

## ***Electoral Competition, Control and Learning***

TORUN DEWAN AND RAFAEL HORTALA-VALLVE\*

This article explores an agency model in which voters learn about both an incumbent and an opponent. They observe the incumbent's policy record and update their beliefs about his opponent via a campaign. Although the former is relatively more informative, it can be costly for the voter to learn about the incumbent from her policy record. This is because policy reforms, which allow a voter to learn an incumbent's ability, are risky and can leave the voter worse off. Then the voter may prefer the incumbent to take safer actions. The efficient level of reform – the one preferred by the voter – balances the value of learning with the expected policy costs/benefits. In a world where the opponent's campaign is uninformative, reform can be too low due to the incumbent's *fear of failure*. Or it can be too high: the incumbent may *gamble on success*. This article shows that the presence of an opponent who can reveal information via a campaign exacerbates these inefficiencies. An incumbent who anticipates the effect of an opponent's campaign on voter beliefs is more likely to make inefficient policy choices. Further, such campaigns can lead to an overall welfare loss when they reveal little about the opponent's ability and yet have an impact on the incumbent's policy choice.

*Keywords:* incumbent vs. opponent; policy reform; informative campaigns

Most formal models of electoral competition assume perfect symmetry between candidates, and so do not consider the different roles played by an incumbent and his or her opponent. They may even systematically ignore the impact of the latter. In the standard spatial model, prospective voters compare the policy programmes of candidates but do not account for the incumbent's past performance. In agency models, voters form expectations based on the incumbent's past performance but do not learn about his opponent.<sup>1</sup> In this article, we develop a new agency model in order to study the interplay between an incumbent and an opponent that incorporates critical aspects of the asymmetric relationship between them.

Central to our model is the fact that voters' beliefs about *both* an incumbent and his opponent depend upon observations of their performance. The idea that voters update their beliefs based on observations of opposition performance is supported by anecdotal evidence. Consider data from the 2010 UK general election that is illustrated in Figure 1. Voters' perceptions of the best choice of prime minister were stable until the dissolution of Parliament and the start of the campaign on 12 April. Thereafter, there was a sharp increase in the percentage of voters supporting Nick Clegg, the leader of the Liberal Democrats (who went on to become deputy prime minister).

\* Government Department, London School of Economics and Political Science (emails: t.dewan@lse.ac.uk, r.hortala-vallve@lse.ac.uk). We thank Scott Ashworth, Steve Callander, Georgy Egorov, Francesco Giovannoni, Jean Guillaume Forand, Stuart Jordan, Navin Kartik, Mik Laver, Ben Lockwood, Pablo Montagnes, David Myatt, Ken Shepsle, Jim Snyder, Francesco Squintani, and seminar audience members at Berkeley, Bristol, Columbia University, the Harris School of Public Policy, the Higher School of Economics in Moscow, Mannheim, NYU, Princeton, Stanford GSB, Warwick, and participants at the Annual Meetings of the American Political Science Association, the II Workshop on Institutions, Individual Behavior and Economic Outcomes in Alghero, and the Midwest Political Science Association. Online appendices are available at <http://dx.doi.org/10.1017/S0007123416000764>.

<sup>1</sup> See recent articles by Dewan and Shepsle (2011) and Ashworth (2012) for a review of the agency and spatial modelling approach to elections.

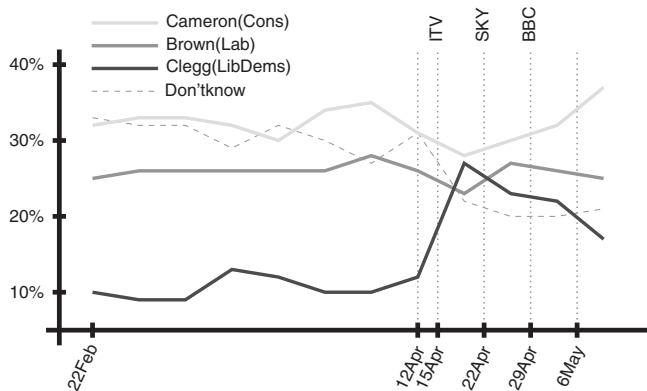


Fig. 1. Voters' perceptions of the best choice of PM, UK general election 2010

There was a sharp corresponding decrease in the percentage of 'don't knows'. This perception continued through the time of the first televised debate between the party leaders: 27 per cent of voters believed that Clegg would make the best prime minister two days after the first televised debate, up from 12 per cent in the same poll taken five days earlier. Again, this corresponded to a sharp decline in the percentage of 'don't knows'. The evidence suggests that Clegg's campaign influenced respondents' beliefs about his suitability to be prime minister. Indeed, by the time of the second televised debate, more people thought he was a better choice than the incumbent, Gordon Brown. This change in attitude occurred despite Clegg's lack of executive experience. More generally, we surmise that successful campaigns can lead to an increase in a candidate's perceived executive competence, whereas unsuccessful ones can have the opposite effect.

These data also illustrate a second aspect of our model. Although the campaign appeared to have an effect on previously undecided voters (who now favoured Clegg), it had no apparent effect on evaluations of the incumbent. This is perhaps unsurprising. After all, by the time of the campaign, voters had already learned about the incumbent from his performance in office. This suggests that voters learn using role-specific technologies. The point is nicely made by Bawn and Somer-Topcu, who argue that 'voters can evaluate government parties on the basis of recent performance, but they must judge opposition parties on the basis of rhetoric and conjecture'.<sup>2</sup>

In this article, we develop and analyse a model that builds on these insights. Here, voters learn about an incumbent from his track record and about the opponent from his campaign. We make the plausible assumption that an incumbent's track record is more informative about his competence than is a campaign about that of his opponent.

We use our model to explore the policy implications of this asymmetry. Despite a commonly held view that campaign talk is cheap, anecdotal evidence suggests a link between opposition campaigns and incumbent policy. In the United Kingdom, for example, respective Labour (Conservative) incumbent governments have adopted opposition policies announced during the campaign. There are different mechanisms that could explain this link. For example, these politicians may have information that such policies are popular, or have become convinced of their merits. However, systematic evidence from Adams and Somer-Topcu suggests that policy proposals of governing and opposition parties are consistent with a best-response logic.<sup>3</sup>

<sup>2</sup> Bawn and Somer-Topcu 2012. Indeed, the difference in the technology available to an incumbent and his opponent has been used to explain the empirical regularity that campaign resources translate into higher vote shares for the latter but not the former (e.g., Jacobson 1978).

<sup>3</sup> Adams and Somer-Topcu 2009.

Correspondingly, we assess the policy implications of the strategic interaction between incumbents, opponents and voters.

In our model, voters learn about the incumbent by observing the policy outcomes associated with his decisions. Specifically, an incumbent chooses whether to implement a project or policy reform that is successful if and only if he is competent enough to implement it; otherwise, it results in failure. Rather than take such ‘risks’, the incumbent can, if he chooses, play it ‘safe’ by doing nothing or sticking with existing policy. In this case nothing can be inferred about his competence. We first explore the incumbent’s choice when a voter has no information about the opponent (as in the canonical model). We then analyse the incumbent’s choice when anticipating that the opponent can reveal information about his (the opponent’s) abilities via a campaign.

We ask whether politicians’ incentives to invest in risky policies are aligned with the social returns of doing so. The efficient level of reform – the one preferred by the voter – balances the value of learning with the expected policy costs/benefits. We note that even negative policy outcomes (or blunders) have social value due to learning: the successful or unsuccessful execution of a risky policy is positively correlated with the outcomes of future policy experiments. Despite this, we show that (relative to an efficient benchmark), when the opponent cannot reveal information via his campaign, levels of risk taking can be too low. Our analysis agrees with arguments made by Harford, who highlights differences in risk taking across sectors.<sup>4</sup> In the business world, he argues, success is built on previous failure, whereas, ‘in politics where are the bad ideas that have been tested, found wanting, and replaced with something better?’ Our model relates such underinvestment in risk to electoral accountability and the ‘fear of failure’ that is induced by the winner-take-all nature of electoral competition.

Perhaps surprisingly, however, we show that electoral accountability can also lead to overinvestment in reform in equilibrium. When subject to an election, an incumbent may be *too* willing to take the risky option. That is, he implements reforms that, on balance, do not benefit the voter. We label this effect ‘gambling on success’.

To illustrate, consider the case of an incumbent who chooses whether to build a public project like a bridge or a major exhibition centre. The voter and incumbent agree on the value of such a project (neither is ideologically committed to it). Suppose that voters learn something about the incumbent’s ability from the execution of this project and believe it to be (positively) correlated with his ability to implement other policies. Suppose further that when taking the expected returns into account the voter is not in favour of the project. The incumbent might accede to the voter’s wishes. Should he do so, however, then the voter will be unable to distinguish between him and his opponent. By contrast, if the incumbent builds the bridge and is successful, a sequentially rational voter will retain him. And so, anticipating this, the incumbent builds the bridge.

Beyond our parable, recently gathered empirical evidence suggests that learning may indeed drive overinvestment in public projects. Voigtlander and Voth study the electoral benefits of the world’s first nationwide highway network built by the Nazis.<sup>5</sup> They show that the German autobahn, ‘a canonical case of public infrastructure development’, contributed to Nazi electoral success and to the perception of the regime’s competence, despite the fact that neither the economic or military case for road development was well developed.<sup>6</sup>

<sup>4</sup> See <http://www.freakonomics.com/2011/05/10/why-is-failure-a-sign-of-a-healthy-economy>

<sup>5</sup> Voigtlander and Voth 2014.

<sup>6</sup> Comparing constituency-level differences in Nazi support between the 1933 parliamentary elections and the 1934 plebiscite, the authors find that opposition to the Nazis was lower in areas where the autobahn had been built. In making the case that highway construction enhanced the Nazis’ reputation for competence, they dismiss alternative explanations. The economic rationale for the autobahn was unclear, as ‘car ownership rates in Germany in 1933 were low, approximately one quarter of those in England or France’; few used the new

Other policy interventions illustrate the notion of ‘gambling on success’. For example, consider British Prime Minister Margaret Thatcher’s decision to send a task force to the Falkland Islands in 1982. The islands were populated by 1,800 British subjects and under occupation by Argentinean forces. During hostilities, 900 people, including 255 British military personnel, lost their lives. The war cost the exchequer around £3 billion. At the outset, only the most hawkish, such as Admiral Sir Henry Leach, head of the Navy, gave the task force much chance. *Ex ante* the choice of war could not be justified. The rest is history. As noted by David Cannadine in his obituary to Thatcher: ‘She gambled everything on getting back the Falkland Islands from Argentina, and she scored a spectacular triumph. But it was a close-run thing: If she had lost she would have been the most derided British Prime Minister since Anthony Eden.’<sup>7</sup>

Our analysis shows that the incumbent places either too little weight on learning relative to the social return (‘fear of failure’) or too much (‘gambling on success’). In both cases, such under/overinvestment is driven by the incumbent’s concern for his career and subsequent elections. Our key question is: What effect do opposition campaigns have on the incumbent’s policy? Moving beyond our baseline model we explore a world in which, as in our motivating example, voters can update their beliefs about an opponent on the basis of his campaign. A conjecture is that access to such information should realign incumbents’ incentives. This is not the case, however. Far from realigning the incumbents’ incentives with those of voters, we show that opponents’ campaigns can exacerbate problems of the ‘fear of failure’ and ‘gambling on success’. That is, policy reform will always be either too high or too low relative to the efficient benchmark. Moreover, comparing cases reveals that when the voter is unable to learn about the opponent’s policy, the resulting policy is more in line with her wishes.

Performing a fuller welfare analysis, we consider a voter’s ability to make more informed choices under asymmetric competition and its policy implications. We show, surprisingly, that such competition can have negative overall consequences. These relative welfare losses arise precisely when little information is gleaned about executive ability from an opponent’s campaign, though it is sufficient to affect the incumbent’s incentives.

Our article is organized as follows. After discussing related literature, we introduce our model and benchmark results with an opponent who is a random draw from a known distribution. Our main results, with an active opponent, then follow. Finally, in the Appendix, we perform several robustness checks: we analyse voters who can commit to the use of their optimal retention strategies within the class of equilibrium strategies, focus on specific fixed rules that are akin to voter biases in favour of (or against) the incumbent, and relax other core assumptions. These different model specifications yield results that are qualitatively similar to those obtained in our core model.

#### RELATED LITERATURE

We explore learning in the context of the multi-armed bandit model, used in policy analysis by Aghion et al.,<sup>8</sup> developed further by Banks and Sundaram,<sup>9</sup> and adapted by Strumpf<sup>10</sup> to look at

(*F* note continued)

autobahns, and road building was a weak aspect of demand stimulus in post-Weimar Germany; though road building delivered local benefits in the form of employment, ‘areas with high unemployment did not show a greater increase in Nazi support when roads were built’; finally, the Nazis had no particular ideological conviction in favour of road building.

<sup>7</sup> ‘How Should we Rank Margaret Thatcher?’ *New York Times*, 14 April 2013.

<sup>8</sup> Aghion et al. 1991.

<sup>9</sup> Banks and Sundaram 1993, 1998.

<sup>10</sup> Strumpf 2002.

policy innovation and its relation to government decentralization and by Strulovici<sup>11</sup> to analyse experimentation by groups of decision makers. Our focus on electoral competition relates most closely to Banks and Sundaram;<sup>12</sup> they analyse an infinite armed bandit problem in which a principal selects a candidate with a single action that yields a reward (to the principal) according to the agent's type. In a recent contribution, Hirsch analyses learning where the principal and agent share the same intrinsic motivation but have potentially different preferred policy instruments.<sup>13</sup> The policy environment in our model is closest to that used by Lizzeri and Persico, who study the impact of different electoral systems on risk control.<sup>14</sup> In our model, elections allow for learning about an incumbent's type via her policy choice. Others focus on learning about policy: Volden, Ting and Carpenter<sup>15</sup> analyse a situation in which governments learn from their experiences and those of other governments, Callander<sup>16</sup> looks at learning by trial and error in a business environment, Callander<sup>17</sup> explores elections, and Callander<sup>18</sup> investigates learning in bureaucracies.

The novelty in our model is the role played by an opponent. This aspect relates our work to that of Ashworth and Shotts, in which an incumbent engages in costly information acquisition before choosing a policy.<sup>19</sup> The role of the challenger is to assess and criticize the incumbent's platform, which also depends on information acquisition. When the opponent's claims are verifiable, the incumbent exerts more effort and so challengers improve welfare. Our focus is on adverse selection, where voters learn about the incumbent's type (competence) rather than the correctness of her decision.

In our model, the incumbent and opponent have access to different technologies, which they use to convince voters of their competence. The assumption builds on a vast literature that highlights that translating campaign resources into votes is much easier for opponents than incumbents. In a related model, Daley and Snowberg look at the constrained choice of an incumbent between a safe and a risky option in a multi-task model.<sup>20</sup> They show that the optimal allocation of resources by a high-quality incumbent does not maximize voter welfare (campaigning is too high), though the voter cannot commit to disregarding information that campaigns convey. Our contribution is to model the interaction between competing politicians with different technologies.<sup>21</sup> In our model, voters cannot commit to ignoring information from the opponent's campaign, and this induces a higher-than-optimal level of reform by the incumbent. Over-reform due to competitive interaction in a campaign between politicians is also a feature in Prato and Wolton.<sup>22</sup> Relatedly, Bektimirov and Montagnes<sup>23</sup> extend the classic model of electoral accountability to explore aspects of a specific set of reforms, namely privatization.<sup>24</sup>

<sup>11</sup> Strulovici 2010.

<sup>12</sup> Banks and Sundaram 1990.

<sup>13</sup> Hirsch 2011.

<sup>14</sup> Lizzeri and Persico 2009.

<sup>15</sup> Volden, Ting, and Carpenter 2008.

<sup>16</sup> Callander 2011a.

<sup>17</sup> Callander 2011b.

<sup>18</sup> Callander 2008.

<sup>19</sup> Ashworth and Shotts 2011.

<sup>20</sup> Daley and Snowberg 2011.

<sup>21</sup> Relatedly, Gul and Pessendorfer (2014) analyse electoral competition with asymmetric access to campaign funds.

<sup>22</sup> Prato and Wolton 2015.

<sup>23</sup> Bektimirov and Montagnes 2014.

<sup>24</sup> Barro 1973; Ferejohn 1986.

Our analysis reveals inefficiencies that arise due to incumbents' career concerns. This relates our model to the literature on 'pandering', in which informed politicians implement populist choices that do not benefit voters.<sup>25</sup> Extensions of this framework look at different aspects of elections and constitutional design.<sup>26</sup> While these models incorporate policy differences, office-seeking concerns and asymmetric information, one of our contributions is to show that similar inefficiencies arise even with symmetric information and shared preferences: the career concerns of the incumbent are sufficient for him to implement an inefficient level of reform. We relate such inefficiencies to information about an opponent and thus contribute to the analysis of accountability where voters have access to multiple information sources. In Ashworth and Shotts, media comments on incumbent policy can either attenuate or exaggerate populist pandering.<sup>27</sup>

#### A MODEL OF LEARNING AND ELECTIONS

We develop an agency model in which an incumbent politician of an unknown type takes one of two possible actions in each of two periods. As in the standard agency model, a voter can learn about the incumbent's type by her first-period action and has the option of either retaining him or replacing him with an alternative. Our key innovation is in modelling the strategic interaction between the first-period incumbent and her opponent. Specifically, we model the strategic interaction between two politicians  $j \in \{i, 0\}$ , where  $i$  is the incumbent at the beginning of the game and  $o$  her opponent, and a representative voter who chooses between them according to their perceived ability. Neither the voter nor the politicians know the ability of the latter prior to any actions taken.

In each of two periods, one of two policies can be adopted: the first is a 'safe' policy, whilst the second is 'risky'. To illustrate, we can think of the risky policy as the implementation of a reform, whilst the safe policy involves sticking with the status quo. One example is where the first-period salient issue is whether to implement budget reforms or not. In the second, it is whether to implement constitutional reform. The absence of reform entails the current budget and constitutional arrangements remaining in place. The outcome of implementing the risky policy can either be a success or a failure. The determining factor in this outcome is the politician's competence at implementing such a policy. The first-period policy reform is successful if the incumbent has the requisite skills to execute it. And the same is true of the second-period reform. Competence across risky policies is positively correlated. This implies that a success or failure in the first period is informative about the incumbent's competence in period 2's risky task. A successful outcome of a risky policy is always beneficial to the voter and the politician implementing it, relative to the pay-off from the status quo.

The innovation in our model is the introduction of an active opponent in the first period. After the incumbent has taken her first-period action and the outcome is revealed, the opponent launches a campaign that has no direct cost attached. As with the policy choice of the incumbent, a campaign can be either safe or risky. An example of a safe campaign is where the opponent emphasizes values, such as patriotism or the need for a strong economy, that are broadly shared, though does not provide detailed policy prescriptions. We think of such a campaign as safe as it is unlikely to have a strong impact, either positive or negative, on how the voter perceives her competence for office. A risky campaign, by contrast, is one in which the opponent provides detailed policy alternatives or a strong criticism of the incumbent's record.

<sup>25</sup> Canes-Wrone, Herron, and Shotts 2001; Maskin and Tirole 2004.

<sup>26</sup> Fox 2007; Fox and Stephenson 2011.

<sup>27</sup> Ashworth and Shotts 2010.




Fig. 2. Timeline

Such a campaign is more revealing in that it showcases her ability to plan policy effectively and to communicate clearly. A risky campaign is successful in that the opponent is able to make the case for reform or launch a successful attack on the incumbent; otherwise it ends in failure. As with the incumbent, the determining factor in a successful campaign is the competence of the politician. Once again we assume that being competent at running a campaign is positively correlated with competence in period 2’s risky task. After the outcome of the opponent’s campaign is revealed, the voter chooses whether to retain the incumbent or not.

Figure 2 describes the timeline of the game. The incumbent chooses the first-period policy. If he has chosen the risky policy, the voter learns whether it has been a success or failure. The opponent then chooses his campaign strategy and the voter learns (if he has chosen the risky policy) about his executive ability. Voters then evaluate both the incumbent and the challenger before electing one or the other in a winner-take-all contest. The winner of this contest then implements the second-period policy and the game ends. In what follows we focus on the perfect Bayesian equilibria of this game.

A sequentially rational voter cares about her expected second-period pay-off. She benefits from implementing the second-period reform only if it is implemented by a competent politician. Her posterior beliefs, based on her first-period observations in each of the relevant sub-games, determine her optimal retention rule. The voter’s concern is the correlation between a successful first-period performance and the successful implementation of the second-period reform. Whilst some skills, such as the effective consultation, communication and orchestration of a bill through the legislative chamber, are required for the successful implementation of all reforms, others require specific skills. This observation suggests a positive (albeit imperfect) correlation between the incumbent’s competence in both periods. Similarly, successful campaigning by the opponent also involves skills, such as effective policy planning and communication, that are relevant but imperfectly correlated with the ability to successfully implement reform in period 2.

To formalize the idea of task- and period-specific qualities, we assume that the type of politician  $j$  in period  $t$  is  $\theta_j^t \in \{1, 0\}$ . When  $\theta_j^t = 1$ , the politician is competent and therefore successful when taking the risky action in period  $t$ . When  $\theta_j^t = 0$  she is not. For notational simplicity, but without loss of generality, we assume that the politician’s competence is perfectly correlated with the outcome of the risky action.<sup>28</sup> The prior probability that a politician is competent on any task is  $p$ .

Our assumption that the incumbent’s competence across periods is positively correlated then leads to the following specification for the voter’s beliefs:

$$Pr\{\theta_i^2 = 1 \mid \theta_i^1 = 1\} = p^H > p = Pr\{\theta_i^2 = 1\}$$

$$Pr\{\theta_i^2 = 1 \mid \theta_i^1 = 0\} = p^L < p = Pr\{\theta_i^2 = 1\}.$$

<sup>28</sup> For our results to hold, we simply need to assume that the probability of a successful risky policy is strictly greater when  $\theta = 1$  than when  $\theta = 0$ .



If the incumbent implements the safe policy, then the prior assumption about his competence with respect to the second-period policy is unaffected. The posterior probability that the incumbent is competent in the implementation of the second-period reform is greater when running on a successful record of implementing reform ( $p^H > p$ ). It is smaller when he runs on a record of previous failure ( $p^L < p$ ).

A successful campaign shows that the opponent can plan and communicate policy. As discussed above, this also makes it more likely that he will successfully deliver the reform policy should he be elected and choose to do so. Incorporating this assumption, we then have that:

$$Pr\{\theta_o^2 = 1 | \theta_o^1 = 1\} = p^h > p = Pr\{\theta_o^2 = 1\}$$

$$Pr\{\theta_o^2 = 1 | \theta_o^1 = 0\} = p^l < p = Pr\{\theta_o^2 = 1\}.$$

When the opponent runs a safe campaign the voter's prior is unaffected. The posterior probability that the opponent is competent to implement the reform is higher when she has showcased her ability in the campaign, and lower when she was unsuccessful at running the risky campaign. Our setup draws attention to an interesting aspect of electoral competition, namely the ability of the opponent to influence the voter's perception of his ability via his endogenous choice of campaign. We assume that information revealed during the course of the campaign is less informative than that obtained from observing the incumbent's track record: Formally,  $p^L < p^l < p < p^h < p^H$ .<sup>29</sup>

Given these beliefs, it is straightforward to describe the sequentially rational actions of the voter in each of the relevant sub-games. In several sub-games, the action of the opponent is irrelevant: the voter always retains an incumbent who successfully implements reform irrespective of the opponent's action, and she never retains one who fails when implementing reform.<sup>30</sup> In the remaining sub-games, however, the actions of the opponent will influence her decision. When the incumbent plays safe, a risky campaign can determine the election outcome. If the opponent succeeds, the voter replaces the incumbent; if he fails, she does not. Finally, when both play safe then the voter is strictly indifferent between retaining the incumbent or not. As we shall see, our results are robust to assuming different re-election rates in this case.

Table 1 summarizes the sequentially rational retention rule that is implemented by the voter in each of the sub-games in which she takes an action. For most of our analysis we assume that when the voter has no information from the incumbent's record or from his opponent, then she re-elects the former with fixed probability  $x \in (0, 1)$  though to illustrate some of our results we set  $x = 1/2$ . However, later on when checking the robustness of our main results, we allow for the endogenous choice of  $x$ . In our design the role of the opponent is limited due to the relatively weak technology available to him. Nevertheless, as we shall see, outcomes are markedly different in a world where the opponent can reveal information about himself via his campaign strategy.

Finally, we complete the specification of pay-offs. Implementation of the status quo yields a pay-off of 1 to the politician implementing it and to the voter. A successful reform yields

<sup>29</sup> In the Appendix we derive these inequalities from the fact that the voter learns more from the track record than the campaigns.

<sup>30</sup> To simplify, we assume that the voter elects the candidate with the highest posterior probability of being competent (with respect to policy implementation in period 2). While this is the optimal course of action for a sequentially rational voter, in some circumstances she might be indifferent between her choices (e.g., when the opponent plays safe and the incumbent has been unsuccessful, thus yielding an expected pay-off of 1 in period 2). All of our results are, however, robust to relaxing this convenient assumption.



TABLE 1 Probability of Re-electing the Incumbent

Incumbent	Opposition		
	Safe campaign	Success in risky campaign	Failure in risky campaign
Safe policy	$x$	0	1
Success in risky policy	1	1	1
Failure in risky policy	0	0	0

pay-off  $r > 1$  to both. A failed reform, however, yields a pay-off of 0. The key difference between the politicians and the voter is that the former obtain a positive pay-off only when they implement the policy. Thus the second-period pay-off of a former incumbent removed from office, or that of an opponent who is not elected, is zero.

NONINFORMATIVE CAMPAIGNS

We begin our analysis by focusing on the case where the opponent’s action is not relevant to the voter’s choice. This is so when his campaign provides no information about his competence for office. The voter then believes him to be competent with prior probability  $p$ . In this context, it is useful to establish the efficient outcome: the policy choice of the incumbent that maximizes voter utility. If the incumbent chooses safe, then he learns nothing and so implements safe again in the next period (yielding a total pay-off of 2). Playing risky allows him to learn his competence with respect to the first-period reform and so update his beliefs that he is competent when implementing the second-period policy. Straightforwardly, an incumbent who maximizes voter welfare plays:

$$\text{risky in period 1} \Leftrightarrow p(r + \max\{1, p^H r\}) + (1-p) > 2. \tag{1}$$

A simple calculation provides the following result:

LEMMA 1: An incumbent who maximizes voter welfare implements the risky policy in period 1 when  $pr > \frac{1+p}{1+p^H}$  and  $p^H r > 1$  and plays safe otherwise.

The efficient policy choice balances the voter’s objectives: she wants the successful implementation of reforms and to learn about the incumbent’s competence. An interesting feature is that, from the voter’s perspective, reform is desirable even though its expected pay-off is less than that obtained when playing safe ( $pr < 1$ ). This highlights the importance of learning. An incumbent who chooses the risky policy in the first period and is successful will do so again in the second period. Learning about the incumbent’s type can improve the voter’s expected pay-off.

Comparative static analysis is straightforward. An increase in the correlation in task-specific competencies increases the value of first-period learning: fixing  $r$ , a welfare-maximizing incumbent engages in more first-period risk taking for larger  $p^H$ . Fixing  $p^H$ , the incumbent engages in more risk taking when the policy pay-off conditional on success is larger. As a result, beneficial reforms (large  $r$ ) are implemented more often.

Next, we consider the actions of an incumbent who is concerned with retaining office rather than maximizing voter welfare. His career concern may lead him to distort his behaviour. The voter has a blunt tool, namely her retention rule, that she uses to realign the incumbent’s incentives. Straightforwardly, given her beliefs, the voter retains an incumbent who successfully

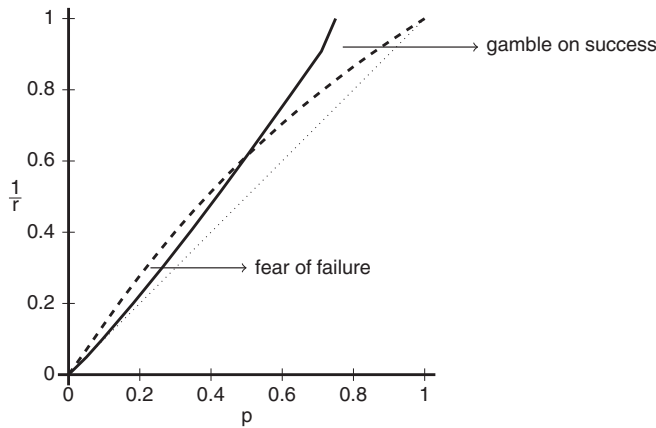


Fig. 3. Equilibrium with a non-informative campaign and  $x = 1/2$

implements first-period reform and replaces one who tries but is unsuccessful. However, when the incumbent plays safe the voter is strictly indifferent between her strategies and so adopts her fixed retention rule  $x \in (0, 1)$ . Given this retention rule, an incumbent who cares about maintaining office plays

$$\text{risky in period 1} \Leftrightarrow p(r + \max\{1, p^H r\}) > 1 + x.$$

Figure 3 illustrates this strategy for the most intuitive rule,  $x = 1/2$ .<sup>31</sup> The 45° (dotted) line defines  $pr = 1$ , where the expected returns from safe and risky are equal. Above this line, the first-period expected pay-offs from choosing risky are lower than those from choosing safe. Two more curves are depicted. The dashed curve illustrates the efficient choice: to its right the voter wants the incumbent to take risks and to its left to play safe. The solid curve separates the parameter space according to the incumbent’s equilibrium actions – to its right she chooses risky, to its left she plays safe – and cuts the efficiency curve from below just once. A clear implication is that the level of risk is either too high (at high  $p$ ) or too low (at low  $p$ ) relative to the efficient benchmark.

At low levels of competence the politician underinvests in the reform policy relative to the efficient benchmark established in Lemma 1. She fears that the risky action will result in failure and loss of reputation compounded by loss of office. We refer to this effect as a ‘fear of failure’ that may lead to inefficiently low levels of reform. At high levels of competence, by contrast, the politician will overinvest in the risky policy relative to the efficient benchmark. In these situations, although on average the incumbent’s competence is high, the value of the risky option is relatively low. While voters would rather she play safe, the incumbent anticipates success when implementing reform. Moreover, a (sequentially rational) voter will then re-elect her. Thus the incumbent has an incentive to take inefficient risks. We refer to this incentive as ‘gambling on success’. These effects are labelled in the appropriate parameter regions in Figure 3.

Both effects (‘fear of failure’ and ‘gambling on success’) stem from the incumbent’s career concern. His probability of retaining office following the implementation of reform depends on the voters’ posterior beliefs. When successfully implementing reform, these jump from  $p$  to  $p^H$  and the incumbent is sure to be elected. When the incumbent is unsuccessful, by contrast, the voter’s posterior decreases from  $p$  to  $p^L$ . In this equilibrium, the incumbent places too much

<sup>31</sup> In the figure we consider  $p^H = p^{0.3}$ .

weight on the impact of the policy outcome on voter beliefs (and hence her chances of re-election). Correspondingly, he places insufficient weight on the benefits of policy.

We note that for the specific values of  $p^L$  and  $p^H$ , underinvestment occurs in the risky policy for  $p < 1/2$ , whereas investment is too high for  $p < 1/2$ . We end this section by showing that this is true whenever  $x = 1/2$  and, moreover, that our insights extend to a situation with any arbitrary  $x$ . An interpretation is that if  $x > 1/2$  a voter has a systematic bias in favour of the incumbent: although she believes both politicians to be of the same competence, she favours the incumbent: this is a form of incumbency advantage. By contrast, if  $x < 1/2$  then she is biased in favour of the opponent – a form of incumbency disadvantage.

**PROPOSITION 1:** When campaigns are uninformative ( $p^h = p$ ) about the opponent's competence and the voter re-elects the incumbent with probability  $x \in (0, 1)$  then the incumbent underinvests in the risky policy (*fear of failure*) for  $p < x$  and overinvests (*gambling on success*) for  $p > x$ .

#### INFORMATIVE CAMPAIGNS

Next we consider the impact of a campaign by political opponents that may influence voters' perceptions about their competence. We study a world in which such information, while relevant to the voter's choice, is less informative than that obtained via the track record of the incumbent (specifically,  $p < p^h < p^H$ ). Following our earlier line of enquiry, we ask: what is the effect of an opponent's campaign on the efficient level of risk taking by the incumbent? We then explore the equilibrium level of risk taking with an active opponent. In dealing with the first question we find the following:

**PROPOSITION 2:** For any value  $x \in (0, 1)$ , and although a voter prefers less risk taking when campaigns are informative, the incumbent is more likely to implement the risky policy.

The first part of the result follows from the fact that the opponent's campaign provides a costless source of information. Therefore, quite naturally, the (marginal) benefit of learning about the incumbent through the implementation of the risky policy is lower than when campaigns are uninformative. As a consequence the voter prefers less risk taking.

The second part of the result stems from the incumbent's career concern. To demonstrate, consider again our 'bridge-building' parable, in which the incumbent could impress a voter by successfully constructing a bridge. Now add an opponent to the mix. His campaign might impress the voter. Specifically, it might persuade her that his executive ability is higher than her prior opinion would suggest. This shapes the incumbent's incentives: anticipating an opponent's campaign makes it more likely that the incumbent will use risky policies to inform the voter of his executive ability. A key insight is that information about an opponent that allows voters to make more informed choices can indirectly affect the incumbent's policy choice.

Next we look at the equilibrium policy outcomes under the incumbent and compare these with the efficient ones.

**PROPOSITION 3:** When campaigns are informative about the opponent's executive competence, then investment in the risky policy is too low (*fear of failure*) for small  $p$  and too high (*gambling on success*) for large  $p$ .

The first part of the proposition reveals that the introduction of informative campaigns does not resolve the problem of underinvestment in risk when the *a priori* competence of politicians is low.

To demonstrate, consider a case in which, for low  $p$ , there is nevertheless a high correlation between the opponent's campaigning ability and her executive competence. In this case the opponent's best response when the incumbent plays safe is to run a risky campaign. The incumbent's initial choice is optimal: she anticipates the opponent's response and, moreover, that he will fail to convince voters of his competence. Thus strategic interaction between the incumbent and her opponent reinforces the 'fear of failure' (an inefficiently low level of reform).

The second part of the proposition establishes that informative campaigns can also reinforce incentives to 'gamble on success', thus inducing an inefficiently high level of reform. The result follows directly from previous ones. As shown in Proposition 2, the efficient level of risk taking is lower when campaigns are informative and yet the incumbent is incentivized to take more risks. It follows that investment in the risky policy will be too high.

Indeed, a straightforward comparison reveals that inefficiently high levels of risk are more prevalent when campaigns are informative than when they are not. As an illustration, and returning to our bridge-building parable, note that, as a consequence of Proposition 2, there are situations in which the voter would prefer the incumbent to build the bridge when the opponent's campaign is informative and that he refrain from doing so when it is not. Suppose that the incumbent acts according to the voter's wish not to build the bridge. Should the opponent play safe, then the voter is unable to distinguish between them. In fact (for large  $p$ ) the opponent's best response to safe play by the incumbent is to seize the opportunity to impress the voter with his campaign. Should he be successful when doing so, then, absent any evidence to the contrary, the voter is persuaded that he is in fact the better choice for office. Anticipating this, a sequentially rational incumbent will (for high  $p$ ) not play safe. The problem for the voter is that she cannot commit to ignore information about the opponent's ability.<sup>32</sup>

With informative campaigns, reform is then either too high or too low relative to an efficient benchmark. Moreover, a qualitative welfare comparison between a world in which the opponent can influence voters' beliefs about her competence and one where she cannot reveals that, in the former case, inefficiencies are greater. This implies that strategic interaction between an incumbent and opponent, in an environment where the voter can learn about the latter, exacerbates the implementation of inefficient policies.

While the result can plausibly explain real-world examples, it does not imply that informative campaigns are bad. Voters benefit from the information provided by the opponent as they are able to make a more informed decision. Moreover, even unwanted reforms (such as in our

<sup>32</sup> A plausible real-world example of our equilibrium is the 2010 British general election discussed earlier. As we have seen, live televised debates, used for the first time in April 2010, allowed the British electorate to update their beliefs about the executive ability of the prime minister's opponents. In our model, the anticipation that voters do so can have an (endogenous) impact on policy choice resulting in inefficiencies. Consider the last act of the outgoing Parliament, which banned the drug mephedrone and imposed severe penalties on its sale. It was passed on 7 April 2010, and came into force just one day after the first televised leader debate (on 15 April). The ban was imposed following the death of two teenagers on 15 March 2010 who, according to police reports, had taken the drug. The ban led to the resignation of Eric Carlin, a senior member of the UK's Advisory Council on the Misuse of Drugs (ACMD) – an independent panel of experts that advises the government on drug-related issues. A scathing editorial in the medical journal *The Lancet* noted several key concerns and 'documented the very scanty evidence on mephedrone, including the absence of a direct causal link between the reported deaths and the drug'. The editorial went on to say that 'the events surrounding the ACMD signal a disappointing finale to the government's relationship with science. Politics has been allowed to contaminate scientific processes and the advice that underpins policy'. (Toxicology tests later revealed there were no traces of mephedrone in the blood of the two deceased teenagers.) The example does not fit the exact timing of our model, since the outcomes of the ban were not realized prior to the election. Moreover, we cannot say, on the basis of our model, that the ban would not have been imposed without informative campaigns. Nevertheless, the example illustrates how inefficient policy choices are made more likely with an impending and informative opposition campaign.

bridge-building parable) yield social returns that are due to learning. This suggests that a fuller welfare comparison is in order. Thus in the following result, we explore the overall welfare impact of introducing informative campaigns.

PROPOSITION 4: Starting from a position with no informative campaigns ( $p^h = p$ ), their introduction can decrease voter welfare. The welfare loss from an informative campaign is monotonically decreasing in  $p^h$ .

As we have seen, even when  $p^h$  is small, the information available to voters from an opponent's campaign can affect the incumbent's strategic choice of policy. Specifically, he may then be more willing to implement the risky policy. Thus even small amounts of information can exacerbate 'gambling on success'. Of course, the learning value of such information is also small. It follows that the overall welfare effect of informative campaigns can be negative. As  $p^h$  increases, the fact that the voter is able to make a more informed choice when campaigns are informative compensates for these unintended policy consequences. So if campaigns are sufficiently informative, their impact on welfare is positive.

In sum, we have seen that information about an opponent's executive competence increases the incumbent's incentive to choose risky policies and can exacerbate 'gambling on success'. In equilibrium, and as in the case without informative campaigns, a voter will perceive reform to be either too high or too low. Moreover, the introduction of informative campaigns that allow voters to infer, albeit very imperfectly, an opponent's competence can have negative consequences. This is due to the indirect effect of such campaigns on the incumbent's policy choice.

#### ROBUSTNESS CHECKS

In this section we consider the robustness of our key insight that informative campaigns can affect the policy choices of incumbents, reinforcing electoral incentives to under and overinvest in risky policies.

#### *Endogenous Choice of $x$*

A possible criticism of our model is that the equilibria we study may be suboptimal. Until now we fixed  $x$ , the retention probability when the voter is indifferent. Endogenizing her choice of  $x$  allows us to instead study the best possible equilibrium from the voter's perspective. We find that, even in this (somewhat optimistic) optimal world, our core insights remain.

PROPOSITION 5: When campaigns are uninformative, the efficient level of reform is attained under an optimal retention rule. By contrast, when campaigns are informative the voter can prevent 'fear of failure' but is unable to prevent 'gambling on success'.

The first part of the result shows that in the absence of information about the opponent, the voter can use her retention rule to perfectly align the incumbent's incentives with her own. The intuition for this result is straightforward. When there is no information about the opponent and the incumbent's risk taking is excessive relative to the efficient benchmark, the voter would prefer that he plays safe and so should reward him when he does. Of course, since the voter is indifferent after the incumbent plays safe she can choose any  $x$ . Specifically, choosing  $x = 1$  eliminates the incumbent's incentive to engage in excessive first-period reform. Following the same logic, whenever the incumbent has an incentive to shy away from implementing reform (because of her fear of failure) the optimal rule entails  $x = 0$ . Then, the incumbent cannot be re-elected *unless* he implements reform. We see that an optimal retrospective voting strategy attains the efficient outcome so long as campaigns are uninformative.

The features of the optimal rule are worth noting. The best rule does not treat the incumbent and opponent equally, although (from the voter's perspective) they are identical with respect to their expected competence. In fact, the optimal rule discriminates either in favour of or against the incumbent, depending on the circumstances.

The second part of the proposition states that efficiency is not attained even under an optimal rule when campaigns are informative. In fact, whereas the level of risk is never too low (there is no fear of failure) it may be too high. Fear of failure is eradicated, as (when  $p < 1/2$ ) the optimal rule has the incumbent re-elected only when he takes the risky option and is successful. However, although in the absence of an informative campaign the optimal retention rule eradicates gambling on success, this is no longer the case when campaigns are informative. Since it is never rational to retain an incumbent who plays safe when an opponent runs a successful campaign, anticipating this, an incumbent will not do so. In sum, even in the best possible equilibrium (from the voter's perspective) the incumbent will overinvest in reform (gamble on success).

### *Analysis for $x = 0$ and $x = 1$*

As a further robustness check, we explore whether our results hold in extreme cases when  $x = 1$  and  $x = 0$ . In the former case, the voter has an extreme bias in favour of the incumbent so that she always re-elects him when indifferent between the candidates. In the latter case, the voter has an extreme bias so that she never re-elects him under the same set of circumstances. The effects are illustrated in the left-hand panel of Figure 4.

In the absence of informative campaigns, an incumbent who is advantaged ( $x = 1$ ) has no incentive to implement reform. Unsurprisingly, reform is always too low in this case. When campaigns are informative, the opponent will always choose a risky campaign in order to influence voter beliefs. Anticipating this, and following our earlier logic, the opponent is more likely to choose risky than in the absence of such information and, as shown in the left-hand panel of Figure 4, risk taking may exceed the desired level. The result starkly illustrates the impact of informative campaigns on the incumbent's incentives.

In the case when  $x = 0$ , the incumbent can only be elected when taking the risky option. He will do so irrespective of how informative the opponent's campaign is. The result is illustrated in the right-hand panel.

In sum, only in the limited case ( $x = 0$ ) do we find that the introduction of informative campaigns does not change the incumbent's incentives.<sup>33</sup>

### *The Incumbent Can Campaign*

We have supposed that the voter can only learn about the incumbent via his track record. How, if at all, do our results change when allowing for an informative campaign by a political opponent? We can show that our core results hold. While suppressing the main details, we provide a brief overview.

<sup>33</sup> A related concern is that our results may hinge upon a non-robust indifference from the voter's perspective with respect to the second-period policies. This core feature of the Baro (1973) and Ferejohn (1986) models, in which the voter is strictly indifferent between her actions under all circumstances, has been subject to criticism. Indeed, both Fearon (1999) and Besley (2006) have shown that, in those models, a small difference in voter pay-offs will lead the incumbent's incentive scheme to unravel. While indifference also arises here, albeit only under some circumstances, our results do not hinge upon this. Straightforwardly, a small change to the primitives of our model involving a higher (lower) pay-off to the voter when the incumbent (opponent) implements the safe policy in the second period has the same first-period policy consequences as an incumbency bias ( $x = 0.1$ ) as studied here.




Fig. 4. Informative campaigns in the presence of incumbency (dis)advantage  
 Note: here  $p^H = p^{0.3}$  and  $p^h = p^{0.7}$ . Reform is either too low or too high when there is an incumbency advantage ( $x = 1$ ). When there is an incumbency disadvantage ( $x = 0$ ), reform is always too high irrespective of whether the opponent’s campaign is informative.

It is immediately apparent that our central findings survive when the incumbent can engage in campaigning that is costless to the voter for some range of parameters. In particular, if the incumbent’s campaign is less informative than that of the opponent, as appears to have been the case in our motivating example depicted in Figure 1, then for high  $p$  the incumbent will overinvest in reform. The assumption is not unrealistic. Since the incumbent has had a term in office, one might assume that the voter will discount her campaign promises in greater measure than she would those of his opponent.<sup>34</sup>

Going further, we might suppose that the incumbent has access to the same campaign technology as his opponent and that, in line with our key assumptions, his campaign influences voter beliefs only when he has not implemented the risky policy in the first period (that is, when no information about his competence has been revealed). A conjecture is that this would reduce over-experimentation, and indeed straightforward analysis reveals that this is indeed so. Nevertheless, we find that the key qualitative predictions of our analysis survive. Although there is less reform in the first period when the incumbent can campaign after implementing the safe policy, we still observe both under and overinvestment for low and high values of  $p$ , respectively. These results are easily proved when considering the simultaneous choice of campaign strategy, but extend to a world with different sequencing.<sup>35</sup>

CONCLUSION

We studied an environment in which a voter influences an incumbent’s policy choice via her retention rule. In our multi-arm bandit setting, certain policies (which we refer to as either risky or reform) allow voters to learn about an incumbent’s executive abilities. The novelty in our

<sup>34</sup> Relatedly, as noted earlier, there is a large body of empirical evidence starting with Jacobson (1978) that shows incumbent campaign spending is less effective than opponent spending.

<sup>35</sup> Details available upon request. Further robustness checks with respect to changes in our model primitives are discussed in Appendix A.



setting involves information that stems from an opponent's campaign that is correlated (albeit, relatively weakly) with his own executive ability. A key result is that such campaigns exacerbate the inefficient implementation of reforms by an incumbent who thereby showcases her talent. Our result contrasts with expectations when agents take risks and are compensated with a salary that increases linearly with the agent's reputation. Then, as shown by Holstrom, an agent who does not know his own ability underinvests in risky projects.<sup>36</sup> The contrast is due to the coarseness of the agent's reward structure. In our environment he is either retained in office or replaced. Our key results show that inefficiencies stemming from the coarse retention rule are exacerbated by the role of an opponent. Our central lesson is that incumbents place too much emphasis on what is learned through implementing policy and not enough on the voter's pay-off from such policies. A change in the competitive environment that informs voters of an opponent's abilities (such as via the introduction of televised debates) can worsen this situation and the overall welfare effect can be negative. Our model thus illustrates a trade-off between incentives and learning: electoral competition is desirable from a learning perspective, but weakens voter control over the incumbent.

## REFERENCES

- Adams, James, and Zeynep Somer-Topcu. 2009. Do Parties Adjust Their Policies in Response to Rival Parties' Policy Shifts? Spatial Theory and the Dynamics of Party Competition in Twenty-Five Postwar Democracies. *British Journal of Political Science* 39:825–46.
- Aghion, Philippe, Patrick Bolton, Cristopher Harris, and Bruno Julien. 1991. Optimal Learning by Experimentation. *Review of Economic Studies* 58 (4):621–54.
- Ashworth, Scott. 2012. Electoral Accountability: Recent Theoretical and Empirical Work. *Annual Review of Political Science* 15:183–201.
- Ashworth, Scott, and Kenneth W. Shotts. 2010. Does Informative Media Commentary Reduce Politicians' Incentives to Pander? *Journal of Public Economics* 94 (11–12):838–47.
- . 2011. Challengers, Democratic Contestation, and Electoral Accountability, working paper.
- Banks, Jeffrey S., and Rangarajan K. Sundaram. 1990. Incumbents, Challengers, and Bandits: Bayesian Learning in a Dynamic Choice Model. *Rochester Center for Economic Research working paper*.
- . 1993. Chapter 12: Adverse Selection and Moral Hazard in a Repeated Elections Model. In *Political Economy: Institutions, Competition, and Representation*, edited by William A. Barnett, Melvin J. Hinich and Norman J. Schofield, 295–312. Cambridge: Cambridge University Press.
- . 1998. Optimal Retention in Agency Problems. *Journal of Economic Theory* 82 (2):293–323.
- Barro, Robert. 1973. The Control of Politicians: An Economic Model. *Public Choice* 14 (1):19–42.
- Bawn, Kathleen, and Zeynep Somer-Topcu. 2012. Government Versus Opposition at the Polls: How Governing Status Affects the Impact of Policy Positions. *American Journal of Political Science* 56 (2):433–46.
- Bektimirov, Baur, and Pablo B. Montagnes. 2014. Political Incentives to Privatize, working paper.
- Besley, Tim J. 2006. *Principled Agents? The Political Economy of Good Government*. Oxford: Oxford University Press.
- Callander, Steven. 2008. A Theory of Policy Expertise. *Quarterly Journal of Political Science* 31 (2):123–40.
- . 2011a. Searching and Learning by Trial and Error. *American Economic Review* 101 (3):2277–308.
- . 2011b. Searching for Good Policies. *American Political Science Review* 105 (4):643–63.
- Canes-Wrone, Brandice, Michael C. Herron, and Kenneth W. Shotts. 2001. Leadership and Pandering: A Theory of Executive Policymaking. *American Journal of Political Science* 45 (3):332–50.
- Daley, Brendan, and Erik Snowberg. 2011. Even If Its Not Bribery: The Case for Campaign Finance Reform. *Journal of Law, Economics, and Organization* 37 (2):324–49.

<sup>36</sup> Holstrom 1999.

- Dewan, Torun, and Kenneth A. Shepsle. 2011. Political Economy Models of Elections. *Annual Review of Political Science* 14:311–30.
- Fearon, James D. 1999. Electoral Accountability and Control of Politicians: Selecting Good Types Versus Sanctioning Poor Performance. In *Democracy, Accountability, and Representation*, edited by Adam Przeworski, Susan C. Stokes and Bernard Manin, 55–97. Cambridge: Cambridge University Press.
- Ferejohn, John. 1986. Incumbent Performance and Electoral Control. *Public Choice* 50 (1):5–25.
- Fox, Justin. 2007. Government Transparency and Lawmaking. *Public Choice* 131:23–44.
- Fox, Justin, and M. Stephenson. 2011. Judicial Review as a Response to Political Posturing. *American Political Science Review* 105:397–414.
- Gul, Faruk, and Wolfgang Pessenhofer. 2014. The War of Information. *Review of Economic Studies* 79:707–34.
- Hirsch, Alexander V. 2011. Experimentation and Persuasion in Political Organizations. Paper Presented to the Annual Meetings of the American Political Science Association, 31 August 2011.
- Holstrom, Ben. 1999. Managerial Incentive Problems: A Dynamic Perspective. *Review of Economic Studies* 66 (1):169–82.
- Jacobson, Gary C. 1978. The Effects of Campaign Spending in Congressional Elections. *American Political Science Review* 72:469–91.
- Lizzeri, Alessandro, and Nicola Persico. 2009. Electoral Incentives, Political Risk-Taking and Policy Reform. In *The Political Economy of Democracy*, edited by Enriqueta Aragones, Carme Bevia, Humberto Llavador and Norman Schofield, 193–218. Fundación BBVA.
- Maskin, Eric, and Jean Tirole. 2004. The Politician and the Judge. *American Economic Review* 94: 1034–54.
- Prato, Carlo, and Stephane Wolton. 2015. The Voters Curses: Why We Need Goldilocks Voters. *American Journal of Political Science* 60 (3):726–37.
- Strulovici, Bruno. 2010. Learning While Voting: Determinants of Collective Experimentation. *Econometrica* 78:933–71.
- Strumpf, Koleman S. 2002. Does Government Decentralization Increase Policy Innovation. *Journal of Public Economics* 4:207–41.
- Voigtlander, Nico, and Hans-Joachim Voth. 2014. Highway to Hitler. Working paper. Cambridge, MA: NBER.
- Volden, Craig, Michael M. Ting, and Daniel P. Carpenter. 2008. A Formal Model of Learning and Policy Diffusion. *American Political Science Review* 102 (3):319–32.