly distributed valence shock. In this respect, the use of a linear function is not without loss of generality, but has been chosen for tractability and follows the vast majority of the career concerns literature.

where ϕ measures the weight of the reputation concern: $\phi = 0$ represents the case of a purely policy-motivated incumbent, whereas $\phi \to +\infty$ represents the case of a fully office-motivated incumbent. Notice that since the ex-ante expected value of reputation is equal to the prior belief λ following Bayes' Rule, social welfare coincides with the welfare of the representative voter.

The incumbent's strategy is, in each time period $t \in \{1, 2\}$, a mapping between the private signals received until that period (s_1 at $t = 1$ and (s_1, s_2) at $t = 2$), the actions taken until that period and a probability distribution across the available actions $a_t \in \{0, 1\}$. This means that the signal received in period $t = 1$ and the action chosen then matter for the mapping between s_2 and a_2.

The equilibrium concept I use is Perfect Bayesian Equilibrium (PBE). Notice that at least for sufficiently high values of reputation concerns ϕ, in the set of PBEs there are always uninformative equilibria that are not interesting for the sake of my analysis[16]. In order to rule out these equilibria, I focus on equilibria that are robust to a fraction $\epsilon \approx 0$ of the incumbent having $\phi = 0$ and thus always playing to maximize the probability of matching the action to the state of the world. When I talk about equilibrium of the game, therefore, I consider the PBE in the limit $\epsilon \to 0$.

[16] An example would be an equilibrium where reputation is zero unless the actions played are $(a_1, a_2) = (0, 0)$ and the incumbent always plays $(0, 0)$. For sufficiently high ϕ, this would be an equilibrium, though based on unreasonable off-equilibrium beliefs. Notice that other kinds of perverse equilibria are ruled out since the incumbent does derive some direct utility from the action.

Results

I start the analysis of the model with the information available to the incumbent. After receiving the first signal realization s_1, the incumbent assigns probability $Pr(\omega_1 = s_1|s_1) = \rho_1$ to the state of the world being equal to the signal[17]. It can immediately be seen that since $q_H > q_L \geq \frac{1}{2}$, it holds that $\rho_1 > \frac{1}{2}$. Given the symmetric prior, ρ_1 is the same for both $s_1 = 0$ and $s_1 = 1$; analogously, the posterior probability that the state of the world is equal to the second realization s_2 only depends on whether $s_2 = s_1$ or $s_2 \neq s_1$ and not on the realizations of s_1 and s_2. In other words, what matters is whether the signal realizations observed by the incumbent are consistent (for which I use the subscript c), or whether they flipped (subscript f). I denote $Pr(\omega_2 = s_2|s_1, s_2 = s_1)$ as $\rho_{2,c}$ and $Pr(\omega_2 = s_2|s_1, s_2 \neq s_1)$ as $\rho_{2,f}$ [18].

Notice that both $\rho_{2,c} \geq \rho_{2,f} \geq 1/2$ (with both inequalities holding strictly for $\gamma \in$

[17]The posterior ρ_1 can be written as follows:

$$\rho_1 = \frac{\frac{1}{2}[\lambda q_H + (1-\lambda)q_L]}{\frac{1}{2}[\lambda q_H + (1-\lambda)q_L] + \frac{1}{2}[\lambda(1-q_H) + (1-\lambda)(1-q_L)]} = \lambda q_H + (1-\lambda)q_L, \tag{3}$$

where $\lambda q_H + (1-\lambda)q_L = Pr(s_1|\omega_1 = s_1)$ and analogously $\lambda(1-q_H) + (1-\lambda)(1-q_L) = Pr(s_1|\omega_1 \neq s_1)$.

[18]The following holds:

$$\rho_{2,c} = \frac{\frac{1}{2}A_c}{\frac{1}{2}A_c + \frac{1}{2}B_c} \tag{4}$$

where $A_c \equiv Pr(s_2 = s_1|\omega_2 = s_2, s_1)$ and $B_c \equiv Pr(s_2 = s_1|\omega_2 \neq s_2, s_1))$ which can in turn be written out as follows:

$$A_c = \lambda[\gamma q_H^2 + (1-\gamma)q_H(1-q_H)] + (1-\lambda)[\gamma q_L^2 + (1-\gamma)q_L(1-q_L)]$$
$$B_c = \lambda[\gamma(1-q_H)^2 + (1-\gamma)q_H(1-q_H)] + (1-\lambda)[\gamma(1-q_L)^2 + (1-\gamma)q_L(1-q_L)]$$

The derivation of $\rho_{2,f}$ is analogous:

$$\rho_{2,f} = \frac{\frac{1}{2}A_f}{\frac{1}{2}A_f + \frac{1}{2}B_f} \tag{5}$$

where again I use $A_f \equiv Pr(s_2 \neq s_1|\omega_2 = s_2, s_1)$ and $B_f \equiv Pr(s_2 \neq s_1|\omega_2 = s_2, s_1)$ which can be written out in a way analogous to that for the previous posterior probability:

$$A_f = \lambda[(1-\gamma)q_H^2 + \gamma q_H(1-q_H)] + (1-\lambda)[(1-\gamma)q_L^2 + \gamma q_L(1-q_L)]$$
$$B_f = \lambda[(1-\gamma)(1-q_H)^2 + \gamma q_H(1-q_H)] + (1-\lambda)[(1-\gamma)(1-q_L)^2 + \gamma q_L(1-q_L)]$$

$(0.5, 1)$), meaning that the most likely state of the world at time t is always the one coinciding with the most recent signal s_t. Therefore, the probability of matching the action to the state of the world is maximized by choosing $a_t = s_t$. The private signals received by the incumbent also represent the information from which the voter can learn something about the incumbent's competence: this means that $a_t = s_t$ also maximizes the amount of learning by the representative voter, that is the expected value of the selection component of the representative voter's utility as per condition (1). Given that the representative voter's utility also coincides with social welfare, the candidate equilibrium in which the optimal incumbent's strategy for all s_t in both periods $t \in \{1, 2\}$ is $a_t = s_t$ maximizes social welfare. I denote this equilibrium as the truthful equilibrium.

If the incumbent's incentives were solely policy-related, $a_t = s_t$ would be the dominant strategy and the truthful equilibrium would always be the outcome of the game. However, reputation concerns can distort the incumbent's behavior away from the truthful equilibrium benchmark. In order to understand how this can happen, consider the four possible action records that can occur in the game, $(a_1, a_2) \in \{(0, 0), (0, 1), (1, 0), (1, 1)\}$. I denote $(0, 0)$ and $(1, 1)$ as consistent action records (or track records); $(1, 0)$ and $(0, 1)$ are instead flip-flopping action records, since the incumbent changes his policy stance from period 1 to period 2. For each realization of (a_1, a_2), the incumbent's reputation is calculated as follows:

$$r(a_1, a_2) = Pr(\theta = H|(a_1, a_2)) = \frac{\lambda Pr((a_1, a_2)|\theta = H)}{\lambda Pr((a_1, a_2)|\theta = H) + (1 - \lambda)Pr((a_1, a_2)|\theta = L)} \quad (6)$$

As a result, an action record delivers a reputation higher than the prior as long as the probability to end up with that action record conditional on being a high type is higher

than the probability conditional on being a low type, or formally $\frac{Pr((a_1,a_2)|\theta=H)}{Pr((a_1,a_2)|\theta=L)} \geq 1$. The opposite holds for a reputation lower than the prior. Notice that given the symmetric initial prior, for each type θ the probability of obtaining each of the two consistent and flip-flopping signal sequences is the same: therefore, we can simply focus on the probability of having a consistent action record, resulting in a reputation that I denote as r_c, versus a flip-flopping one, which yields reputation r_f.

Since the incumbent is free to choose a policy action independently of the signal he received, $Pr((a_1, a_2)|\theta)$, and therefore r_c and r_f, crucially depend on the equilibrium strategy chosen by the incumbent. To fix ideas I first consider the benchmark case of a truthful equilibrium, in which $(a_1, a_2) = (s_1, s_2)$, giving rise to reputations that I denote as r_c^T and r_f^T. Given type θ and the resulting signal accuracy q_θ, the probability of having each consistent track-record is the following:

$$Pr((0,0)|\theta) = Pr((1,1)|\theta) = \frac{\gamma}{2}q_\theta^2 + \frac{\gamma}{2}(1-q_\theta)^2 + (1-\gamma)q_\theta(1-q_\theta) \tag{7}$$

and it can easily be checked that the value of (7) is increasing in q_θ. Therefore, $\frac{Pr((a_1,a_2)|a_1=a_2,\theta=H)}{Pr((a_1,a_2)|a_1=a_2,\theta=L)} > 1$, which means that in the truthful equilibrium, the reputation r_c^T after a consistent action record is better than the prior. Conversely, the probability of a flip-flopping action-record is decreasing in q_θ, resulting in a reputation r_f^T worse than the prior, that is $r_f^T < \lambda < r_c^T$[19].

The intuition for the bad reputation associated with flip-flopping is that given that the state of the world is persistent ($\gamma > \frac{1}{2}$), the more precise is the incumbent's signal,

[19]The probability of each flip-flopping action-record is:

$$Pr((0,1)|\theta) = Pr((1,0)|\theta) = \frac{1-\gamma}{2}q_\theta^2 + \frac{1-\gamma}{2}(1-q_\theta)^2 + \gamma q_\theta(1-q_\theta) \tag{8}$$

which is decreasing in q_θ.

the less it flips. At the extremes, a perfect signal flips with probability equal to the state of the world flipping, i.e. $1-\gamma$, whereas a completely noisy signal flips with probability $\frac{1}{2}$, and $\frac{1}{2} > 1-\gamma$. This also means that r_f^T decreases as γ increases: when the state is completely persistent, i.e. $\gamma = 1$, a flipping signal conveys that the policy choice was wrong in one of the two periods[20].

Of course, the truthful behavior of the incumbent cannot be taken for granted: in order to prevent the reputation loss from a flip-flopping action record, the incumbent can ignore the signal and act consistently. This of course comes at a cost, both for the politician distorting his action and for the representative voter. Consider an incumbent about to take his second policy action, with private information at hand that suggests to change the policy compared to what he chose in the first period. Following the signal delivers the highest probability of matching the state of the world (equal to $\rho_{2,f} > 1/2$ versus $1-\rho_{2,f} < 1/2$ if he acts against the signal), but at the same time this makes the incumbent appear incompetent, which costs him $\phi(r_c - r_f)$. Mathematically, the incumbent has the incentive to follow his signal as long as the following holds:

$$2\rho_{2,f} - 1 \geq \phi(r_c - r_f) \tag{9}$$

For high enough reputation concerns, as a result, the incumbent has the incentive to avoid the flip-flop and choose a consistent policy stance despite his private information pointing in the opposite direction. Notice that by the same logic, in the first period, the incumbent has the incentive to follow his signal: on top of being the best choice to match

[20]The intuition of flipping being associated with being wrong is nicely captured by the following quote by Mickey Kaus: *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.*

the state of the world, following the first signal also maximizes the probability of not facing a trade-off between flip-flopping and consistency in the second period[21].

Therefore, a truthful (undistorted) equilibrium is only sustainable if, everything else equal, reputation concerns ϕ are sufficiently low. Using $r_c = r_c^T$ and $r_f = r_f^T$ and letting condition (9) hold as equality yields a threshold on ϕ, which I denote as $\phi_1 \equiv \frac{2\rho_{2,f}-1}{r_C^T - r_F^T}$.

The natural question to ask is now what happens when reputation concerns exceed ϕ_1. As I have just shown, the incumbent does not have the incentive to follow his signal if $s_2 \neq a_1$. To this end, denote by σ the probability $Pr(a_2 = s_2 | s_2 \neq s_1, a_1 = s_1)$. Notice that in the truthful equilibrium, $\sigma = 1$. For a given σ, the reputation from a consistent action record r_c takes the following value:

$$r_c = \frac{\lambda[1 - Pr(f|H)\sigma]}{\lambda[1 - Pr(f|H)\sigma] + (1-\lambda)[1 - Pr(f|L)\sigma]} \tag{10}$$

where, with a slight abuse of notation, I denote $Pr(s_2 \neq s_1 | s_1, \theta) = (1-\gamma)[q_\theta^2 + (1-q_\theta)^2] + \gamma 2 q_\theta (1 - q_\theta)$ as $Pr(f|\theta)$ and its complement $Pr(s_2 = s_1 | s_1, \theta)$ as $Pr(c|\theta)$. From condition (7), then, we obtain that $Pr(c|H) > Pr(c|L)$ and consequently $Pr(f|L) > Pr(f|H)$. Therefore, (10) is increasing in σ. Unlike r_c, the reputation from a flip-flopping track-record, represented by the following expression, does not depend on σ, which cancels out of the expression:

$$r_f = \frac{\lambda Pr(f|H)\sigma}{\lambda Pr(f|H)\sigma + (1-\lambda) Pr(f|L)\sigma} = r_f^T \tag{11}$$

Notice that at the extreme of $\sigma = 0$, $r_c = \lambda$, that is a consistent policy record contains

[21] Notice that if the prior probability of the state of the world was unbalanced in favor of one state, there would be an additional trade-off in the first period when receiving the less likely signal. The aversion to flip-flopping, however, would discipline the incumbent towards following the signal in the first period.

no information on the incumbent's competence.

Suppose now that the truthful equilibrium is not sustainable but that $\sigma = \sigma^*$ is in $(0, 1)$, meaning that the incumbent mixes after receiving $s_2 \neq s_1$. Clearly, σ^* has to make the incumbent indifferent, which requires, following from (9):

$$2\rho_{2,f} - 1 = \phi(r_c - r_f^T) \tag{12}$$

Denoting by $Pr(f) = \lambda Pr(f|H) + (1 - \lambda)Pr(f|L)$ the unconditional probability of the signal flipping and by $k = \frac{2\rho_{2,f}-1}{\phi}$ the cost of not following the signal normalized by the value of reputation, the indifference condition can be rewritten as:

$$\frac{1 - \sigma^* Pr(f|H)}{1 - \sigma^* Pr(f)} = \frac{k + r_f^T}{\lambda} \tag{13}$$

which finally allows me to pin down the value of σ^* after some algebric manipulations[22]. In an equilibrium with $\sigma = \sigma^* \in (0, 1)$, the mixing done by the incumbent decreases r_c, lowering the reputation premium from a consistent action record just enough to make the incumbent indifferent between the two actions. Notice that the left-hand side of condition (13) is strictly larger than 1 as long as $\sigma > 0$: therefore, a necessary condition for the existence of a truthful or mixed equilibrium is that $\frac{k+r_f^T}{\lambda} > 1$, that is that $k > \lambda - r_f^T$. This means the cost k of not following the signal has to be sufficiently high, such that when $\sigma = 0$, i.e. when $r_c = \lambda$, the incumbent prefers to follow his signal. For sufficiently high

[22]Further denoting $\chi = \frac{k+r_f^T}{\lambda}$, (13) can be solved to yield:

$$\sigma^* = \frac{\chi - 1}{\chi Pr(f) - Pr(f|H)} \tag{14}$$

values of ϕ, however, this condition is violated and there is no mixing probability that makes the incumbent indifferent between following the signal and receiving a consistent action record. The only possible outcome is then an equilibrium with $\sigma = 0$, that is with complete flip-flop avoidance. The threshold above which this happens is denoted by ϕ_2, which takes the following value: $\phi_2 = \frac{2\rho_{2,f}-1}{\lambda - r_f^T}$.

The game has a unique Perfect Bayesian Equilibrium, with the following properties:

- For $\phi \leq \phi_1 = \frac{2\rho_{2,f}-1}{r_c^T - r_f^T}$, the unique equilibrium is the truthful equilibrium, that is $\sigma = 1$.

- For $\phi_1 < \phi \leq \phi_2$, the unique equilibrium displays partial flip-flopping avoidance, with $\sigma = \sigma^* \in (0, 1)$, with σ^* defined by condition (13).

- For $\phi > \phi_2 = \frac{2\rho_{2,f}-1}{\lambda - r_f^T}$, the unique PBE displays complete flip-flopping avoidance, that is $\sigma = 0$.

The equilibrium value of reputation following a flip-flopping track-record is equal to r_f^T defined in (11); reputation following a consistent track-record is r_c defined in condition (10). Notice that for all values of ϕ, $Pr(a_1 = s_1|s_1) = 1$ and $Pr(a_2 = s_2|s_2 = s_1) = 1$.

To recap, the first result of my model is that even in the absence of any information on the quality of the policy chosen by the incumbent, some information on his competence can be inferred from his action record. Consistent incumbents are on average more competent than flip-flopping ones. Therefore, the model provides a rationale for the stigma associated with flip-flopping, consistent with the empirical evidence pointing out that flip-floppers are negatively evaluated beyond the effects of partisanship and even in environments in which the optimal policy is supposed to change over time[23].

[23]See for example Carlson and Dolan (1985) and Tomz and Van Houweling (2012). For more details

With respect to the incentives arising from the negative perception of policy changes, Theorem 1 states that when reputation concerns are lower than ϕ_1, reputation concerns have no effect and the equilibrium outcome is equivalent to that of an environment with a policy-motivated incumbent; when instead reputation concerns are higher than ϕ_2, the distortion is maximal and the incumbent plays as if he were fully office motivated. For intermediate values of ϕ, mixing occurs.

In other words, when reputation concerns are high, strategic politicians have the incentive to avoid flip-flops in order to protect their reputation. This leads to an insufficient amount of flip-flopping (and policy change) in equilibrium. Notice that even in an equilibrium with partial flip-flopping avoidance, flip-flops are detrimental to an incumbent's reputation, but the reputation premium for consistency is lower than in a truthful equilibrium.

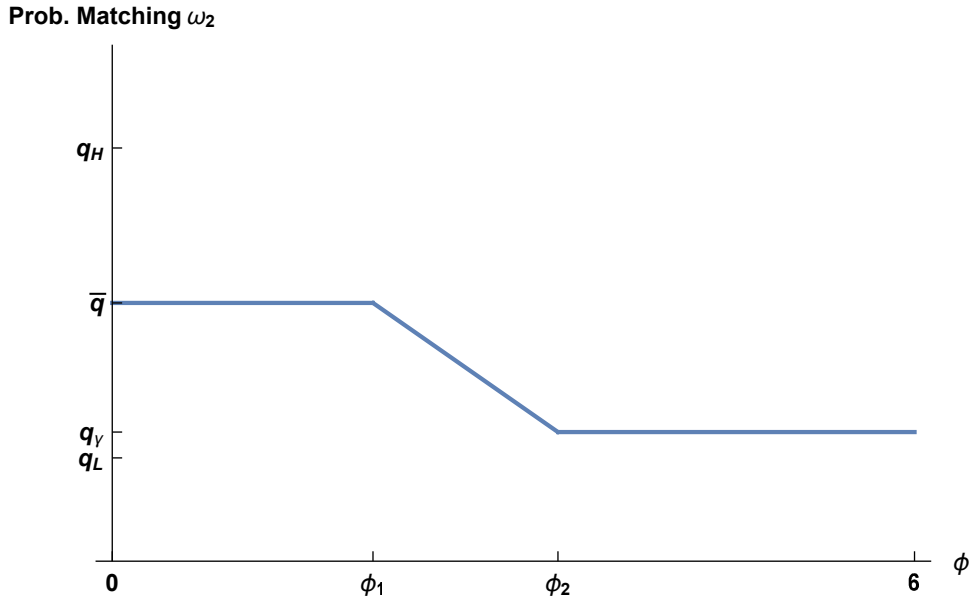

Figure 1: Policy quality in the second period with parameters $\lambda = 0.5$, $\gamma = 0.75$, $q_L = 0.6$, $q_H = 0.9$. Notice that $\bar{q} = \lambda q_H + (1-\lambda)q_L$ and $q_\gamma = \gamma\bar{q} + (1-\gamma)(1-\bar{q})$.

The consequences of flip-flopping avoidance for welfare are two-fold: first of all, the

refer to the literature review.

second policy action is distorted, as the incumbent sometimes acts disregarding his informative signal. This is depicted graphically in Figure 1. Moreover, the more the incumbent avoids flip-flopping, the closer r_c moves to λ, that is the incumbent's track record becomes less and less informative. This decreases the selection component of social welfare[24]. In this model, therefore, reputation concerns such as electoral incentives can only have a negative effect on welfare.

In the analysis summarized by Theorem 1 I have characterized the equilibrium based on the reputation concerns parameter ϕ, keeping all other parameters fixed. However, the model offers other interesting comparative statics. In particular, a crucial role is played by the parameter γ, which measures the persistence of the state of the world between $t = 1$ and $t = 2$. The higher γ, the more informative the first signal (and therefore the first decision) is for the decision to be taken in the second period. As a consequence, the higher is γ the more informative flip-flopping is about the incumbent's incompetence: that is r_f^T is decreasing in γ.

This result has several empirical predictions, that make my model match many of the empirical regularities discovered by the empirical literature. I discuss these implications in the next section.

Analogously to what happens with ϕ, as γ increases the incentives to avoid flip-flops become larger. This means that the thresholds ϕ_1 and ϕ_2 featured in Theorem 1 are decreasing in γ. In other words, beyond the value of γ that I denote as γ_1 at which $\phi_1(\gamma_1) = \phi$, the truthful equilibrium stops being sustainable[25]. As γ grows further, after the threshold where that I denote as γ_2, where $\phi_2(\gamma_2) = \phi$ not even mixing is feasible and

[24]To see this, notice that since $\mathbb{E}r = \lambda$ and $\nu(\cdot)$ is convex, $\mathbb{E}\nu(r)$ decreases if r_c decreases from r_c^T, keeping r_f fixed at r_f^T.

[25]$\phi_1(\gamma)$ reflects the fact that ϕ_1 is a function of γ.

the incumbent completely avoids flip-flops.

Proposition 1. *The reputation from flip-flopping r_f^T is decreasing in γ. However, σ is (weakly) decreasing in γ. That is to say, the thresholds ϕ_1 and ϕ_2 and the equilibrium value of σ are decreasing in γ.*

This has two effects on welfare: the first is that the quality of the second policy decision is non-monotonic in γ: it starts off constant in γ; between γ_1 and γ_2 it is decreasing; finally it increases above γ_2, thanks to the fact that as γ grows, the probability that $s_1 = a_2 = \omega_2$ increases. As a matter of fact, when $\gamma = 1$, i.e. the state of the world is fixed, the probability of matching the state in the second period does not depend on σ[26].

The second effect on welfare has to do with the selection component. For $\gamma \in (1/2, \gamma_1]$, selection welfare increases thanks to the increasing informativeness of consistency as a signal of the incumbent's competence. However, above γ_1 the ability of the representative voter is hindered by the flip-flopping avoidance, until γ_2 above which no learning takes place in equilibrium. Summing up the two components of welfare, as a consequence of Proposition 1 we obtain that $\gamma = \gamma_1$ maximizes social welfare. This is the highest value of γ for which the truthful equilibrium is sustainable. Therefore, the quality of policy-making is maximal, as is the selection component of welfare.

Proposition 2. *The value of γ that maximizes social welfare, both in terms of policy quality and in terms of selection, is equal to γ_1 and strictly smaller than* 1.

Also with respect to this result, in the next section I will go into details about its implications. Concerning other parameters, it is worth mentioning some interesting and

[26]The reason for this is that when $\gamma = 1$, choosing the second policy according to the first signal causes no distortion, since when the signal flips, the signal received in the first period is just as good as that received in the second to determine the second period policy.

slightly counterintuitive non-monotonicities. Take for example the fraction of competent politicians λ: the incumbent is more likely to act in an undistorted way when λ takes values close to the extremes. The intuition is that if the representative voter is very confident about the level of competence of the incumbent, the effect of the track-record on reputation is small. This means that it is possible for social welfare to decrease when λ increases, because the increased proportion of competent incumbents is more than offset by the increased distortions in their behavior. In other words, the *the devil you know* (that is an incumbent with a lower λ) might be better than one who is more likely to be competent (i.e. has a higher λ) but who acts in a less informed way to foster his reputation. A similar effect can also occur with respect to q_H: increasing q_H can in fact lead to lower welfare, since the incentive for flip-flop avoidance increases (given that the accuracy difference between high and low type increases, a flip-flop becomes more telling of the incumbent being low type). In other words, in this model, given an average quality of politicians, the quality of policies is always (weakly) better when the ability of politicians is more equally distributed: in other words, at least from the policy perspective, it is better to have many relatively good politicians rather than a few stars among a majority of incompetents.

Discussion and Implications

Having presented the main results, this section serves the purpose of providing a broader discussion of the implications of the model.

Effects and Interpretation of Persistence

- *Timing of Decisions:* Instead of thinking of γ as the speed at which conditions change, one could think of it as time between decisions. The longer the time, one might argue, the smaller the dependence we should expect between ω_1 and ω_2. Therefore, γ can be thought as being a decreasing function of the time passed after the previous decision: therefore, γ is close to $1/2$ for two decisions far away in time from each other and close to 1 for decisions very close to each other. This has several interesting implications: first of all, also in light of Proposition 1, it fits with the evidence (see Doherty, Dowling, and Miller (2015)) that the closer a policy reversal is to the previous decision, the more negatively it is evaluated by voters. Moreover, it suggests that Theorem 1 can be interpreted as defining how responsive the incumbent politician is as a function of the time passed from his previous decision: because of reputation concerns, a minimum amount of time has to pass in order for the politician to become again responsive to information. An interesting extension for future research would be to follow this intuition to make the timing of decisions endogenous[27].

- *Task Allocation:* along with persistence of the state of the world over time, γ can be more generally interpreted as measuring the degree of similarity of the

[27] More details are available in the Online Appendix.

two tasks addressed by the incumbent politician in his term in office. The more similar two issues are, the more likely the appropriate policy response to both problems is to be the same. In this respect, following Proposition 1, my model predicts that the more related two issues are, the more negatively a policy reversal is judged. Moreover, following Proposition 2, the model suggests that limits to specialization (meant as dealing with the same problem over time) or rotations across sufficiently differentiated tasks, can be welfare improving when dealing with agents with reputation concerns. This might also be seen as a cautionary tale with regards to excessive centralization: bundling multiple decisions in the hands of a single decision-maker can lead to an inefficient *one size fits all* approach motivated by reputation concerns.

Along similar lines, γ could also be interpreted as an index of issue complexity: when an issue is very complex and there are many moving parts, it is intuitive to expect γ to be close to $1/2$, leading to a small effect of flip-flopping on reputation. Conversely, for simpler issues γ is close to 1 and flip-flopping is more detrimental. This is consistent with the empirical evidence reported in Doherty, Dowling, and Miller (2015) and Tavits (2007).

- *Public Information:* it is often the case that voters have a relatively good sense of the fact that conditions changed, but much less so of what is the appropriate policy conduct conditional on a given set of circumstances. The model I describe in this paper is equivalent to the subgame of a larger model, in which after the first policy decision has been taken, players observe the realization of a public signal $\tilde{s} \in \{c, f\}$ telling them whether the state of the world remained constant (c) or

changed (f). The parameter γ can be interpreted as the accuracy of that signal, i.e. $\gamma = Pr(\tilde{s} = c|\omega_2 = \omega_1) = Pr(\tilde{s} = f|\omega_2 \neq \omega_1)$. The subgame following the realization of $\tilde{s} = c$ would be equivalent to my model, whereas the subgame following $\tilde{s} = f$ is equivalent to my model under the assumption of persistence being lower than 1/2. In that subgame, the voter would update upwards instead of downwards following a flip-flop. The incumbent would thus have the incentive to flip-flop, leading to excessive policy volatility. Following Proposition 1, the more public information points to the fact that conditions did not change (respectively, that they changed), the more negatively flip-flopping is evaluated (respectively, the less negatively). This is consistent with the empirical evidence on the effect of information supporting the decision of the incumbent to change policy[28]: the more publicly available information points to the fact that conditions stayed the same, the more negative the reputation effect of a flip-flop (and the other way around).

Importantly, the welfare maximizing value of γ derived in Proposition 2 corresponds the accuracy of the signal on the change of state of the world maximizing social welfare. My model can therefore be interpreted as suggesting that excessive information on whether the state of the world changed is detrimental to social welfare.

Testable Implications from Variation in Electoral Concerns

- *Behavior of Term-Limited Politicians*: It is reasonable to assume that term-limited politicians have a lower ϕ, since they are not up for re-election. Therefore, according to the model we should expect term-limited politicians to change their behavior

[28]See Robison (2017) and Levendusky and Horowitz (2012b).

with respect to previous decisions more than those up for re-election. Moreover, the model suggests that this increased propensity to flip-flop is not due to the lack of accountability caused by term-limits, but to the fact that such politicians are free from reputation incentives that might distort their behavior. As a result, the increased propensity of term-limited politicians to change their policies can be seen in a positive light[29]. In addition to this, the finding that term-limited politicians are more willing to challenge the status quo and enact welfare improving policy changes is a point in favor of term-limits, to be weighed against the cost of not being able to re-appoint incumbents who are believed to be competent.

- *Consistency at unsuspecting times (or issues)*: as I discuss in the previous paragraph, an implication of my model is that politicians should be more willing to reverse their previous decisions when reputation concerns ϕ are lower. In addition to the effect of binding term limits, the level of ϕ can also vary across the career of a politician, or across the type of issue. In this respect, my model predicts that consistent behavior occurring away from the public scrutiny, for example at a moment in time when the politician was not running for an important election (or on an issue not salient for voters), should be rewarded more than consistent behavior at times of strong electoral concerns. The flip-side of this implication is that flip-flops far away in time (or on issues previously off the public radar) are more likely to be used against a politician running for re-election.

[29]Somewhat consistently with my model, List and Sturm (2006) find that US governors are more likely to reverse environmentally friendly policies when term-limited. Lopes da Fonseca (2020) finds that term-limited mayors act in a more fiscally conservative way.

Consistency as Indirect Information on Competence

In the model analyzed in the previous sections, the information on the consistency of the incumbent's actions is the only information available to the representative voter. The following question therefore arises: if other sources of information about the incumbent were available, when would we expect consistency to play an important role for the formation of public opinion? Although my baseline model does not directly address this question[30], it does shed light on the circumstances under which consistency is a stronger signal of competence (and flip-flopping of incompetence). In this sense, an indirect prediction of the model is to characterize the types of environments in which consistency/flip-flopping is more likely to be in the spotlight of the public opinion.

In particular, these circumstances include i) the information contained in the track-record is more accurate, i.e. for example for large values of γ (subject to the value of ϕ, as per Theorem 1 and Proposition 2); ii) additional information on the quality of the decisions is not available, too costly or of bad quality (for example, because there is no government transparency, or because it takes time to acquire information on the state of the world); iii) there is some other obstacle towards conveying direct information on the quality of an incumbent's decision (for example, when the media is perceived as biased, or when the public opinion is polarized).

[30]I do it instead in the Online Appendix, where I extend the model to allow for ω_1 to be revealed before the second action is taken. When the incumbent is proved wrong in the first period, he has an incentive to reverse his policy to match what is likely to be the correct decision also in period 2.

Conclusion

In this paper I present a political agency model of sequential policy-making, which shows how revising previous policy decisions can be detrimental to a politician's reputation. This stems from the fact that even in the absence of any other information on the quality of the decisions taken by an incumbent politician, the sequence of actions taken by the incumbent can provide a signal of incompetence. As a result, politicians facing strong reputation concerns have the incentive to avoid policy changes (that is, to avoid flip-flops). This leads to an excessive degree of policy persistence and also hampers the ability of voters to select competent representatives.

My model is consistent with several stylized facts about the effects of policy changes on reputation: first of all, empirical evidence suggests that consistency is rewarded by voters, often independently of partisanship, and that incompetence is one of the reasons behind the negative evaluations of flip-floppers. Moreover, flip-flops on complex issues are better received than flip-flops on simpler ones. Similarly, there is evidence that policy changes close to each other are more damaging, whereas additional information supporting the decision to reverse the policy has the opposite effect. These are all circumstances in which there are elements in favor of the fact that conditions might have changed, warranting a change in policy response. My model captures such forces through the persistence parameter γ, and is to the best of my knowledge the first study investigating the interaction between persistence and the incentives of politicians taking repeated decisions.

Along with the evidence on flip-flopping, my model also has several additional implications, many of which could be either tested empirically or be the object of further

theoretical investigations. For example, roll-call voting records - together with approval ratings or primary and general election data - could be used to test whether representatives voting in an inconsistent way are more likely to face a decline in approval ratings, primary election challenges or decreases in the shares of votes won in elections[31]. Another directly testable implication of the model is that term-limited policy-makers should be more willing to revise previous policy choices. Similarly, politicians should be more willing to change their mind on issues that are off the public spotlight. Along with the possibility of testing these hypotheses, my model suggests that such behavior need not be considered a worrying sign of the lack of accountability, but to the contrary a welcome consequence of politicians' willingness to update policies in response to new information.

On the theory side there are several interesting open issues. For example, an avenue of future research would be to build on the results of this paper to further explore the substitutability relation between direct and indirect information about the incumbent's competence and the use of consistency as a tool in electoral campaigns. Along similar lines, another relevant question to answer concerns whether along with environmental characteristics such as persistence, there are also personal characteristics that affect the evaluation of a politician's consistency.

[31]I thank an anonymous referee for suggesting this possibility.

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.

Allgeier, A., D. Byrne, B. Brooks, and D. Revnes (1979): "The Waffle Phenomenon: Negative Evaluations of Those Who Shift Attitudinally," *Journal of Applied Social Psychology*, 9(2), 170–182.

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.

Bernhardt, M. D., and D. E. Ingerman (1985): "Candidate reputations and the 'incumbency effect'," *Journal of Public Economics*, 27(1), 47–67.

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.

Carlson, J. M., and K. Dolan (1985): "The waffle phenomenon and candidates' image," *Psychological Reports*, 57(3), 795–798.

Croco, S. E. (2016): "The flipside of flip-flopping: Leader inconsistency, citizen preferences, and the war in Iraq," *Foreign Policy Analysis*, 12(3), 237–257.

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. (1994): "Domestic political audiences and the escalation of international disputes," *American political science review*, 88(3), 577–592.

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

Hoffman, H. S., and C. S. Carver (1984): "Political Waffling: Its Effects on the Evaluations of Observers 1," *Journal of Applied Social Psychology*, 14(4), 375–385.

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

Levendusky, M. S., and M. C. Horowitz (2012b): "When Backing Down Is the Right Decision: Partisanship, New Information, and Audience Costs," *The Journal of Politics*, 74(02), 323–338.

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.

List, J. A., and D. M. Sturm (2006): "How elections matter: Theory and evidence from environmental policy," *The Quarterly Journal of Economics*, 121(4), 1249–1281.

Lopes da Fonseca, M. (2020): "Lame ducks and local fiscal policy: Quasi-experimental evidence from Portugal," *The Economic Journal*, 130(626), 511–533.

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.

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.

Robison, J. (2017): "The role of elite accounts in mitigating the negative effects of repositioning," *Political Behavior*, 39(3), 609–628.

Schultz, K. A. (2001): *Democracy and coercive diplomacy*, vol. 76. Cambridge University Press.

Sigelman, L., and C. K. Sigelman (1986): "Shattered expectations: Public responses to "out-of-character" presidential actions," *Political Behavior*, 8(3), 262–286.

Smith, A. (1998): "International crises and domestic politics," *American Political Science Review*, 92(3), 623–638.

Sorek, A. Y., K. Haglin, and N. Geva (2018): "In capable hands: An experimental study of the effects of competence and consistency on leadership approval," *Political Behavior*, 40(3), 659–679.

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. Van Houweling (2012): "Political Repositioning," *Unpublished Manuscript, Stanford University*, pp. 1–34.

Tomz, M. (2007): "Domestic audience costs in international relations: An experimental approach," *International Organization*, 61(4), 821–840.

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

Tomz, M., and R. P. Van Houweling (2014): "Political repositioning: A conjoint analysis," *Unpublished manuscript, Stanford University*.