ECONOMICS

Working Paper Series

No. 10-2026

Political Economy of National Security under the threat of Terrorism

Sambit Bhattacharyya

University of Sussex

s.bhattacharyya@sussex.ac.uk

Stephen Calabrese

University of Sussex

S.M.Calabrese@sussex.ac.uk

Adnan Fakir

University of Sussex

A.Fakir@sussex.ac.uk

Simon Lodato

Middlesex University

s.lodato@mdx.ac.uk

Abstract: Protecting citizens from terrorism often requires opaque security institutions that entail higher agency costs and divert resources from other public goods. Yet politicians face strong electoral incentives to satisfy voters' demand for security. We study this trade-off within a political economy model with politicians, voters, and an unelected security agency featuring endogenous terrorist attacks, two elections, credibility-as-commitment, and bureaucratic inefficiency. Using a calibrated version of this model, we show that electoral incentives lead politicians to overpromise and overprovide security when voter demand for the same is high, but underpromise and underprovide when demand is low. These distortions disappear when security provision is as efficient as other public goods or when bureaucratic inefficiency in security becomes sufficiently severe, as in a failed state. We also present cross-country evidence on security spending, political rhetoric, and institutional quality consistent with the mechanisms highlighted by the model.

JEL codes: D74; C33; C72

Key words: Political Economy; National Security; Terrorism

Political Economy of National Security under the threat of Terrorism[†]

Sambit Bhattacharyya, Stephen Calabrese, Adnan M. S. Fakir and Simón Lodato[‡]

September 21, 2026

Abstract

Protecting citizens from terrorism often requires opaque security institutions that entail higher agency costs and divert resources from other public goods. Yet politicians face strong electoral incentives to satisfy voters' demand for security. We study this trade-off within a political economy model with politicians, voters, and an unelected security agency featuring endogenous terrorist attacks, two elections, credibility-as-commitment, and bureaucratic inefficiency. Using a calibrated version of this model, we show that electoral incentives lead politicians to overpromise and overprovide security when voter demand for the same is high, but underpromise and underprovide when demand is low. These distortions disappear when security provision is as efficient as other public goods or when bureaucratic inefficiency in security becomes sufficiently severe, as in a failed state. We also present cross-country evidence on security spending, political rhetoric, and institutional quality consistent with the mechanisms highlighted by the model.

JEL Classification: D74, C33, C72

Keywords: Political Economy; National Security; Terrorism

[†]We gratefully acknowledge comments by and discussions with the seminar participants at the University of Sussex Economic Theory Group. All remaining errors are our own. **Data Availability Statement:** Data used in this article are from publicly available sources. Replication codes are available upon request. **Funding and Conflict of Interest Statement:** The authors received no external funding for this paper. The authors do not have any conflicts of interest to disclose.

[‡]Bhattacharyya: Department of Economics, University of Sussex, E-mail: S.Bhattacharyya@sussex.ac.uk. Calabrese: Department of Economics, University of Sussex, E-mail: S.M.Calabrese@sussex.ac.uk. Fakir: Department of Economics, University of Sussex, E-mail: A.Fakir@sussex.ac.uk. Lodato: Department of Accounting, Finance and Economics, Middlesex University, E-mail: S.Lodato@mdx.ac.uk

1 Introduction

Political discourse in established democracies often associate the risk of terrorism with national security. Such discourse is not new and goes back to at least President George W. Bush Administration's war on terror. More recently President Donald J. Trump's Presidential Memorandum on 'Countering Domestic Terrorism and Organized Political Violence'[1] is also a case in point. Examples of similar discussions can also be found in the political landscape of Australia, Britain, France, Germany and elsewhere. Governments in established democracies facing terrorist attacks often respond to these threats by increasing national security budget. Politicians in these countries have strong electoral incentives to appear tough on terrorism and to foreground national security. Yet the effectiveness of greater spending ultimately depends on the security agencies responsible for translating these resources into security provision. These agencies prefer larger budgets because they provide greater scope for bureaucratic slack,[2] part of which they can appropriate. Voters dislike bureaucratic slack, even when they are unaware of its extent, because they ultimately bear its costs through the crowding out of other public services. As a consequence politicians also want them to be minimized but not to the extent that it significantly limits the ability of the security agency in delivering its core mission on counterterrorism. Despite the significance of these issues and the interconnectedness of incentives, institutions, and risks, formal theoretical models exploring strategic interactions and mechanisms are rare.

In this paper, we explore inefficiencies in the provision of security and counterterrorism within a theoretical model that captures the strategic interaction between voters, politicians, and the security agency. The model features two elections, which give politicians incentives to keep their promises while, at the same time, they must respond to the need for action generated by the electoral consequences of endogenous terrorist attacks. To compete in the first election, political parties choose their platforms given voters' initial preferences over security and other public goods. Parties can perfectly pander to voters before the election, but doing so may be costly. If elected, the winning party must allocate resources to security provision, which is subject to bureaucratic slack. When deciding whether to re-elect the incumbent, voters assess its credibility by comparing the level of security promised during the campaign with the level ultimately implemented.

The incumbent's problem is further complicated by the possibility of a terrorist attack. The probability of an attack depends on security expenditure but also contains a stochastic

[1] See The White House (2025) for full text.

[2] Security agencies are particularly vulnerable to bureaucratic inefficiency, as the secrecy inherent in national security provision makes agency performance difficult to monitor, creating scope for bureaucratic slack. This problem is compounded in the security and counterterrorism sphere due to the weak relationship between observable agency activities and ultimate outcomes (Cayford and Pieters, 2020; Wilson, 1989). For inefficiencies in the provision of Defense, see Hartley (2011)

component, so the incumbent has only partial control over the realized threat. After taking office, the incumbent chooses security spending, which is implemented by the security agency subject to bureaucratic slack. Before the second election, voters observe whether a terrorist attack has occurred and may update their preferences accordingly. Following an attack, voters place greater weight on national security relative to civil liberties[3] and other public goods. Electoral responses and party platforms are therefore jointly shaped by the institutional environment, including political credibility and incumbency advantage.

We distinguish among three solutions for the provision of security and counterterrorism. First, we have the political solution given by the Nash Equilibrium of the political game where parties compete in the two elections. The political solution is compared to a Central Planner solution, where the Central Planner maximizes voter utility, albeit still subject to the security agency's slack as well as a potential terrorist attack, which renders this a second best. The central planner has the advantage of not facing any credibility constraints or threat from any challenger since both arise purely from electoral battles. Therefore the absence of political competition allows the Central Planner to fully internalize changes in citizens' preferences for security should a terrorist attack occur. Finally, the combination of this last feature and the absence of slack yields the first-best solution.

The model has multiple non-linear components that prevent us from deriving closed-form solutions. Our strategy is then to obtain numerical solutions and assess their welfare implications. For the comparative statics, we take as our baseline a calibration of the model parameters using country-level international data. Furthermore, we provide empirical support for the main assumptions and mechanisms of the model using strong association between approximate variables in both cross-national and national datasets.

The analysis delivers four main results. First, under security stability, characterized by low perceived terrorism risk and a strong preference for civil liberties, political parties represented by Downsian politicians underpropose security spending in Nash equilibrium relative to the level that is implemented. On the other hand, when voters' preferences for security are high, parties propose more than the amount implemented. These outcomes arise because parties pander to the preferences of a myopic voter in order to win the first election.

Second, relative to the Central Planner, parties underprovide security when voters' perceptions of terrorism risk are low but overprovide it when perceived risk is high. In this sense, political competition amplifies the threat cycle. The second result reflects the interaction between bureaucratic inefficiency and politicians' credibility concerns. Once elected, politicians face two opposing incentives. Technical inefficiency pushes them to

[3]The trade-off between security and civil liberties is central to the political economy of security and counterterrorism, typically under the label of "The Dilemma of Liberal Democracies," (see, for instance, Enders and Sandler, 2011.)

reduce security expenditure and reallocate resources toward other public goods, whereas credibility concerns push them toward the implementation of the security expenditure promised before the election. When voters perceive terrorism risk to be low, parties can win the first election by promising relatively little security. Because the promised level is low, deviations from that promise induced by bureaucratic slack generate only modest credibility costs. The incentive to limit inefficient security expenditure therefore dominates, leading politicians to underprovide security relative to the Central Planner. By contrast, when perceived terrorism risk is high, parties must promise substantially more security to win the first election. In this case, since the amount of slack is positively correlated with security expenditure, deviations from the promised level caused by bureaucratic slack entail larger credibility losses. Credibility concerns consequently become stronger, inducing parties to promise and allocate more security than under the Central Planner solution.

The third result shows that when the state security provision is perfectly efficient, characterized by no bureaucratic waste, promised and implemented security expenditures in the political solution converge to the Central Planner's allocation as voters' perceptions of insecurity become sufficiently high. In this case, the distortions generated by bureaucratic slack disappear, eliminating the wedge between the political and planner allocations and eventually eroding any risk of a terrorist attack. The flip side is our fourth result: when bureaucratic waste becomes sufficiently severe, as in a failed state, security expenditure under both the Central Planner and the political equilibrium converges to the same low level, rendering a terrorist attack certain. At such high levels of inefficiency, limiting the resources lost to bureaucratic waste becomes the overriding consideration for both politicians and the Central Planner. Taken together, these two results highlight the crucial role played by the inefficiency characteristic of security and counterterrorism provision.

The paper makes the following contributions. Game theoretic treatment of the relationship between voters, politicians, and the security bureaucracy under the threat of terrorism and institutional constraints such as credibility is relatively new. This paper contributes to this literature which has largely evolved from the study of 9/11 terrorist attack. Even though this literature offers a number of empirical insights through cross national econometric modeling, the study of strategic behavior within government as well as its interaction with the voter is to the best of our knowledge new.

There are however notable exceptions. Our paper is close to the recent work by Di Lonardo and Dragu (2025), who show that electoral incentives induce an overprovision of counterterrorism intervention and that greater electoral salience magnifies this distortion. Their mechanism operates through voters inferring incumbent competence from the realization or not of terrorist attacks, giving politicians an additional electoral incentive to prevent attacks. Their bureaucratic extension separates the politician's choice of counterterrorism powers from an agency's choice of effort, with greater powers lowering the marginal cost of effort. We instead model a slack-maximizing security agency whose inefficient implementation makes fulfilling electoral promises costly. Combined with credibility concerns and attack-induced changes in voter preferences, our mechanism yields a fundamentally different result: electoral competition can produce both too little and too much security. Whereas in Di Lonardo and Dragu (2025) electoral distortion is consistently toward excessive intervention, ours reverses direction as initial security demand rises: from underprovision to overprovision relative to a planner facing the same bureaucratic constraints. Electoral incentives therefore amplify both insufficient provision in low-demand environments and excessive provision in high-demand environments. We then explain this reversal, through the interaction of electoral commitments, which is in line with the descriptive evidence of political rhetoric we provide, and bureaucratic inefficiency.

Less recently, Bueno de Mesquita (2007) shows how the opacity of counterterrorism spending creates agency problems between voters and government, inducing an inefficient composition of expenditure that favors observable measures over covert operations. In our model, the opacity of bureaucratic inefficiency to voters contributes to the divergence between voter preferences and both proposed and actual security expenditures.

Using simulations and comparative statics, we are also able to investigate welfare implications within different institutional settings as well as security threats. Finally, we complement the theoretical analysis with descriptive empirical patterns that motivate selected assumptions and mechanisms in the model.

Literature on the economics of national security and terrorism originates from the seminal work of Landes (1978) studying US aircraft hijackings. More recently, Abadie (2006) presents an econometric study of the relationship between civil liberties, political rights and terrorism. In a similar vein, a large empirical literature focuses on the interlinkages between economic freedom, civil liberties, and economic growth. Helliwell (1994), Heckelman (2000), and Harms and Ursprung (2002) are notable examples. Broad focus of the above-mentioned literature has been four-fold. They are: (1) macroeconomic effects of terrorism and national security, especially on economic growth, (2) causal determinants of terrorism, (3) dilemma of liberal democracies in balancing the tradeoff associated with counter-terrorism, and (4) effectiveness of counter-terrorism policies. Sandler (2014) presents an excellent summary of this literature, and Daxecker and Hess (2013) and Daxecker (2017) present a summary of the literature on the economics of counter-terrorism. Needless to say, our study relates to this literature. In particular, it explores strategic interaction between the three key agents in counter-terrorism policies in established democracies.

On the theory side, our paper builds on the probabilistic voting model (Persson and Tabellini, 2002) and separates government decision-making into political and bureaucratic components, with the latter in the hands of the security agency. In doing so, we merge the ideas of a slack maximizing bureaucracy (Mueller, 2003; Wyckoff, 1990) into a probabilistic voting model. This framework also allows us to study political accountability (Adserà et al., 2003) and inherent inefficiency in government (Brownfield, 1977).

Our paper also relates to the economics of conflict literature to the extent that it studies strategic interaction between agents, even though the nature of interactions is starkly different in the case of terrorism. Esteban and Ray (2008) is a notable theoretical contribution in this literature that models conflict as an outcome of strategic interactions between agents. Esteban et al. (2012) presents empirical applications of their theoretical result. Numerous other empirical studies analyze the causes and consequences of conflict. Berman et al. (2017) and Bhattacharyya and Mamo (2021) are two recent empirical studies that focus on the local determinants of conflict. Notable earlier empirical studies focusing on cross-national variation are Collier and Hoeffler (1998), Collier and Hoeffler (2004), and Miguel et al. (2004). Blattman and Miguel (2010), Anderton and Carter (2019), and Anderton and Brauer (2021) present excellent reviews of this literature.

The remainder of the paper is organized as follows. Section 2 presents the theoretical model, together with the associated propositions and corresponding proofs. It also provides empirical evidence supporting the assumption that preferences over security provision respond to terrorist attacks. Section 3 calibrates the model using real-world data. Using this calibration as the benchmark equilibrium, Section 4 presents the comparative statics and welfare analysis. To motivate modeling assumptions and validate benchmark and comparative statics calibrations, we use observational data from the Global Terrorism Database (GTD), the Stockholm International Peace Research Institute (SIPRI) database, and the Manifesto Project Database (MPD). Using these databases, we are able to observe relevant variables of interest across country-year samples over 1990–2020, with coverage varying by source. Section 5 concludes. Some proofs, theoretical analysis, and additional empirical results are provided in the Appendix.

2 The Model

We study the political economy of security and counter-terrorism within a model with three classes of agents: political parties D and R, a single voter representing the whole electorate, and an unelected Security Agency, the SA. The political parties propose and allocate (as an incumbent) a proportion of the government budget on security and anti-terrorism, and are exclusively focused on electoral victories in two sequential elections.

The voter, motivated by their own welfare, elects one party in each election. The *SA* decides how much of a given budget allocation to capture through slack.

The model is rich in that it encompasses the realization of an endogenously random[4] terrorist attack, two elections to incorporate credibility-as-commitment technology, and institutional and bureaucratic inefficiency. Within this framework, we explore questions such as the welfare implications of electoral incentives, including on the probability of a terrorist attack, under different levels of voter initial preference for security and different levels of relative inefficiency in the provision of security compared to other government provisions. The analysis is particularly relevant because the secrecy involved in some security and counter-terrorism activities, such as intelligence or defense procurement, may weaken accountability and monitoring, thus increasing the scope for inefficiency relative to other government activities.

In what follows, we first present a general description of the model through the timeline of events in the political game on Security Provision, then elaborate further on the behavior of the security agency, the voter's preferences, and finally derive the utility functions of the candidates.[5]

The first event in the political game, shown in Figure 1, is that the parties D and R choose their platforms to compete in the first election. The platforms are in terms of the proportion of the government budget parties propose to spend on Security and Anti-terrorism, namely $s_1^D, s_1^R \in [0,1]$. Since there is no prior history ($\emptyset$), the assumed backward-looking (or experience-based) voter takes at face value the future implementation of the proposed policies. Then Election 1 takes place, where the voter elects either D or R. Once in office, the winner of Election 1, the incumbent $I \in \{D, R\}$, appropriates $s_2^I \in [0,1]$, which is the *upon-approval* allocation to the Security Agency *SA*, which then implements

$$s_{\text{SA}}(s_2^I) = s_2^I\Big(1 - \frac{\omega s_2^I}{2}\Big). \tag{1}$$

In the following section, we derive and formalize the connection between the upon-approval allocation, s_2^I, and the (lower) security expenditure implemented by the SA, $s_{\text{SA}}(s_2^I)$, that results in equation 1. For now, we define ω broadly as the relative level of bureaucratic inefficiency in the provision of security compared to the overall average inefficiency level of other government activities.[6]

After $s_{\text{SA}}(s_2^I)$ is implemented, there may (T) or may not be a terrorist attack (NT), which could change the voter preferences, as we detail in Section 2.2 below. The probability

[4]By endogenously random, we mean that there is a random (stochastic) component, but the probabilities themselves depend on the policy.

[5]Appendix A6 provides a table of model notation.

[6]As a complementary assumption, we assume that inefficiency in the provision of public goods other than security is orthogonal to policy and normalize it to zero.

of a "No Terrorist Attack" (NT) is positively correlated with the security provision implemented by the SA. The challenger observes whether T or NT is realized and can adjust its platform to compete in the second election according to the updated preferences of the voter. The next event on Figure 1 is then the challenger's proposal s_2^{-I}.

The incumbent, by contrast, cannot respond to T or NT, because its platform for the second election is tied to the policy already implemented. The voter, unaware of the bureaucratic role in the policy, assesses the incumbent's performance comparing $s_{\text{SA}}(s_2^I)$ against the promised s_1^I. This comparison determines the perceived credibility of the incumbent and influences the voter's choice in Election 2. Then Election 2 is held, and the game ends.

- $\emptyset$
- Policy platforms s_1^D s_1^R
- Election 1: the voter elects either D or R to be the Incumbent I
- Incumbent $I \in \{D, R\}$ appropriates s_2^I
- The SA implements $s_{\text{SA}}(s_2^I)$
- Realization of the Terrorist Attack state: T or NT
- Policy platform challenger, s_2^{-I}, $-I \triangleq \neg I$
- Election 2: The voter either re-elects or replaces I
- End

Figure 1: Timeline of the political provision of Security and anti-terrorism

This general overview conveys the main features of the model. First, there is an agency problem between the policymaker (or politician) and the bureaucratic agency that implements the policy. Central to this problem is the level of inefficiency involved in the provision of security and anti-terrorism. One interpretation is that the secrecy and lack of transparency required in these activities make monitoring more costly, rendering the provision of security relatively more inefficient than the implementation of other public policies. Second, electoral incentives are captured through two elections with a voter who is backward-looking and ignorant of the "cooking" of policymaking, in the sense that they do not directly observe the inefficiency involved in the implementation of security policy by the SA. In particular, because the voter is backward-looking, preferences adjust only after, and if, a terrorist attack (T) is realized. Finally, this feature links the policy proposed in the first election to the actual implementation of the policy, capturing a credibility mechanism whereby voters punish unfulfilled promises or deviations from institutional commitments. Next, we formalize these features to derive the utility functions of the parties.

2.1 SA's Implementation of Security and Counter-terrorism

As illustrated in Figure 1, the *upon-approval* allocation s_2^I is typically larger than the by-SA implemented s_{SA}. We provide below a rationale for this difference by modeling the implemented level of expenditure on security and counter-terrorism as a reduced-form outcome of the interaction between the elected incumbent party (the principal) and the unelected security agency (the agent).

The SA has no motive other than to maximize "slack," that is, the difference between the budget appropriated to the agency and the amount the agency actually spends. The incumbent, in turn, holds some power in the sense that they can, to some extent, withhold part of the promised allocation (and spend it on other items) which incentivize the SA to increase its effort. This last feature finds a real-world counterpart in U.S. institutions. While the Constitution clearly establishes that the president cannot spend money without a congressional allocation/appropriation, there is nonetheless room for the president to refuse to spend funds that have been approved by Congress, a practice known as *rescission*. For example, the Rescissions Act of 2025, passed by the 119th United States Congress, rescinds \$7.9 billion in funding from international assistance programs and \$1.1 billion from the Corporation for Public Broadcasting (CPB).

Our approach, which features a slack-maximizing SA and a government with rescission powers, captures many of the behavioral assumptions and conceptual foundations used in the literature that models the relationship between governing parties and bureaucratic agencies: a slack-driven bureaucracy (Busch, 2002; Park and Jang, 2021; Wyckoff, 1990); the negative relationship between bureaucratic size and efficiency (Adserà et al., 2003; Alesina and Angeletos, 2005; Becker, 1968; Brownfield, 1977; Goel and Nelson, 1998; Kotera et al., 2012; Rose-Ackerman, 1978,9; Treisman, 2000); hierarchical control (Breton and Wintrobe, 1975,8; Li et al., 2023; Miller, 1977); and the idea that additional effort by a bureaucratic agency enhances observable outputs and thereby justifies higher allocation while the marginal political returns to the government diminish as effort increases (Arce and Sandler, 2005; Holmström and Milgrom, 1991; Niskanen, 1971; Sandler and Enders, 2008).

Formally, the SA chooses the level of effort S_e with the objective of maximizing the

Slack function:

$$SA \text{ Slack} = S_I \left[1 - \frac{GB}{\omega\, S_I}\left(1 - \frac{cS_e}{S_I}\right)^2\right] - cS_e, \tag{2}$$

where $\omega \geq 0$ is an exogenous measure of bureaucratic inefficiency (including corruption) in the provision of Security & anti-terrorism relative to the provision of other government goods and services. A higher values of ω imply greater relative inefficiency or corruption of the Security Agency compared to other government agencies; GB denotes the total government budget; c is the cost per unit of effort by the SA; cS_e represents the SA's total implemented expenditure on security and counter-terrorism; and S_I is the amount that the government allocates to the SA. S_I corresponds to (the upon-approval) $s_2^I \times GB$ in Figure 1. The first term of the RHS of equation 2, $S_I\left[1 - \frac{GB}{\omega\, S_I}\left(1 - \frac{cS_e}{S_I}\right)^2\right]$, the actual allocation to or revenue of the SA, incorporates the main features of the relationship between a principal government and a bureaucratic agent. For example, it is increasing and concave in cS_e, which means that, for a given upon-approval S_I, the actual allocation to the SA is increasing in effort but subject to diminishing returns to the SA. It also captures government's power of rescission every time that $S_I > S_I\left[1 - \frac{GB}{\omega\, S_I}\left(1 - \frac{cS_e}{S_I}\right)^2\right]$.

As S_I results from the electoral contest and therefore is exogenous to the SA, and for a given total government budget GB, the SA implemented expenditure, as a fraction of the GB, that maximizes the slack in equation 2 is:

$$s_{\text{SA}}(s_2^I; c, \omega) = \frac{cS_e^*}{GB} = s_2^I\left[1 - \frac{\omega s_2^I}{2}\right] \tag{3}$$

Hereafter, we sometimes denote $s_{\text{SA}}(s_2^I; c, \omega)$ simply by $s_{\text{SA}}(s_2^I)$ to emphasize that this is the implemented policy corresponding to an upon-approval allocation s_2^I. Finally, it is straightforward to derive the SA's slack and the rescission undertaken by the incumbent that result from the reduced-form government–agency relationship, all as a proportion of the (fixed) total budget GB:

$$\text{Rescission} = \text{Slack} = \frac{\omega(s_2^I)^2}{4}. \tag{4}$$

2.2 Parties' Utilities from Electoral Competition

Electoral competition takes place over two elections in which parties D and R care only about winning. The model is probabilistic in that parties face random popularity shocks. In addition, there is an incumbent advantage which can be narrowed by a (lack of) credibility punishment, and also, albeit partly endogenous, a random occurrence of a terrorist attack.

The political parties maximize the expected utilities of the outcomes of the two sequential elections with the pair of (instantaneous) utilities as shown in Table 1. B can be interpreted as the discount factor applied to the utility of the second election, and in that case, typically $B \in (0, 1]$, although this does not necessarily have to be the case.

Election Outcome	**Candidate Utility Pair**
Lose both elections	$(0, 0)$
Win first, lose second	$(1, 0)$
Lose first, win second	$(0, B)$
Win both elections	$(1, B)$

Table 1: Instantaneous utilities of electoral outcomes

Voter Preferences for Security and Counter-terrorism

The electoral fate of parties D and R depends on how attractive and credible their policies appear to the voter, whose preferences over security and the rest of public goods and services (say education or individual liberties) are represented by the utility function

$$U = \mathcal{U}\Big(CL, P(T), \textit{other government activities}\Big) \tag{5}$$

where CL are civil liberties and $P(T)$ the probability of a terrorist attack.

We transform the fundamental inputs of the voter utility function in (5) into terms of government spending allocated to each component, and embed them within the Cobb–Douglas utility function

$$U = (GB - S)^{\beta} S^{\frac{\gamma_j}{\phi}}, \tag{6}$$

where, as before, GB represents the given total government budget, and S denotes expenditures on security and counter-terrorism. The parameter β captures the voter's preference for expenditures other than security and counter-terrorism spending.

We define $\delta_j = \frac{\gamma_j}{\phi}$, with $j \in \{L, H\}$ and $\delta_H > \delta_L$. The components of δ_j are γ_j, the preference parameter for spending that reduces the probability of terrorism, and ϕ, the preference parameter reflecting the negative effects that security spending may have on civil liberties. Hence, through δ_j, the model captures the voter's perceived relationship between government security and counter-terrorism expenditures and the probability of no terrorist attack. An increase in γ_j (and thus δ_j) implies that voters perceive an increase in the risk of a terrorist attack for a given level of S, and therefore prefer higher SA spending, assuming ϕ remains constant. We denote these preferences by H following a terrorist attack, and by L when no attack has occurred.

We model an attack as increasing the weight placed on security from δ_L to δ_H, rather than changing voters' ideological type. Figure 2 reports cross-country associations between terrorism exposure, military spending, and security-oriented political rhetoric.[7] Appendix A4 defines the variables, Appendix A5.2 presents the ideology comparison and rhetoric decomposition, and Appendix A5.1 discusses institutional quality.

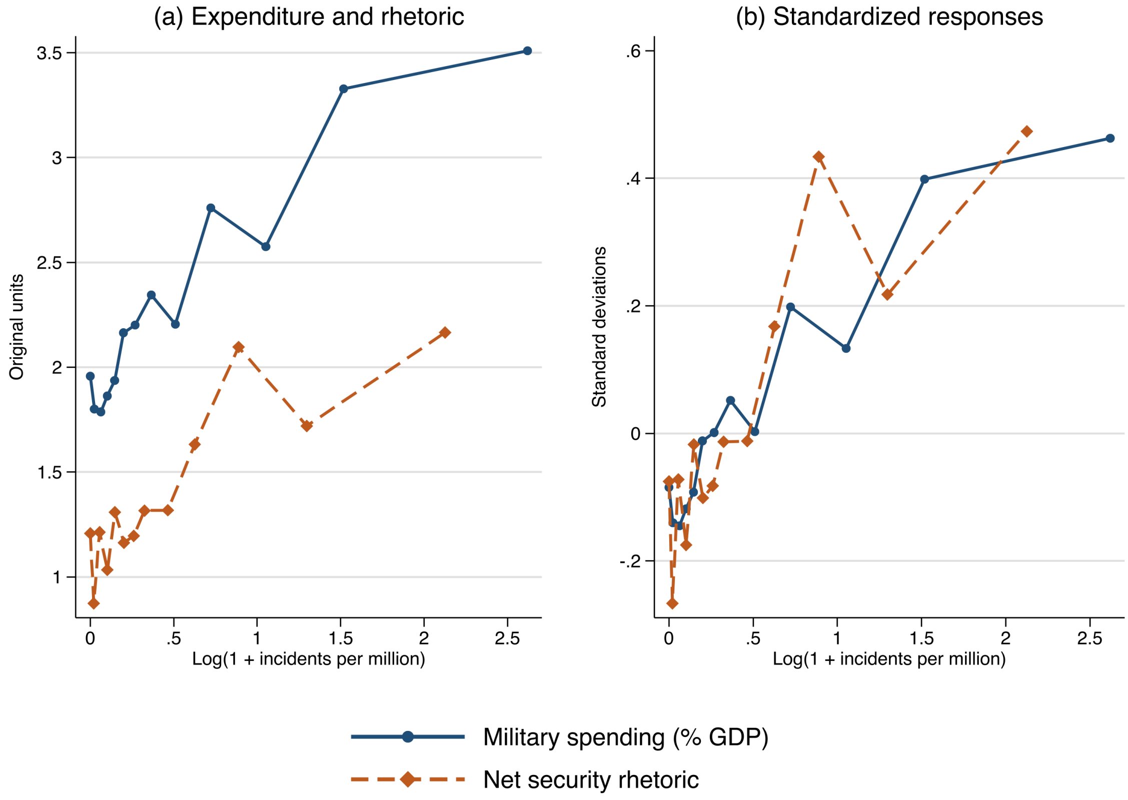

Figure 2: **Terrorism, security spending, and political rhetoric.** Country-year observations from 1990–2020, excluding 1993, are grouped separately for each series into a zero-exposure bin and up to 12 quantile bins of positive terrorism exposure. Each point gives the series-specific mean outcome and mean log(1 + terrorist incidents per million). Military spending uses 4,266 observations from 162 countries/territories; net security rhetoric uses 457 observations from 63. Panel (a) reports original units; panel (b) reports means of outcomes standardized at the country-year level within each series' sample before binning. Sources: GTD, SIPRI and MPD. See Appendix A4 for more details.

As GB is fixed, we want to express the policy as a proportion of GB. For that, we can factorize the utility as $U = GB^{(\beta+\delta_j)}(1-s)^{\beta}s^{\delta_j}$. Furthermore, $GB^{(\beta+\delta_j)}$ can be dropped since it is a constant and, as before, s is the proportion of GB spent on security.

Three main elements in our model render the political contest probabilistic: a random though endogenous terrorist attack, incumbent advantage, and candidate popularity. For the latter, as typical in probabilistic voting models, we introduce a popularity shock χ that favors the party R over D, which follows a uniform distribution with height Ψ and support $\left[-\frac{1}{2\Psi}, \frac{1}{2\Psi}\right]$. Therefore, the voter will support D if $(1-s^D)^{\beta}(s^D)^{\delta_j} > (1-s^R)^{\beta}(s^R)^{\delta_j} + \chi$

[7] These patterns provide descriptive motivation for the preference shift; they do not identify a causal effect or test the mechanism generating the reversal in political distortions.

or will be indifferent when $\chi = (1-s^D)^\beta (s^D)^{\delta_j} - (1-s^R)^\beta (s^R)^{\delta_j} \triangleq \bar{\chi}$.

In a single election, the probability that the voter votes for D and thus D wins the election is,

$$\text{Prob(D wins)} = \text{Prob}(\chi < \bar{\chi}) = \frac{1}{2} + \Psi\left[(1-s^D)^{\beta_j}(s^D)^{\delta_j} - (1-s^R)^{\beta_j}(s^R)^{\delta_j}\right] \tag{7}$$

One minus this probability gives the probability that party R wins the election. As indicated above, a party receives a payoff of one if they win the first election and a payoff of zero if they lose.

Probability of a Terrorist Attack

As suggested by the timeline in Figure 1, the probability that a candidate wins the second election depends on whether a terrorist attack occurs between elections. This is because the occurrence of such an attack strengthens voter preferences for security and counter-terrorism (SA) spending.

In particular, if no terrorist attack occurs, the parameter δ_j equals δ_L, which we assume represents the starting voter preferences in the first election. If, instead, a terrorist attack occurs before the second election, the voters' parameter becomes $\delta_H = \delta_L + \varepsilon$, where $\varepsilon > 0$. The intuition is that, following a terrorist attack, a backward-looking voter perceives a higher-risk environment and therefore place greater weight on security and counter-terrorism spending.

Parties perceive the probability of a terrorist attack as a function of actual counter-terrorism spending, and their belief structure about this relationship is modeled as follows. Assume party D wins the first election, sets $s_2^D(s_1^D)$ and SA implements $s_{\text{SA}}(s_2^D)$. Then, the candidates' perceived relationship between s_2^D and the probability of *no* terrorist attack is:

$$0 \le \text{Prob(NT)} = \nu\, s_2^D \left(1 - \frac{\omega\, s_2^D}{2}\right) \le 1 \quad \text{for all } s_2^D \in [0,1].$$

Where NT stands for No Terrorist Attack. A sufficient bound on ν would be:

$$\nu \;\le\; \frac{1}{\max\limits_{s\in[0,1]} s\left(1 - \dfrac{\omega s}{2}\right)} = \begin{cases} \dfrac{1}{1-\dfrac{\omega}{2}}, & 0 < \omega \le 1, \\ 2\omega, & \omega \ge 1. \end{cases}$$

The probability of a Terrorist Attack (T) is

$$\text{Prob(T)} \;=\; 1 - \text{Prob(NT)} = 1 - \nu\, s_2^D \left(1 - \frac{\omega\, s_2^D}{2}\right).$$

Credibility and Incumbency Advantage

In our model, the winner of the first election, say, D, becomes the incumbent heading into the second election and begins the electoral contest with an exogenous incumbent advantage component, θ, which enters directly as an addend in the incumbent's probability of winning the second election.

However, both parties anticipate that after the incumbent, in this case D, and the SA implement $s_2^D\left(1-\frac{\omega s_2^D}{2}\right)$, voters will observe this policy and compare it to the previously promised s_1^D. We then define a credibility-as-commitment mechanism whereby the voter punishes the incumbent by subtracting from the exogenous component of the incumbent advantage, θ, the term

$$a\left|s_2^D\left(1-\frac{\omega s_2^D}{2}\right)-s_1^D\right|^{expc}, \tag{8}$$

where $expc \geq 1$, $a > 0$, and as $a \to 0$, the connection between elections through credibility fades out.

In Appendix A3, we characterize the properties of the credibility term in expression 8. In particular, for sufficiently large values of ω, it becomes impossible to fully deliver a promised level s_1^D, that is, to choose an appropriation level s_2^D such that the credibility penalty in expression 8 is zero. Moreover, for a fixed promised level s_1^D, the credibility penalty is a convex function of s_2^D and, for sufficiently high values of s_2^D, it increases toward the supremum of the credibility function as s_2^D increases.

In general, credibility limits how far parties can strategically move toward the median voter during elections. While in Alesina (1988) credibility ensures dynamic consistency through forward-looking voters even in a one-election model, here consistency is instead ensured by the presence of two elections.

Two Elections and Parties Expected Utility

We now have all the elements needed to derive the expected utilities of the parties, encompassing two elections with instantaneous utilities as shown in Table 1. We construct the utility functions component by component. The derivation of the parties' utility functions uses the assumption of one (representative) voter and the following result.

Lemma 1. *Let $C \in \{D, R\}$ be a challenger in the second election and $j \in \{L, H\}$. Then, in equilibrium,*

$$s_C^2 = \frac{\delta_j}{\beta + \delta_j}.$$

Proof. Candidate C is not committed to any s_C^2 by credibility. Furthermore, at the moment of choosing s_C^2, candidate i has observed whether or not a terrorist attack has occurred. Therefore, the value of s_C^2 that maximizes i's probability of winning the second election is the one that maximizes voters' instantaneous utility; that is, $s_C^2 = \frac{\delta_j}{\beta+\delta_j}$. □

The probability of winning the first election for party D is

$$P_1(D) = \frac{1}{2} + \Psi\left[\left(1 - s_1^D\right)^\beta \left(s_1^D\right)^{\delta_L} - \left(1 - s_1^R\right)^\beta \left(s_1^R\right)^{\delta_L}\right] \tag{9}$$

The probability that D loses the first election $1 - P_1(D) =$ the probability R wins the first election $= P_1(R)$.

The probability of no terrorist attack if D is incumbent times the probability of D winning the second election if D is incumbent and no terrorist attack equals

$$\begin{aligned} &P_2^{NT}(D; D = \text{incumbent}) = \\ &\nu s_2^D\left[1 - \frac{\omega s_2^D}{2}\right] \times \left(\frac{1}{2} + \Psi\left[\left(1 - s_2^D\left[1 - \frac{\omega s_2^D}{4}\right]\right)^\beta \left(s_2^D\left[1 - \frac{\omega s_2^D}{2}\right]\right)^{\delta_L}\right.\right. \\ &\left.\left. - \left(\frac{\beta}{\beta+\delta_L}\right)^\beta \left(\frac{\delta_L}{\beta+\delta_L}\right)^{\delta_L} + \theta - a\left|s_2^D\left[1 - \frac{\omega s_2^D}{2}\right] - s_1^D\right|^{expc}\right]\right) \end{aligned} \tag{10}$$

The probability of terrorist attack if D is incumbent times the probability of D winning second election if D is the incumbent and the terrorist attack is realized is

$$\begin{aligned} &P_2^{T}(D; D = \text{incumbent}) = \\ &\left(1 - \nu s_2^D\left[1 - \frac{\omega s_2^D}{2}\right]\right) \times \left(\frac{1}{2} + \Psi\left[\left(1 - s_2^D\left[1 - \frac{\omega s_2^D}{4}\right]\right)^\beta \left(s_2^D\left[1 - \frac{\omega s_2^D}{2}\right]\right)^{\delta_H}\right.\right. \\ &\left.\left. - \left(\frac{\beta}{\beta+\delta_H}\right)^\beta \left(\frac{\delta_H}{\beta+\delta_H}\right)^{\delta_H} + \theta - a\left|s_2^D\left[1 - \frac{\omega s_2^D}{2}\right] - s_1^D\right|^{expc}\right]\right) \end{aligned} \tag{11}$$

The probability of no terrorist attack if D is the challenger times the probability D wins the second election if D is the challenger and no terrorist attack is realized is

$$\begin{aligned} &P_2^{NT}(D; D = \text{challenger}) = \\ &\nu s_2^R\left[1 - \frac{\omega s_2^R}{2}\right] \times \left(\frac{1}{2} + \Psi\left[\left(\frac{\beta}{\beta+\delta_L}\right)^\beta \left(\frac{\delta_L}{\beta+\delta_L}\right)^{\delta_L}\right.\right. \\ &\left.\left. - \left(1 - s_2^R\left[1 - \frac{\omega s_2^R}{4}\right]\right)^\beta \left(s_2^R\left[1 - \frac{\omega s_2^R}{2}\right]\right)^{\delta_L} - \theta + a\left|s_2^R\left[1 - \frac{\omega s_2^R}{2}\right] - s_1^R\right|^{expc}\right]\right) \end{aligned}$$

The probability of terrorist attack if D is the challenger times the probability D wins the second election if D is challenger and a terrorist attack is realized is

$$P_2^T(D; D = \text{challenger}) =$$
$$\left(1 - \nu s_2^R\left[1 - \frac{\omega s_2^R}{2}\right]\right) \times \left(\frac{1}{2} + \Psi\left[\left(\frac{\beta}{\beta+\delta_H}\right)^{\beta}\left(\frac{\delta_H}{\beta+\delta_H}\right)^{\delta_H}\right.\right.$$
$$\left.\left. - \left(1 - s_2^R\left[1 - \frac{\omega s_2^R}{4}\right]\right)^{\beta}\left(s_2^R\left[1 - \frac{\omega s_2^R}{2}\right]\right)^{\delta_H} - \theta + a\left|s_2^R\left[1 - \frac{\omega s_2^R}{2}\right] - s_1^R\right|^{expc}\right]\right)$$

Based on the above, below are the utility functions of parties D and R maximize, respectively to solve for the Nash equilibrium

$$\begin{aligned} EU_D = & P_1(D) \\ & + B\Big[P_1(D)(P_2^{NT}(D; D = \text{incumbent}) + P_2^T(D; D = \text{incumbent})) \\ & + P_1(R)(P_2^{NT}(D; D = \text{challenger}) + P_2^T(D; D = \text{challenger}))\Big] \end{aligned} \tag{12}$$

$$\begin{aligned} EU_R = & P_1(R) \\ & + B\Big[P_1(R)(P_2^{NT}(R; R = \text{incumbent}) + P_2^T(R; R = \text{incumbent})) \\ & + P_1(D)(P_2^{NT}(R; R = \text{challenger}) + P_2^T(R; R = \text{challenger}))\Big] \end{aligned} \tag{13}$$

FOC's are symmetrical, hence both candidates converge to the same s_1 before the election and the same s_2 if they win the election in the Nash Equilibrium. The full expression of the expected utility function is provided in equation (A5) in Appendix A1, for party D. Note that immediately after the first election, given s_1^D and s_1^R, and conditional on having won that election, party D chooses $s_2^D(s_1^D)$ to maximize the probability of winning the second election, where s_1^D and s_2^D are related through the credibility component. These two values behave as if they were simultaneously chosen before Election 1. The solution concept is then (symmetric) Nash equilibrium.

3 Benchmark Equilibrium and Parameter Calibration

In this section, we describe and justify the parameter values that characterize our benchmark equilibrium and the resulting endogenous variables, which serves as the baseline for the comparative statics that follow. All parameters and variables appear in the utility function of the parties in equations (12) and (13).

The incumbency advantage has in our setting two parts: an exogenous component θ that increases the re-election probability that the incumbent enjoys purely by virtue of already being in office. Also, an endogenous component that penalizes the incumbent for not implementing the security policy they proposed in their first-election platform. In the utility function (see for example Equation 11), the (lack of) credibility punishment reduces θ by the absolute value difference between the policy implemented by the SA and the period–1 platform, to the power of parameter 'expc' and weighted by the credibility parameter a.

We set a baseline of $\theta = 0.05$, representing an exogenous incumbency advantage that gives a 5-percentage-point boost in the probability of winning the second election for the party holding office. Our calibration follows previous studies: Lee (2008) estimates a U.S. House incumbency advantage of about 6–8 percentage points using a regression-discontinuity design, while Ferreira and Gyourko (2009) find a similar 5 percentage-point effect in U.S. mayoral elections. These magnitudes imply baseline re-election probabilities of roughly 0.45–0.55 for otherwise symmetric candidates. It is worth noting Guriev et al. (2025), which documents a decline in incumbency advantage from about 10 percentage points to just 3—using a regression-discontinuity design—and attributes this fall primarily to the expansion of mobile broadband. Finally, we consider an additional benchmark in which θ takes the value 0.045 after a terrorist attack, and 0.05 if no attack has occurred.

Before each election and after choosing their electoral policy platforms, parties are subject to a popularity shock $\chi_{t,\,t=1,2}$. Both shocks are independently drawn from the same uniform distribution with support $\left(-\frac{1}{2\Psi}, \frac{1}{2\Psi}\right)$. The parameter $\Psi > 0$ governs the sensitivity of electoral outcomes to policy choices, with the effect of policy differentials on the probability of winning an election increasing in Ψ. For example, when $\Psi \to 0^+$, policies have no effect and elections are maximally noisy (maximum entropy). We normalize Ψ to one, which conveniently turns χ's support into $\left(-\frac{1}{2}, \frac{1}{2}\right)$.

Voter's preferences over policies are represented by the Cobb-Douglas utility function (6) with two power parameters: β, the exponent for government expenditure other than the SA's, and $\delta_{j,\,j=L,H}$, for the expenditure on security and counter-terrorism. The values of these two parameters determine the voter's most preferred share of the budget spent on security, given by $\frac{\delta_j}{\delta_j+\beta}$. Importantly, we allow the voter to update their default preference parameter δ_L so that, after a terrorist attack, they become more inclined toward higher spending on security: $\delta_H > \delta_L$. In our benchmark, we set $\beta = 1 - \delta_L$, and the value of δ_L is calibrated from the benchmark equilibrium constraints, as elaborated below. The decomposition of security rhetoric in Appendix A5.2, Figure A7, provides additional descriptive motivation for the shift from δ_L to δ_H.

Still, it is convenient at this point to discuss how the values of δ and β relate to the expenditure shares. Table 2 shows the estimated level of spending on defense as a percentage of both GDP and total government spending in two particular periods. Suppose, for the sake of argument, that these figures represent what the median voter perceives as the optimal share of defense. If we consider only US defense expenditure in the pre-9/11 period, the defense share was about 9%, so we set $\delta_L = 0.09$ and $(1 - \delta_L) = \beta = 0.91$. During the post-9/11 peak the share was approximately 15%, which we impose as

$$\frac{\delta_H}{\delta_H + 0.91} = 0.15,$$

implying $\delta_H \approx 0.16$.

Period	Context	As % of GDP	As % of total spending assump. 38% GDP
1999–2001	Pre-9/11 trough	$\approx 3.5\%$	$\approx 9\,\%$
2010–2011	Post-9/11 peak (Iraq/Afghanistan wars)	$\approx 5.7\%$	$\approx 15\,\%$

Table 2: US defense spending: pre-9/11 trough vs post-9/11 peak

Sources: "Military expenditure (% of GDP)" for United States — World Bank / SIPRI data. "Share of defense in U.S. federal outlays / budget" — data compiled from historical U.S. federal budget tables.

In the *SA* slack function in (2), ω is the exogenous parameter measuring the relative bureaucratic inefficiency of security provision compared with other government activities. We set the benchmark $\omega = 5$. This choice is somewhat *ad hoc*; however, we will vary it further in the comparative statics. To give a more tangible idea of what $\omega = 5$ implies, consider an initial appropriation to the *SA* of 10% of the budget, i.e., $s^I = 0.1$. Then, by equation (4), such an appropriation, given $\omega = 5$, would result in both slack and rescission equal to 12.5% of the initial appropriation.

The final parameter we directly assign is the time-discount factor, which we set to $B = 0.9$. Since $B < 1$, candidates value future (second-term) office benefits less than current-period payoffs, implying a higher payoff from winning the first election relative to the second. The discount factor, therefore, captures the extent to which candidates are forward-looking: as $B \to 1$ candidates become more reputation-oriented, while lower values imply more short-termism.

Our choice of $B = 0.9$ is consistent with existing political-economy models. Besley and Coate (1997) employ a two-period setting with an implicit discount factor close to 0.9; Persson and Tabellini (2002) use values in the 0.7–1.0 range to capture re-election incentives; Alesina and Tabellini (2007) derive effective political discount factors of 0.6–0.9 from typical 1–3 year policy horizons; and Ashworth (2012) summarize empirical studies indicating that incumbents value future rents at roughly 60–90% of current-term rents.

Benchmark Calibration Procedure

We calibrate the values of the parameters $(\delta, \nu, a, expc)$ by fixing the remaining parameters and imposing the representative benchmark equilibrium conditions. Recall that ν is the parameter such that ν times the expenditure actually implemented by the *SA* on anti-terrorism equals the equilibrium probability of a terrorist attack NOT taking place. The parameter a and $expc$ in turn enters the party's utility function through the term $a\left|s_2^I\left(1-\frac{5s_2^I}{2}\right)-s_1^I\right|^{expc}$ which is subtracted from the incumbency-advantage parameter θ. Therefore, a indicates the intensity of the punishment the voter imposes for a lack of credibility.

The benchmark values for these parameters are calibrated as follows:

1. Benchmark equilibrium actual expenditure on SA equals 7% of total government budget, that is $s_2\left(1-\frac{5s_2}{2}\right)=0.07$. This value is lower than the US figures reported in Table 2, although those may partly reflect slack or bureaucratic inefficiencies. Our benchmark is closer to international patterns. For example, Sandler and Enders (2007) review empirical evidence showing that most OECD countries devote roughly 2–5% of GDP to defense and internal security, which corresponds to about 7% of government spending on average. Spending can reach 10% in high-threat countries such as Israel. Similarly, Stockholm International Peace Research Institute (SIPRI) (2023) report global military expenditure averaging 2.3% of GDP in 2023, with several conflict-prone states spending 7–11%. These figures, however, do not account for bureaucratic inefficiency: International Monetary Fund (2018) document average budget-execution ratios of 85–95% in OECD countries and 70–85% in developing economies.

2. The benchmark equilibrium probability of NO terrorist attack is set to 90%, i.e. $p_{NT}=1-p_T=\nu\times s_2\left(1-\frac{5s_2}{2}\right)=0.9$. Since from 1. above $s_2\left(1-\frac{5s_2}{2}\right)=0.07$, then $\nu=12.85717$. This is broadly consistent with cross-country evidence on the annual likelihood of at least one terrorist incident. Cross-country frequency data in Gassebner and Luechinger (2011) indicate that approximately 5–10% of OECD and 10–20% of non-OECD countries experience at least one attack per year, a pattern also reflected in the Global Terrorism Database (National Consortium for the Study of Terrorism and Responses to Terrorism (START), 2023). Earlier risk assessments such as Martz and Johnson (1987) and Levine (2012) model event probabilities between 0.02 and 0.20, while GTD-based global estimates in Clauset and Woodard (2013) and the venue analysis of Drakos and Gofas (2006) report similar incidence ranges. Taken together, these studies justify calibrating p_T in the 0.05–0.15 interval for low- to high-threat environments.

3. The difference between the first election's proposed expenditure on anti-terrorism and actual expenditure on anti-terrorism is 0.5% of the government budget, that is $s_1-s_2\left(1-\frac{5s_2}{2}\right)=0.005$. If the actual government expenditure on SA is constrained to 7% in the benchmark equilibrium (from the constraint above), then the first–election proposed expenditure is constrained to 7.5% by this third constraint. That

is, in the benchmark equilibrium, the actual expenditure is 6.67% lower than the first–election proposal: $\frac{0.005}{0.075} = 0.0667 \approx 6.67\%$.

Furthermore, we use constraints 1 and 3 to solve for the values of $(\delta_L, a, expc)$ that minimise the sum of squared differences when computing the equilibrium. In other words, the Nash equilibrium is obtained simultaneously with the following optimisation problem:

$$\min_{\delta_L,\, a, expc} \text{ Func} = \left(s_2\Big[1 - \frac{5s_2}{2}\Big] - 0.07\right)^2 + (s_1 - 0.075)^2 .$$

The program then finds the values that generate the Nash equilibrium minimizing the above objective function. In the benchmark case, the equilibrium satisfies each of the two parts of the objective function to within a $10e-7$ margin. Table 3 reports the benchmark calibrated equilibrium values for these parameters.

Parameter	Value
δ_L	0.077371
a	6.209705
$expc$	2.000000

Table 3: Benchmark equilibrium values

4 Comparative Statics from Calibration

In this section, we explore the comparative statics of the model by departing from the benchmark parameterization and varying one parameter at a time. We analyze how key endogenous variables, such as parties' first-election proposed and implemented expenditures, slack, and the probability of a terrorist attack, adjust in response, in equilibrium. We also introduce a Central Planner, which maximizes the voter's expected welfare conditional on first-period information, an addition that provides a useful point of comparison.

Table 4 reports the equilibrium values of the endogenous variables for alternative values of δ_L, the voter's initial preference-for-security parameter. In this table, the bureaucratic inefficiency parameter is at its benchmark value $\omega = 5$. The rest of parameters, other than δ_L, are also at their benchmark values: the exogenous component of Prob[NT] at $\nu = 12.857143$; the credibility parameter at $a = 6.2097048$, the exponent in the credibility function $expc = 2.00$, the incumbency advantage parameter $\theta = .05$, the voter's preference for expenditures other than security and counter-terrorism spending $\beta = .923$; the parameter $\Psi = 1$, which governs the sensitivity of electoral outcomes to policy choices; and the relative value that a win in the second election has for the parties at

$B = 0.9$. Among the endogenous variables, the table shows the first-period *Nash Proposed Expenditure*, the *Nash Implemented Expenditure*, and the corresponding probability that no terrorist attack occurs, *Nash Probability NT*; the Central Planner's implementation *Optimal Implemented Expenditure* and the associated probability of no attack, *Optimal Probability NT*; and the resulting slack variables, *Nash Slack*, *Optimal Slack*.

δ_L	Nash Proposed Expenditure	Nash Implemented Expenditure	Nash Probability NT	Optimal Implemented Expenditure	Optimal Probability NT	Nash Slack	Optimal Slack
0.010	0.01095	0.01455	0.18703	0.02079	0.26728	0.00029	0.00061
0.020	0.02176	0.02599	0.33410	0.04157	0.53450	0.00098	0.00278
0.030	0.03219	0.03582	0.46060	0.05811	0.74718	0.00198	0.00622
0.040	0.04217	0.04456	0.57291	0.07004	0.90051	0.00326	0.01024
0.050	0.05165	0.05241	0.67379	0.07778	1.00000	0.00481	0.01397
0.060	0.06062	0.05946	0.76451	0.07778	1.00000	0.00660	0.01397
0.070	0.06909	0.06579	0.84592	0.07778	1.00000	0.00862	0.01397
0.077	0.07500	0.07000	0.90000	0.07778	1.00000	0.01023	0.01397
0.080	0.07704	0.07141	0.91808	0.07778	1.00000	0.01082	0.01397
0.090	0.08447	0.07632	0.98120	0.07778	1.00000	0.01318	0.01397
0.100	0.09041	0.07778	1.00000	0.07778	1.00000	0.01397	0.01397
0.110	0.09638	0.08029	1.00000	0.07778	1.00000	0.01546	0.01397
0.120	0.10238	0.08362	1.00000	0.07778	1.00000	0.01772	0.01397
0.130	0.10792	0.08636	1.00000	0.07951	1.00000	0.01989	0.01498
0.140	0.11304	0.08860	1.00000	0.08181	1.00000	0.02193	0.01645
0.150	0.11777	0.09042	1.00000	0.08381	1.00000	0.02384	0.01786

Table 4: Choices by parties in a Nash equilibrium, and choices by a Central Planner, as a function of δ_L, for $\omega = 5$. Values rounded (main columns: 5 decimals; slack: 6 decimals). NT = No Terrorist Attack.

Before analyzing the results, we find it convenient to explain the trade-offs the political parties face when choosing the proposal of security expenditure for the first period, s_1. One of these trade-offs arises from the interest in winning two sequential elections. On the one hand, parties aim to win the first election and therefore may attempt to appeal to a voter whose preferred expenditure is $\frac{\delta_L}{\beta+\delta_L}$.

On the other hand, the parties want to be competitive in the second election. However, voter's preferences might have changed by then, and the (myopic) voter does not internalize the mechanism through which security expenditure affects the probability of a terrorist attack and her own preferences. As a result, following the realization of a terrorist attack, the voter's preferred level of expenditure may increase to $\frac{\delta_H}{\beta+\delta_H}$ ($\delta_H > \delta_L$), albeit the voter does not anticipate this change at the first election. Yet, in case of winning the first election, it is not the first-period proposal, but rather the expenditure implemented,

$\left(s - \frac{5s^2}{2}\right)$, which determines both the probability of a terrorist attack and the probability that the party wins the second election.

Thus, why do parties not simply set s_1 to win the first election and then independently adjust the implemented expenditure to best compete in the second election? The probability that the incumbent wins the second election is not independent of the policy proposed in the first election. In particular, the voter evaluates the incumbent's credibility by comparing the implemented expenditure level with the promised level, s_1. This credibility assessment directly influences the voter's re-election decision. Ceteris paribus, a larger deviation of the realized expenditure from the promised expenditure s_1 reduces the incumbent's perceived credibility and, consequently, lowers the probability that the voter supports the incumbent in the second election.

In addition, the implemented level of security (SA) spending is affected by the intrinsic inefficiency of anti-terrorism expenditures. Higher allocations to SA generate waste and crowd out other welfare-enhancing public goods. This inefficiency introduces a constraint on the incumbent's implementation decision when seeking to maximize the probability of re-election, as both the voter's utility from public expenditures and the incumbent's perceived credibility depend on the realized level of SA spending. As a result, the incumbent faces a trade-off between enhancing security and preserving voter welfare and credibility.

Finally, the incumbent's implementation decision is influenced by the anticipated strategic behavior of the challenger in the second election. The challenger possesses three distinct informational and strategic advantages. First, the challenger observes the realization of the terrorist threat before proposing a policy platform, whereas the incumbent must implement an expenditure level under uncertainty, conditioning only on the probability of an attack, even though this probability is itself affected by the level of expenditure. Second, the challenger is not constrained by the inefficiency of SA spending when formulating a policy proposal, since the voter does not directly observe the degree of inefficiency. Third, the challenger's proposal is not subject to retrospective credibility evaluation by the backward-looking voter, as the challenger has no prior governing record against which promises can be compared.

In contrast, even when the Central Planner accounts for the inefficiency of SA spending, they fully internalize since inception the possibility that voters' preferences may shift following the realization of a terrorist attack. Moreover, they do not face credibility concerns or threat from any challenger, which arise purely from electoral contests.

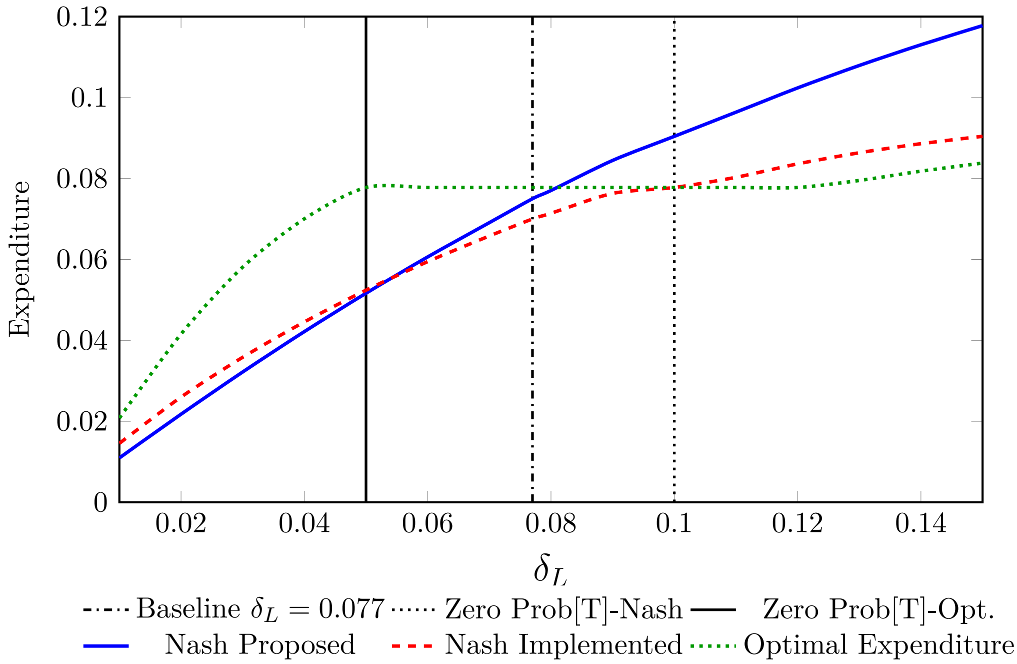

Figure 3: Expenditures on Security & Anti-terrorism as a function of δ_L when $\omega = 5$.

With these in mind, we explore Figure 3, which shows the equilibrium choices of the political parties and the Central Planner as functions of δ_L.

Statement 1 (Political platform bias). *For δ_L sufficiently small, in a Nash equilibrium, political parties propose to spend less on Security & anti-terrorism than what is actually implemented. For δ_L sufficiently large, in a Nash equilibrium, political parties propose more than what is implemented.*

The first noticeable result is that the expenditure proposed by the parties in the first election (blue line) is lower than the expenditure actually implemented (dashed-red line) for very small values of δ_L.

One way to interpret this result is to note that the proposed expenditure is electorally relevant in the first election, whereas the implemented expenditure matters for the second. In the first election, proposing a sufficiently low level of s_1—close to $\frac{\delta_L}{\delta_L+\beta}$—increases the probability of winning. This incentive is reinforced by the fact that the voter does not internalize the possibility of a terrorist attack (T) and the resulting upward adjustment in preferences.

Yet, because low security spending increases the probability of a terrorist attack, the implemented expenditure will have to be larger to appeal to a voter whose preferences may shift upward following a terrorist attack. This explains why the implemented expenditure exceeds the initial proposal. Still, the credibility constraint tightens the distance between proposed and implemented expenditures.

In an environment characterized by an initial low demand for security—for instance, following a prolonged period without terrorist attacks—the first election becomes more salient than the second. To compete in the second election, the incumbent faces the prospect of a likely upward shift in voter preferences, while the challenger is not constrained by past promises. In contrast to the credibility effect, what prevents implemented expenditure from falling even closer to the proposal is the relatively weak effect of waste due to *SA* inefficiencies at low expenditure levels.

This pattern reverses for high values of δ_L, where the first-period proposal exceeds the implemented expenditure. When baseline security concerns are strong, the environment is one with large expected wastes. At the same time, the probability of a terrorist attack becomes negligible—indeed, zero for $\delta_L > 0.1$ (dotted-vertical line)—which limits/eliminates the pull effect of an upward revision of voter preferences. These two effects, high overall waste and reduced risk of preference shifts, more than compensate the electoral cost of credibility loss, inducing parties to implement less than promised.

Statement 2 (Political amplification of the threat cycle)**.** *For δ_L sufficiently small, in a Nash equilibrium with political parties, there is less expenditure on Security & anti-terrorism than in the case with a Central Planner. For δ_L sufficiently large, in a Nash equilibrium, there is more expenditure on Security & anti-terrorism than in the case with a Central Planner.*

A second relevant takeaway from Figure 3 emerges when comparing expenditure under electoral incentives with the efficient allocation chosen by a Central Planner (green-dotted line). The counterpart is Figure 4, which shows that, with electoral competition, society underwastes in (and underprovides) security when baseline preferences for security are low (small δ_L), and overwastes in (and overprovides) security when voters place a high value on security (large δ_L).

In Figure 3, when $\omega = 5$, political parties lead to a security overprovision when the endogenous probability of a terrorist attack (T) is zero for both the Central Planner and the parties. Thus, electoral competition typically[8] leads to a higher expected probability of a terrorist attack, particularly when society does not place an ex-ante high value on security (low δ_L).

[8]Note that this is not always the case. For example, when $\omega = 7$, they cross when $\delta = .66$ and the probability of no terrorist attack is 92 percent.

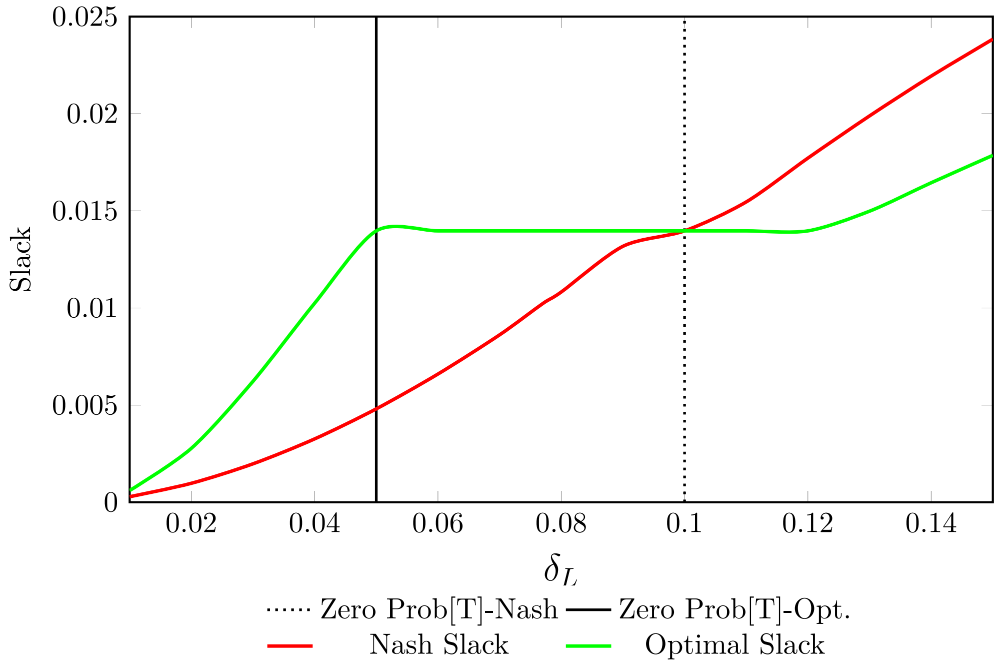

Figure 4: Slack under Nash equilibrium and the Central Planner as a function of δ_L when $\omega = 5$.

The explanation for this pattern is on three key features that distinguish the Central Planner from the political parties: (i) the planner does not need to pander to a myopic voter in the first period and therefore fully internalizes the effect of a potential terrorist attack on subsequent voter preferences; (ii) the planner faces no credibility constraint; and (iii) the planner does not have to compete with a fully informed uncommitted challenger in the second election.

When δ_L is small, parties have incentives to pander to a voter who does not anticipate a potential upward shift in preferences following a terrorist attack. This keeps the first-period expenditure that parties proposed low wrt. a planner's. At the same time, credibility considerations link implemented expenditure to the initial promise, since deviations are punished electorally in the second election. As a result, both promised and implemented expenditures remain below the planner's efficient level.

When δ_L is sufficiently large, the equilibrium level of SA expenditure implemented under electoral competition reduces the probability of a terrorist attack to zero. This outcome first arises when $\delta_L \approx 0.1$. For values of δ_L at or above this threshold, the expenditure under a social planner would likewise eliminate the probability of an attack. Thus, for $\delta_L \geq 0.1$, the incumbent's allocation of SA expenditure to maximize the probability of re-election differs from the allocation of the social planner due to the presence of a credibility constraint.

In particular, because the SA expenditure promised in the first election is increasing in δ_L, the incumbent faces an incentive to implement a level of expenditure consistent with prior commitments in order to preserve credibility.[9] This constraint distorts the incumbent's allocation decision, resulting in an overprovision of SA expenditure relative to the socially optimal level. This explains why the slack is larger under electoral competition than under the optimal solution by a central planner, as Figure 4 shows. Consequently, even though both the incumbent and the social planner would eliminate the risk of a terrorist attack at high values of δ_L, electoral incentives due to credibility considerations induce inefficiently high levels of SA spending under electoral competition.

The incentive to honor electoral promises does not, however, imply that these promises are fully implemented. At high values of δ_L, Figure 3 shows that implemented expenditure exceeds the planner's level but remains below the first-election promise. Credibility concerns encourage the incumbent to fulfill that promise, while bureaucratic inefficiency and the electoral consequences of crowding out other public goods limit implementation. In the calibrated exercise, proposed expenditure rises faster than implemented expenditure, so the proposed-implementation gap widens even as security provision remains above the planner's level.

Figure 5 repeats the security-demand exercise without bureaucratic inefficiency ($\omega = 0$); the corresponding numerical values are reported in Table A1, Appendix A2. Under this assumption, SA generates no slack under either an electoral system or a central planner.

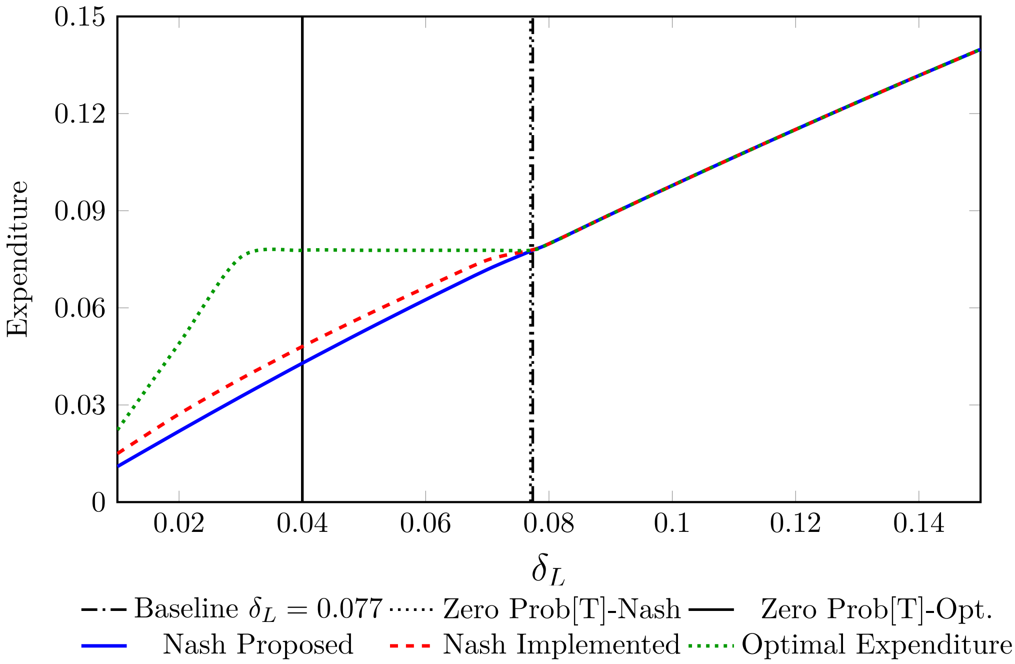

Figure 5: Expenditures on Security & Anti-terrorism as a function of δ_L when $\omega = 0$.

[9]Because of the credibility constraint, there are also incentives to increase implementation to compete with a challenger that does not face the waste constraint.

Figure 5 highlights the role of security and anti-terrorism (SA) bureaucratic inefficiency—or its absence—in shaping the performance of the political system. When security provision is efficient (i.e., involves no waste) and security concerns (δ_L) are sufficiently high—such as in a society that has recently experienced terrorist attacks—electoral competition induces parties to promise and implement the efficient level of security expenditure. This is true in general as we prove in Statement A1, Appendix A1. In this case, the probability of a terrorist attack converges to zero and is exactly zero at the baseline value and for all higher values, $\delta_L \geq 0.077$.

This outcome occurs because, when the probability of a terrorist attack is constrained to zero, there is no possibility that voters' preferences will change. In addition, in the absence of bureaucratic waste, the credibility constraint becomes irrelevant, allowing candidates to freely propose and implement the efficient policy. When the probability of a terrorist attack is zero, the only potential divergence between the expenditure proposed during the first election and the level implemented afterward arises from inefficiencies within the security agency.

Consequently, when δ_L is sufficiently high that the voter prefers a level of expenditure that eliminates the probability of terrorist attacks in the absence of bureaucratic waste, candidates propose the voter-preferred level of spending during the first election. The winning candidate then implements this same level of expenditure after the election, resulting in no credibility loss.

Statement 3 (SA efficiency erodes political distortions)**.** *Assume $\omega = 0$. Then, for δ_L sufficiently large, in a Nash equilibrium with political parties, expenditures on security, both promised and implemented, converge to the case with a Central Planner.*

By contrast, when δ_L falls below the benchmark level and voters do not internalize the risk of future terrorist attacks and the consequent upward shift in preferences, the political system generates underprovision of security. In this sense, electoral competition leads parties to pander to a myopic electorate. Again, credibility concerns keep promises close to implementation; however, implementation exceeds promised levels, as it must target voters who may place a higher value on security.

We next vary relative bureaucratic inefficiency, ω. Higher values reduce the effective security delivered by a given appropriation. Figure 6 presents this comparative-static exercise, holding the other parameters at their benchmark values, including $\delta_L = 0.077371$. The detailed numerical results appear in Table A2 in Appendix A2. Appendix A5.1 provides descriptive empirical motivation: terrorism exposure is greater among countries with high levels of government inefficiency, measured using the reversed standardized government-effectiveness index. This association motivates examining how implementation inefficiency affects security provision and terrorism risk. The governance indicators capture the broader institutional environment in which security agencies operate.

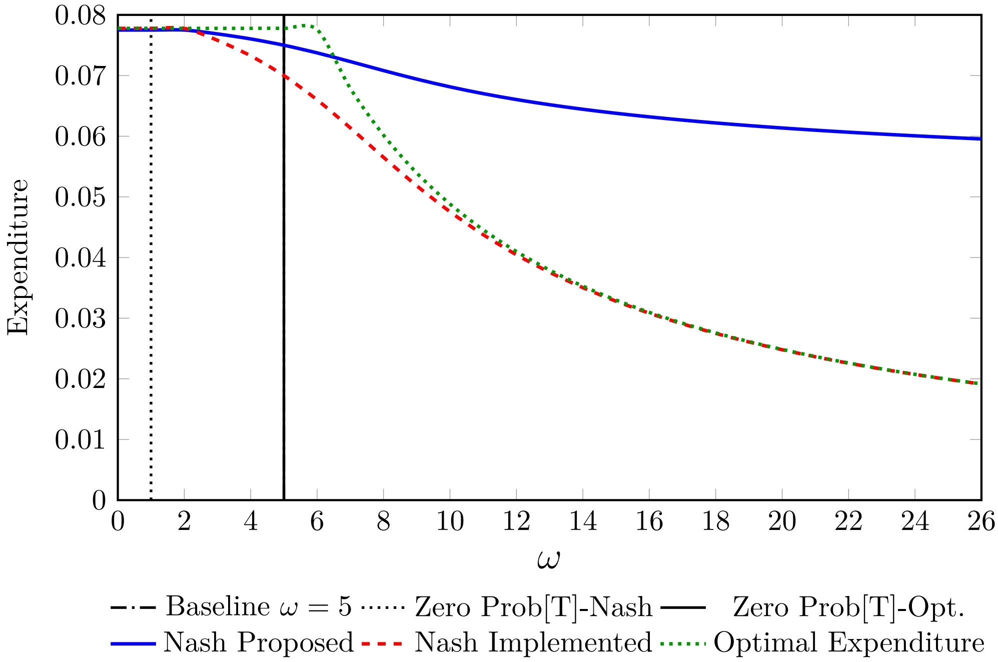

Figure 6: Expenditures on Security & Anti-terrorism as a function of ω when $\delta_L = 0.077371$.

The main societal trade-off in the provision of security and anti-terrorism services can be stated as follows. Relative to other public goods, a Security Agency (*SA*) may require lower transparency to operate effectively, for instance due to intelligence-gathering and related activities necessary to prevent terrorist attacks. However, lower transparency increases monitoring costs, worsening the agency problem, thereby raising waste and making prevention more costly.[10]

Figure 6 illustrates the latter side of this trade-off: beyond a certain threshold, higher *SA* inefficiency reduces incentives to invest in security and increases the probability of a terrorist attack, both under electoral competition and under a Central Planner.

Statement 4 (Toward state failure)**.** *For sufficiently large ω, expenditures on Security & anti-terrorism under a Central Planner and under a Nash equilibrium among political parties converge and both decrease.*

As Figure 6 indicates, for sufficiently large values of ω, *SA* slack—being monotonic in implemented expenditure, other things the same—starts to decrease. At high levels of relative inefficiency of the SA, the contraction in security spending induced by higher opportunity costs of waste dominates the direct increase in inefficiency due to higher ω, leading to a net reduction in slack. Recall that the implemented expenditure is given by $s - \frac{\omega s^2}{2}$, while Slack $= \frac{\omega s^2}{4}$.

[10] In our setting, ω can be interpreted as capturing the differential inefficiency of the Security Agency relative to agencies providing other public goods.

Importantly, Figure 6 shows the fundamental role of ω in shaping political incentives behind parties' promises and implementation of security. At one extreme, where there is zero relative technical inefficiency ($\omega = 0$), elections entail no political cost. Since the probability of a terrorist attack is zero and security spending involves no increase in waste, credibility concerns do not arise. Moreover, voter myopia is immaterial since the probability of terrorist attack is zero and therefore there are no incentives for parties to underpromise in the first period. Therefore, proposed and implemented expenditures on security under elections are very close, and the latter coincides with the optimal implementation.

At the other extreme, for large values of ω, the expenditure implemented under political competition (dashed-red line) converges to the optimal level (dotted-green line) from below as ω increases. Intuitively, honoring first-period promises would impose substantial opportunity costs in terms of foregone spending on other public goods when inefficiency is high. As a result, first-period promises and second-period implementation increasingly disentangle. In particular, the equilibrium electoral strategies involve high promised expenditure to pander a myopic voter and low level of implementation to ditch the waste involved in allocating resources to the SA.

For intermediate values of ω, the Nash-equilibrium levels of proposed and implemented expenditures resulting from candidates' strategic interaction differ from the optimal expenditure that maximizes expected voter welfare (i.e., the central planner's solution). For example, under the Nash equilibrium, the implemented expenditure induces a positive probability of a terrorist attack (to the right of the dotted-black vertical line), while under the planner's allocation the probability is zero (to the left of the solid-black vertical line). In this region, parties typically implement less than they propose, and they propose less than the optimal level of expenditure.

In summary, for intermediate values of ω—for which the probability of a terrorist attack is strictly positive given the proposed, implemented, or optimal SA expenditure—the proposed expenditure exceeds both the implemented and the socially optimal levels. This distortion arises because candidates pander to voters who are unaware of SA inefficiency. At the same time, the socially optimal level exceeds the implemented level due to the electoral incentive constraints faced by the incumbent in the second election.

Although these second election incentive constraints exert opposing pressures on the level of implemented expenditure, the distortion leading to underprovision (implemented $<$ optimal) dominates. While credibility considerations may induce the incumbent to increase expenditure above the socially optimal level, the incumbent must also account for the challenger's policy proposals. These proposals are state-contingent, as voter preferences differ depending on whether a terrorist attack has occurred prior to the second election.

An increase in expenditure above the Nash level raises the probability of avoiding an attack; however, from the perspective of maximizing reelection probability, places excessive weight on the utility differential between the implemented policy and the challenger's policy in the no-attack state.

This source of underprovision vanishes when voter preferences are invariant to the occurrence of a terrorist attack. In such a setting, the credibility constraint becomes dominant, resulting in implemented expenditure exceeding the socially optimal level.

5 Conclusions

National security and anti-terrorism policies are an integral part of the political discourse in established democracies. It is a reasonable claim that policies in this domain are formulated through the strategic interaction of three key stakeholders (Downsian politicians, security agency, and voters) under heterogeneous institutional settings. Yet, systematic theoretical analysis of such strategic interactions under endogenously random terrorist attack and heterogeneous institutions and commitment is rare. This is what we set out to do in this paper.

We present a model where Downsian political parties, whose sole interest is to get elected, choose their platforms given preferences and parameters. In particular, the incumbent chooses their security spending and the security agency implements them with a slack. The security agency maximizes the slack at the expense of the civil liberties of voters and other public goods, creating a trade-off between these goods and national security. In an environment of relative peace, voters' relative preference for civil liberties remains high. However, following a terrorist attack, voters prefer national security more, albeit at the expense of civil liberties and the rest of public services. The model includes two elections under credibility-as-commitment technology and endogenously random terrorist attacks. We calibrate the model using country-level international data and offer numerical solutions and comparative statics. Furthermore, we motivate our modeling assumptions and results using observational data.

We obtain the following findings. First, we find that in an environment of relative peace with relatively high voter preference for civil liberties, Downsian political parties underpropose security spending in a Nash equilibrium. In other words, they propose less spending relative to what is implemented. In contrast, in an environment of relative insecurity, political parties propose more spending than what is implemented.

Second, political parties amplify the threat cycle. They allocate less on security than a central planner would during periods of relative peace, when voters place a high value on civil liberties. In contrast, when perceived threats and voters' preference for security are high, political parties allocate more on security than a central planner would. Such overreaction is not unusual. In fact, our model tracks the reality of security spending fairly well. Security spending by the EU and non-US members of NATO is a case in point where there is a history of significant underspend during times of relative peace and a significant increase in urgency to increase security spending following the Ukraine war, US threats over Greenland, and US threats of withdrawal from NATO.

Third, we find that when security provision is efficient and involves no slack, the political parties promised and SA implemented security expenditures converge with the Central Planner for sufficiently high levels of perceived threat. This result points out that inefficiency in security provision is central to our argument.

Finally, the same argument yields a contrasting result: when government is highly inefficient and bureaucratic waste is very high, as in a failed state, security spending under both the central planner's allocation and the political parties' Nash equilibrium converges to the same extremely low level, rendering a terrorist attack certain.

What are the potential implications of our key findings? It is apparent that a key feature of social choice on security spending under constraints and information asymmetry is a game of second best. Both the voters and the politician have information asymmetry over the security agency in terms of the rationale for a particular level of spending, because micro-level disclosure of spending items could damage the operational capacity of the agency and counter-terrorism. However, such negative effects could perhaps be avoided with some degree of disclosure and audit at a relatively macro level. Macro audit and disclosure under an independent agency, along the lines of the UK's Office for Budget Responsibility (OBR), might help address security agency slack.

Another potential challenge for a sustainable security spending policy is its high level of volatility in response to swings in voter sentiment following a terrorist attack. Security perception of voters might follow cyclical patterns or even experience regime switches, whereas security spending would need to be smooth in order to reduce threats. Political competition with Downsian politicians will not address such a challenge, as the politicians would try to overreact in order to impress voters. Therefore, political competition will aggravate spending volatility rather than reducing it when it comes to security spending. This is in sharp contrast to the observations made by Alesina (1988). A possibility could be the introduction of long-term spending rules with an allowance for reviewing the rules every 3-4 years. The spending rules could remain confidential to the highest levels of government, but the security agency and the government could be required to file independent annual reports on adherence to such rules to an institution like the OBR. The OBR can then inform the public on whether the security agency and the government are adhering to the spending rules.

References

Abadie, A. (2006). Poverty, political freedom, and the roots of terrorism. *American Economic Review*, 96(2):50–56.

Adserà, A., Boix, C., and Payne, M. (2003). Are you being served? political accountability and quality of government. *Journal of Law, Economics, and Organization*, 19:445–490.

Alesina, A. (1988). Credibility and policy convergence in a two_party system with rational voters. *The American Economic Review*, 78(4):796–805.

Alesina, A. and Angeletos, G.-M. (2005). Corruption, inequality, and fairness. *Journal of Monetary Economics*, 52:1227–1244.

Alesina, A. and Tabellini, G. (2007). Bureaucrats or politicians? part i: A single policy task. *American Economic Review*, 97(1):169–179.

Anderton, C. H. and Brauer, J. (2021). Mass atrocities and their prevention. *Journal of Economic Literature*, 59(4):1240–1292.

Anderton, C. H. and Carter, J. R. (2019). *Principles of conflict economics: the political economy of war, terrorism, genocide, and peace*. Cambridge University Press.

Arce, D. G. and Sandler, T. (2005). Counterterrorism: A game-theoretic analysis. *Journal of Conflict Resolution*, 49(2):183–200.

Ashworth, S. (2012). Electoral accountability: Recent theoretical and empirical work. *Annual Review of Political Science*, 15(1):183–201.

Becker, G. S. (1968). Crime and punishment: An economic approach. *Journal of Political Economy*, 76:169–217.

Berman, N., Couttenier, M., Rohner, D., and Thoenig, M. (2017). This mine is mine! how minerals fuel conflicts in africa. *American Economic Review*, 107(6):1564–1610.

Besley, T. and Coate, S. (1997). An economic model of representative democracy. *The quarterly journal of economics*, 112(1):85–114.

Bhattacharyya, S. and Mamo, N. (2021). Natural resources and conflict in africa: what do the data show? *Economic Development and Cultural Change*, 69(3):903–950.

Blattman, C. and Miguel, E. (2010). Civil war. *Journal of Economic literature*, 48(1):3–57.

Breton, A. and Wintrobe, R. (1975). The equilibrium size of a budget-maximizing bureau: A note on niskanen's theory of bureaucracy. *Journal of Political Economy*, 83:195–207.

Breton, A. and Wintrobe, R. (1982). *The Logic of Bureaucratic Control*. Cambridge University Press, Cambridge.

Brownfield, A. (1977). The inherent inefficiency of government bureaucracy. Foundation for Economic Education.

Bueno de Mesquita, E. (2007). Politics and the suboptimal provision of counterterror. *International Organization*, 61(1):9–36.

Busch, T. (2002). Slack in public administration: conceptual and methodological issues. *Managerial Auditing Journal*, 17(3):153–159.

Cayford, M. and Pieters, W. (2020). Effectiveness fettered by bureaucracy : why surveillance technology is not evaluated. *Intelligence and National Security*, 35(7):1026–1041.

Clauset, A. and Woodard, R. (2013). Estimating the historical and future probabilities of large terrorist events. *The Annals of Applied Statistics*, pages 1838–1865.

Collier, P. and Hoeffler, A. (1998). On economic causes of civil war. *Oxford economic papers*, 50(4):563–573.

Collier, P. and Hoeffler, A. (2004). Greed and grievance in civil war. *Oxford economic papers*, 56(4):563–595.

Daxecker, U. (2017). Dirty hands: Government torture and terrorism. *Journal of Conflict Resolution*, 61(6):1261–1289.

Daxecker, U. E. and Hess, M. L. (2013). Repression hurts: coercive government responses and the demise of terrorist campaigns. *British Journal of Political Science*, 43(3):559–577.

Di Lonardo, L. and Dragu, T. (2025). When salience undermines representation: democratic dilemmas in security and counterterrorism policy. *Political Science Research and Methods*, page 1–9.

Drakos, K. and Gofas, A. (2006). In search of the average transnational terrorist attack venue. *Defence and Peace Economics*, 17(02):73–93.

Enders, W. and Sandler, T. (2011). *The Dilemma of Liberal Democracies*, page 28–60. Cambridge University Press.

Esteban, J., Mayoral, L., and Ray, D. (2012). Ethnicity and conflict: An empirical study. *American Economic Review*, 102(4):1310–1342.

Esteban, J. and Ray, D. (2008). On the salience of ethnic conflict. *American Economic Review*, 98(5):2185–2202.

Ferreira, F. and Gyourko, J. (2009). Do political parties matter? evidence from us cities. *The Quarterly journal of economics*, 124(1):399–422.

Gassebner, M. and Luechinger, S. (2011). Lock, stock, and barrel: A comprehensive assessment of the determinants of terror. *Public Choice*, 149(3):235.

Goel, R. K. and Nelson, M. A. (1998). Corruption and government size: A disaggregated analysis. *Public Choice*, 97:107–120.

Guriev, S., Melnikov, N., Silva, J., and Zhuravskaya, E. (2025). Mobile broadband and the decline of incumbency advantage.

Harms, P. and Ursprung, H. W. (2002). Do civil and political repression really boost foreign direct investments? *Economic inquiry*, 40(4):651–663.

Hartley, K. (2011). *The Economics of Defence Policy: A New Perspective (1st ed.)*. Routledge.

Heckelman, J. C. (2000). Economic freedom and economic growth: A short-run causal investigation. *Journal of Applied Economics*, 3(1):71–91.

Helliwell, J. F. (1994). Empirical linkages between democracy and economic growth. *British journal of political science*, 24(2):225–248.

Holmström, B. and Milgrom, P. (1991). Multitask principal-agent analyses: Incentive contracts, asset ownership, and job design. *Journal of Law, Economics, & Organization*, 7:24–52. Special Issue.

International Monetary Fund (2018). Public Expenditure and Financial Accountability (PEFA) Handbook. Technical report, International Monetary Fund, Washington, DC.

Kotera, G., Okada, K., and Samreth, S. (2012). Government size, democracy, and corruption: An empirical investigation. *Economic Modelling*, 29(6):2340–2348.

Landes, W. M. (1978). An economic study of us aircraft hijacking, 1961-1976. *the Journal of Law and Economics*, 21(1):1–31.

Lee, D. S. (2008). Randomized experiments from non-random selection in us house elections. *Journal of Econometrics*, 142(2):675–697.

Levine, E. S. (2012). Estimating conditional probabilities of terrorist attacks: modeling adversaries with uncertain value tradeoffs. *Risk Analysis: An International Journal*, 32(2):294–303.

Li, C. M., Sasso, G., and Turner, I. R. (2023). Hierarchical control. Vanderbilt University Working Paper.

Martz, H. F. and Johnson, M. E. (1987). Risk analysis of terrorist attacks. *Risk Analysis*, 7(1):35–47.

Miguel, E., Satyanath, S., and Sergenti, E. (2004). Economic shocks and civil conflict: An instrumental variables approach. *Journal of political Economy*, 112(4):725–753.

Miller, G. J. (1977). Bureaucratic compliance as a game on the unit square. *Public Choice*, 29:48–50.

Mueller, D. C. (2003). *Public choice III*. Cambridge University Press.

National Consortium for the Study of Terrorism and Responses to Terrorism (START) (2023). Global Terrorism Database (GTD) [data file]. https://www.start.umd.edu/gtd. Data retrieved from the 2023 Global Terrorism Database release (covering incidents 1970–2022).

Niskanen, W. A. (1971). *Bureaucracy and Representative Government*. Aldine-Atherton, Chicago, IL.

Park, S.-J. and Jang, S. (2021). Asymmetric information and excess budget: the influence of performance-based budgeting on budgetary slack in us states. *International Review of Public Administration*, 26(4):353–372.

Persson, T. and Tabellini, G. (2002). *Political economics: explaining economic policy*. MIT press.

Rose-Ackerman, S. (1978). *Corruption: A Study in Political Economy*. Academic Press, Cambridge.

Rose-Ackerman, S. (1999). *Corruption and Government: Causes, Consequences, and Reform*. Cambridge University Press, Cambridge.

Sandler, T. (2014). The analytical study of terrorism: Taking stock. *Journal of Peace Research*, 51(2):257–271.

Sandler, T. and Enders, W. (2007). An economic perspective on transnational terrorism. In *The economic analysis of terrorism*, pages 29–44. Routledge.

Sandler, T. and Enders, W. (2008). *The Political Economy of Terrorism*. Cambridge University Press, Cambridge, UK.

Stockholm International Peace Research Institute (SIPRI) (2023). Sipri military expenditure database. https://www.sipri.org/databases/milex. Data retrieved from https://www.sipri.org/databases/milex.

The White House (2025). Countering domestic terrorism and organized political violence. https://www.whitehouse.gov/presidential-actions/2025/09/countering-domestic-terrorism-and-organized-political-violence/. Presidential Action, Accessed 30 Dec 2025.

Treisman, D. (2000). The causes of corruption: A cross-national study. *Journal of Public Economics*, 76:399–457.

Wilson, J. (1989). *Bureaucracy: What Government Agencies Do and why They Do it*. Basic Books.

Wyckoff, P. G. (1990). The simple analytics of slack-maximizing bureaucracy. *Public Choice*, 67(1):35–47.

A Appendix

Appendix contents

A1 Additional theoretical results

Statement A1. *Consider an environment where the probability of a terrorist attack is zero. Assume $\omega = 0$ and $0 < B \leq 1$. Then, the unique symmetric NE is given by $s_1^i = s_2^i = \frac{\delta_L}{\delta_L+\beta}$, $i = D, R$.*

Proof. Suppose PTA=0 and $\omega = 0$, then

$$\begin{aligned} EU_D =& P_1(D) \\ &+B\Big[P_1(D)P_2^{NT}(D; D = \text{incumbent}) \\ &+[1 - P_1(D)]P_2^{NT}(D; D = \text{challenger})\Big] \end{aligned} \tag{A1}$$

The first order conditions (FOCs) are,

$$\begin{aligned} \frac{\partial EU_D}{\partial s_1^D} &= 0 = P_1(D)_{s_1} \\ &\quad + BP_1(D)_{s_1}P_2^{NT}(D : D = inc.) + BP_1(D)P_2^{NT}(D : D = inc.)_{s_1} \\ &\quad - BP_1(D)_{s_1}P_2^{NT}(D : D = chal.) \\ \frac{\partial EU_D}{\partial s_2^D} &= 0 = BP_1(D)P_2^{NT}(D : D = inc.)_{s_2} \end{aligned} \tag{A2}$$

where

$$\begin{aligned} &P_1(D)_{s_1} = \Psi\Big[-\beta(1-s_1)^{\beta-1}(s_1)^{\delta} + (1-s_1)^{\beta}\delta(s_1)^{\delta-1}\Big] \\ &P_2^{NT}(D : D = inc.)_{s_2} = -\beta(1-s_2)^{\beta-1}(s_2)^{\delta} + (1-s_2)^{\beta}\delta(s_2)^{\delta-1} - expc \times a|s_2 - s_1|^{expc-1} \\ &P_2^{NT}(D : D = inc.)_{s_1} = expc \times a|s_2 - s_1|^{expc-1} \end{aligned}$$

By symmetry, $s_n^D = s_n^R \equiv s$, $n = 1, 2$, the FOCs boil down to,

$$\begin{aligned} &expc \times a|s_2 - s_1|^{expc-1} = -\beta(1-s_2)^{\beta-1}(s_2)^{\delta} + (1-s_2)^{\beta}\delta(s_2)^{\delta-1} \\ &expc \times a|s_2 - s_1|^{expc-1} = \Big(\beta(1-s_1)^{\beta-1}(s_1)^{\delta} - (1-s_1)^{\beta}\delta(s_1)^{\delta-1}\Big) \times \frac{1+B[\cdot]}{B \times P_1(D)} \end{aligned} \tag{A3}$$

where, $\frac{1+B[\cdot]}{2} = \frac{1}{2} + B\Big[\Psi \times \Big((1-s_2)^{\beta}(s_2)^{\delta} - (\frac{\beta}{\beta+\delta})^{\beta}(\frac{\delta}{\beta+\delta})^{\delta}\Big)\Big] \geq 0 \quad \text{if} \quad 0 < B \leq 1.$

One solution of the system A3 is

$$s_1 = s_2 = \frac{\delta}{\delta+\beta}. \tag{A4}$$

Furthermore, this is the unique solution:

Suppose $s_2 > s_1$, thus $a(s_2 - s_1)^{expc-1} > 0$ and

$$-\beta(1-s_2)^{\beta-1}(s_2)^{\delta} + (1-s_2)^{\beta}\delta(s_2)^{\delta-1} > 0$$

which implies

$$s_2 < \frac{\delta}{\delta+\beta}.$$

As $s_2 > s_1$, and then $a(s_2 - s_1)^{expc-1} > 0$, it must be that

$$\beta(1-s_1)^{\beta-1}(s_1)^{\delta} - (1-s_1)^{\beta}\delta(s_1)^{\delta-1} > 0$$

which implies

$$s_1 > \frac{\delta}{\delta+\beta}.$$

which contradicts $s_2 > s_1$.

A similar argument applies to the case in which $s_2 < s_1$. $\square$

Full version of party D's Utility Function

$$
\begin{aligned}
U^D &= \frac{1}{2} + \Psi\left[\left(1-s_1^D\right)^\beta \left(s_1^D\right)^{\delta_L} - \left(1-s_1^R\right)^\beta \left(s_1^R\right)^{\delta_L}\right] \\
&+ \left(\frac{1}{2} + \Psi\left[\left(1-s_1^D\right)^\beta \left(s_1^D\right)^{\delta_L} - \left(1-s_1^R\right)^\beta \left(s_1^R\right)^{\delta_L}\right]\right) \times \nu s_2^D \left[1 - \frac{\omega s_2^D}{2}\right] \times \\
&B\Bigg(\frac{1}{2} + \Psi\Bigg[\left(1 - s_2^D\left[1-\frac{\omega s_2^D}{4}\right]\right)^\beta \left(s_2^D\left[1-\frac{\omega s_2^D}{2}\right]\right)^{\delta_L} - \left(\frac{\beta}{\beta+\delta_L}\right)^\beta \left(\frac{\delta_L}{\beta+\delta_L}\right)^{\delta_L} \\
&+ \theta - a\left|s_2^D\left[1-\frac{\omega s_2^D}{2}\right] - s_1^D\right|^{expc}\Bigg]\Bigg) \\
&+ \left(\frac{1}{2} + \Psi\left[\left(1-s_1^D\right)^\beta \left(s_1^D\right)^{\delta_L} - \left(1-s_1^R\right)^\beta \left(s_1^R\right)^{\delta_L}\right]\right) \times \left(1 - \nu s_2^D \left[1 - \frac{\omega s_2^D}{2}\right]\right) \times \\
&B\Bigg(\frac{1}{2} + \Psi\Bigg[\left(1 - s_2^D\left[1-\frac{\omega s_2^D}{4}\right]\right)^\beta \left(s_2^D\left[1-\frac{\omega s_2^D}{2}\right]\right)^{\delta_H} - \left(\frac{\beta}{\beta+\delta_H}\right)^\beta \left(\frac{\delta_H}{\beta+\delta_H}\right)^{\delta_H} \\
&+ \theta - a\left|s_2^D\left[1-\frac{\omega s_2^D}{2}\right] - s_1^D\right|^{expc}\Bigg]\Bigg) \qquad \text{(A5)} \\
&+ \left(\frac{1}{2} + \Psi\left[\left(1-s_1^R\right)^\beta \left(s_1^R\right)^{\delta_L} - \left(1-s_1^D\right)^\beta \left(s_1^D\right)^{\delta_L}\right]\right) \times \nu s_2^R \left[1 - \frac{\omega s_2^R}{2}\right] \times \\
&B\Bigg(\frac{1}{2} + \Psi\Bigg[\left(\frac{\beta}{\beta+\delta_L}\right)^\beta \left(\frac{\delta_L}{\beta+\delta_L}\right)^{\delta_L} - \left(1 - s_2^R\left[1-\frac{\omega s_2^R}{4}\right]\right)^\beta \left(s_2^R\left[1-\frac{\omega s_2^R}{2}\right]\right)^{\delta_L} \\
&- \theta + a\left|s_2^R\left[1-\frac{\omega s_2^R}{2}\right] - s_1^R\right|^{expc}\Bigg]\Bigg) \\
&+ \left(\frac{1}{2} + \Psi\left[\left(1-s_1^R\right)^\beta \left(s_1^R\right)^{\delta_L} - \left(1-s_1^D\right)^\beta \left(s_1^D\right)^{\delta_L}\right]\right) \times \left(1 - \nu s_2^R \left[1 - \frac{\omega s_2^R}{2}\right]\right) \times \\
&B\Bigg(\frac{1}{2} + \Psi\Bigg[\left(\frac{\beta}{\beta+\delta_H}\right)^\beta \left(\frac{\delta_H}{\beta+\delta_H}\right)^{\delta_H} - \left(1 - s_2^R\left[1-\frac{\omega s_2^R}{4}\right]\right)^\beta \left(s_2^R\left[1-\frac{\omega s_2^R}{2}\right]\right)^{\delta_H} \\
&- \theta + a\left|s_2^R\left[1-\frac{\omega s_2^R}{2}\right] - s_1^R\right|^{expc}\Bigg]\Bigg).
\end{aligned}
$$

Note on Existence

The options of the political parties D and R are as follows. D chooses $\mathbf{s^D} = (s_1^D, s_2^D) \in [0,1]^2$, and R chooses $\mathbf{s^R} = (s_1^R, s_2^R) \in [0,1]^2$. The payoff functions are $U^D(\mathbf{s^D}, \mathbf{s^R})$ and $U^R(\mathbf{s^R}, \mathbf{s^D})$. Within the value of parameters in our calibration exercise, the neighborhood, and the comparative statics, the best response functions are continuous. This is locally true, but not necessarily for all the possible values of the multiple parameters in our complex utility functions. The best response functions map as follows,

$$BR^D : [0,1]^4 \to [0,1]^2 \text{ and } BR^R : [0,1]^4 \to [0,1]^2.$$

Define the joint best-response function $B : [0,1]^4 \to [0,1]^4$ by

$$B(\mathbf{s^D}, \mathbf{s^R}) = (BR^D(\mathbf{s^R}, \mathbf{s^D}), BR^R(\mathbf{s^D}, \mathbf{s^R})).$$

If BR^D and BR^R are continuous, then the joint best-response function B is also continuous. Moreover, since both parties' strategies belong to $[0,1]^2$, the joint strategy space is

$$[0,1]^2 \times [0,1]^2 = [0,1]^4.$$

This set is compact, convex, and non-empty. Also, B maps this set into itself. Therefore, by Brouwer's Fixed Point Theorem, there exists a point $(\mathbf{s^{D*}}, \mathbf{s^{R*}}) \in [0,1]^4$ such that $(\mathbf{s^{D*}}, \mathbf{s^{R*}}) = B(\mathbf{s^{D*}}, \mathbf{s^{R*}})$. Equivalently, $\mathbf{s^{D*}} = BR^D(\mathbf{s^{R*}}, \mathbf{s^{D*}})$ and $\mathbf{s^{R*}} = BR^R(\mathbf{s^{D*}}, \mathbf{s^{R*}})$. Hence, $(\mathbf{s^{D*}}, \mathbf{s^{R*}})$ is a Nash equilibrium.

Furthermore, the game is symmetric in that

$$U^D(\bar{\mathbf{s}}, \mathbf{s}) = U^R(\mathbf{s}, \bar{\mathbf{s}})$$

for all strategy profiles. If we only pay attention to the diagonals $\mathbf{s^D} = \mathbf{s^R} = \mathbf{s}$ and if, as in the calibration, the best responses are unique, whenever both parties face the same strategy profile $\mathbf{s}$, they share the same best response. Define the common best-response function

$$BR(\mathbf{s}) = BR^D(\mathbf{s}, \mathbf{s}) = BR^R(\mathbf{s}, \mathbf{s}).$$

When the original best-response functions are continuous, BR is also continuous. Moreover, $BR : [0,1]^2 \to [0,1]^2$.

The set $[0,1]^2$ is non-empty, compact, and convex. Therefore, if the Best Responses are continuous, by Brouwer's Fixed Point Theorem, there exists a symmetric NE.

A2 Additional numerical comparative statics

This section reports the numerical values underlying the comparative-static figures in Section 4. Each exercise varies the indicated parameter while holding the others at the benchmark settings described in the main text. NT denotes no terrorist attack.

A2.1 Security demand without bureaucratic inefficiency

Table A1 below accompanies Figure 5, with $\omega = 0$.

δ_L	Nash Proposed Expenditure	Nash Implemented Expenditure	Nash Probability NT	Optimal Implemented Expenditure	Optimal Probability NT	Nash Slack	Optimal Slack
0.010	0.01097	0.01496	0.19238	0.02228	0.28644	0	0
0.020	0.02190	0.02720	0.34968	0.04898	0.62977	0	0
0.030	0.03256	0.03807	0.48941	0.07563	0.97244	0	0
0.040	0.04289	0.04806	0.61790	0.07778	1.00000	0	0
0.050	0.05288	0.05742	0.73829	0.07778	1.00000	0	0
0.060	0.06251	0.06629	0.85233	0.07778	1.00000	0	0
0.070	0.07182	0.07476	0.96117	0.07778	1.00000	0	0
0.077	0.07750	0.07778	1.00000	0.07778	1.00000	0	0
0.080	0.07979	0.07979	1.00000	0.07979	1.00000	0	0
0.090	0.08888	0.08888	1.00000	0.08888	1.00000	0	0
0.100	0.09778	0.09778	1.00000	0.09778	1.00000	0	0
0.110	0.10652	0.10652	1.00000	0.10652	1.00000	0	0
0.120	0.11509	0.11509	1.00000	0.11509	1.00000	0	0
0.130	0.12350	0.12350	1.00000	0.12350	1.00000	0	0
0.140	0.13175	0.13175	1.00000	0.13175	1.00000	0	0
0.150	0.13984	0.13984	1.00000	0.13985	1.00000	0	0

Table A1: Choices by parties in a Nash equilibrium, and choices by a Central Planner, as a function of δ_L, for $\omega = 0$. Values rounded to 5 decimals.

A2.2 Variation in bureaucratic inefficiency

Table A2 below accompanies Figure 6, with $\delta_L = 0.077371$.

ω	Nash Proposed Expenditure	Nash Implemented Expenditure	Nash Probability NT	Optimal Implemented Expenditure	Optimal Probability NT	Nash Slack	Optimal Slack
0	0.07750	0.07778	1.00000	0.07778	1.00000	0.00000	0.00000
1	0.07750	0.07778	1.00000	0.07778	1.00000	0.00164	0.00164
2	0.07750	0.07778	0.99999	0.07778	1.00000	0.00361	0.00361
3	0.07686	0.07581	0.97466	0.07778	1.00000	0.00571	0.00606
4	0.07604	0.07325	0.94174	0.07778	1.00000	0.00795	0.00928
5	0.07500	0.07000	0.90000	0.07778	1.00000	0.01022	0.01397
6	0.07373	0.06599	0.84845	0.07757	0.99739	0.01232	0.02263
7	0.07230	0.06137	0.78905	0.06788	0.87270	0.01393	0.02156
8	0.07082	0.05651	0.72659	0.06016	0.77345	0.01489	0.02032
9	0.06941	0.05182	0.66632	0.05393	0.69336	0.01524	0.01908
10	0.06814	0.04755	0.61139	0.04882	0.62770	0.01510	0.01791
11	0.06702	0.04377	0.56275	0.04457	0.57307	0.01481	0.01683
12	0.06604	0.04046	0.52017	0.04099	0.52699	0.01434	0.01585
13	0.06519	0.03756	0.48296	0.03793	0.48764	0.01380	0.01496
14	0.06444	0.03503	0.45036	0.03529	0.45367	0.01325	0.01416
15	0.06378	0.03280	0.42167	0.03298	0.42408	0.01270	0.01343
16	0.06319	0.03082	0.39627	0.03096	0.39807	0.01218	0.01277
17	0.06266	0.02906	0.37367	0.02917	0.37504	0.01168	0.01216
18	0.06218	0.02749	0.35345	0.02757	0.35452	0.01121	0.01161
19	0.06175	0.02608	0.33526	0.02614	0.33610	0.01076	0.01110
20	0.06136	0.02480	0.31882	0.02485	0.31950	0.01034	0.01064
21	0.06099	0.02364	0.30390	0.02368	0.30445	0.00996	0.01021
22	0.06066	0.02258	0.29029	0.02261	0.29074	0.00960	0.00982
23	0.06035	0.02161	0.27784	0.02164	0.27822	0.00926	0.00944
24	0.06006	0.02072	0.26641	0.02074	0.26672	0.00894	0.00910
25	0.05980	0.01990	0.25587	0.01992	0.25613	0.00864	0.00878
26	0.05955	0.01914	0.24613	0.01916	0.24635	0.00836	0.00849

Table A2: Choices by parties in a Nash equilibrium, and choices by a Central Planner, as a function of ω. δ_L is at the benchmark 0.077371. Values rounded to 5 decimals.

A3 Bureaucratic inefficiency and credibility

In the model, credibility enters the utility of the parties by subtracting the following term from the incumbency advantage parameter θ:

$$a \left| s_2^D \left(1 - \frac{\omega s_2^D}{2}\right) - s_1^D \right|^{expc}, \tag{A6}$$

where $a > 0$ and $expc \geq 1$.

To see the effect of credibility, note first that due to the inefficiency of SA, the cost of fulfilling the electoral promise s_1^D increases with the level of s_1^D. Figure A1 shows that the appropriation to the SA, s_2^D, required to fully meet the promised level s_1^D increases more than proportionally as s_1^D rises. Thus, for $\omega > 0$, the (blue) curve

$$s_1^D = s_2^D \left(1 - \frac{\omega s_2^D}{2}\right)$$

bends downward. By contrast, when $\omega = 0$, the relationship is simply the 45° line.

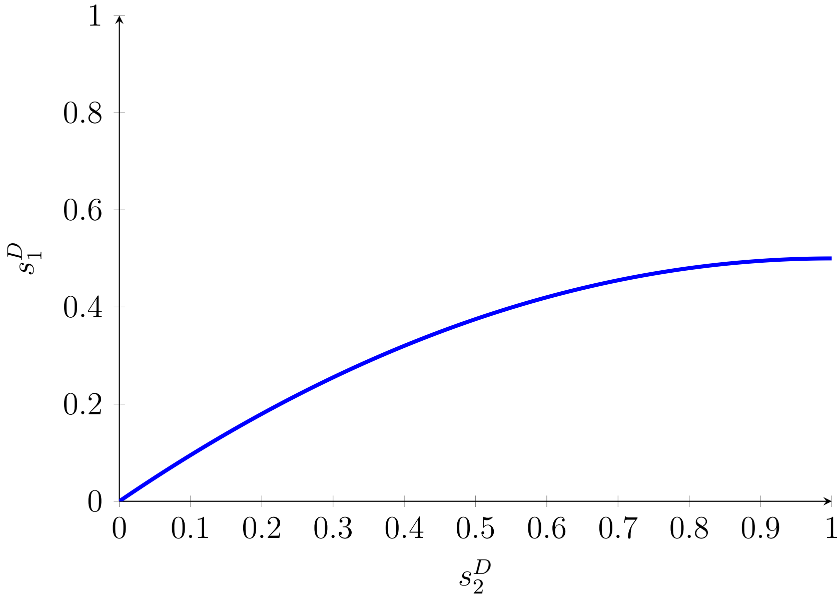

Figure A1: Partition of the unit square by the curve $s_1^D = s_2^D(1 - \frac{\omega s_2^D}{2})$. Above the curve $s_1^D > s_2^D(1 - \frac{\omega s_2^D}{2})$; below $s_1^D < s_2^D(1 - \frac{\omega s_2^D}{2})$.

Figure A2 illustrates the effect of ω on credibility in terms of expression A6. In the figure, we set $a = 1$, $expc = 1$, and fix the electoral promise at $s_1^D = 0.4$. For small values of ω, for example when $\omega = 0$ or $\omega = 1$, it is possible to choose an appropriation level s_2^D that fully delivers the promised level s_1^D. By contrast, for larger values of ω, such as the benchmark value $\omega = 5$ or above, fully delivering the promise becomes impossible. Moreover, the credibility penalty becomes a convex function of s_2^D and increases toward the supremum of the credibility function as s_2^D rises.

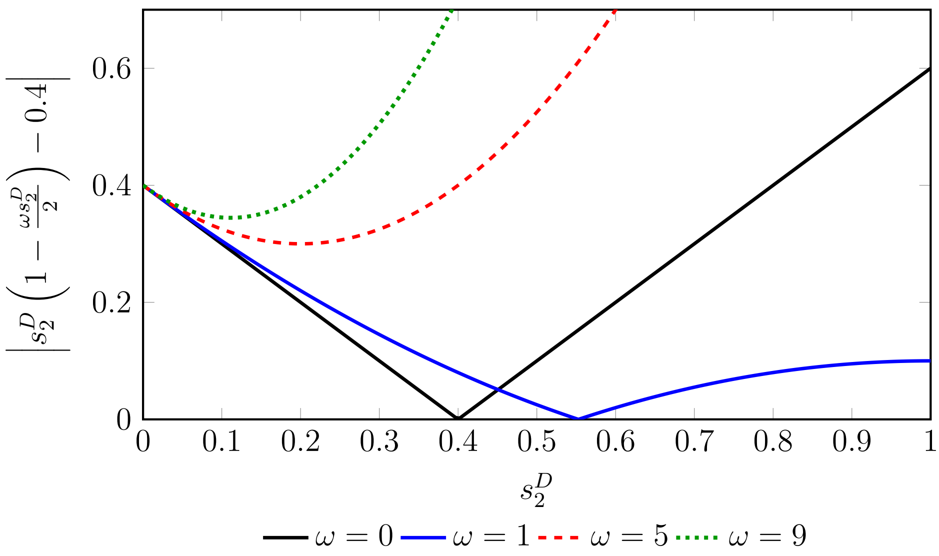

Figure A2: Absolute value function $a \left| s_2^D \left(1 - \frac{\omega s_2^D}{2}\right) - s_1^D \right|^{expc}$ for different values of ω. Illustration with $a = 1$, $expc = 1$, and $s_1^D = 0.4$.

A4 Data and descriptive statistics

This section describes the data and variables used in Figure 2 and the descriptive figures in Appendix A5. The analysis combines country-year observations on terrorism, military expenditure, political rhetoric and institutional quality. These comparisons provide descriptive context and motivation for the model rather than causal estimates of its mechanisms. Data availability and the observations retained vary across figures.

Sample construction and coverage The study period is 1990–2020. The analysis combines the union of available SIPRI, WGI and MPD country-year records with population and incident data, rather than restricting all outcomes to a common complete-case sample. Each plotted outcome requires its own nonmissing value and valid terrorism exposure. Territories are retained where the sources permit, and no exposure-tail trimming or winsorization is applied. The governance series begin in 1996 and lack observations in 1997, 1999 and 2001.

Incident counts are reconstructed from GTD event identifiers. The year 1993 is treated as missing because GTD documents incomplete incident coverage for that year. For eligible country-years with a recognized GTD country name and no recorded event, the count is set to zero conditional on coverage; this denotes no recorded incident, not proof that no attack occurred. Unresolved country-name or coverage cases and missing or nonpositive population are not assigned zero exposure. Conservative first-full-year restrictions are applied to successor states to avoid treating pre-independence or transition-year mismatches as zero incidents.

Table A3 summarizes the nine variables used in these figures.

Table A3: Variable definitions and data sources.

Variable	Definition and interpretation	Source
Log(1 + terrorist incidents per million)	The natural logarithm of one plus the annual number of recorded terrorist incidents in a country divided by its population and multiplied by one million. Adding one makes the transformation defined at zero incidents.	

Continued on next page

Table A3 (continued)

Variable	Definition and interpretation	Source
Military spending (% of GDP)	Military expenditure divided by gross domestic product and multiplied by 100. SIPRI's definition covers current and capital spending on armed forces, defense agencies, military activities and qualifying paramilitary forces, including personnel, operations, procurement and military research and development. This is a GDP share, distinct from the government-budget shares used in the theoretical model.	SIPRI.
Security rhetoric (net)	The vote-weighted difference between the manifesto percentages assigned to Military: Positive (`per104`) and Military: Negative (`per105`). Positive values indicate more favorable than unfavorable military content.	MPD.
Left–right ideology	The vote-weighted Manifesto Project `rile` index: the sum of right-coded manifesto percentages minus the sum of left-coded percentages. Higher values indicate a more right-wing position. The ideology figure uses its standardized form.	MPD.
Pro-military rhetoric	The vote-weighted percentage of manifesto quasi-sentences assigned to Military: Positive (`per104`); this measures favorable military emphasis, not a net difference.	MPD.
Peace rhetoric	The vote-weighted percentage of manifesto quasi-sentences assigned to Peace (`per106`), measuring peace-oriented content.	MPD.
Government inefficiency (standardized)	The negative of the standardized WGI Government Effectiveness estimate. The underlying indicator concerns public services, civil-service quality and independence, and policy formulation, implementation and credibility. Higher values indicate weaker government effectiveness.	WGI.
Weak control of corruption (standardized)	The negative of the standardized WGI Control of Corruption estimate, which concerns the use of public power for private gain, including corruption and state capture. Higher values indicate weaker control of corruption.	WGI.
Political instability (standardized)	The negative of the standardized WGI Political Stability estimate. Higher reversed values indicate greater perceived instability.	WGI.

Sources: GTD = Global Terrorism Database; SIPRI = Stockholm International Peace

Research Institute; MPD = Manifesto Project Database; WGI = Worldwide Governance Indicators.

Aggregation and transformations Manifesto measures are aggregated across parties within each election using positive vote-share weights, with parliamentary vote shares preferred and presidential shares used as a fallback. Weights are normalized separately over nonmissing values for each measure. Net rhetoric is calculated from paired military-positive and military-negative values before aggregation; missing categories are not converted to zero. Where several elections occur in a country-year, the latest usable measure is retained separately for each variable. No values are carried forward into non-election years. These aggregates describe the observed electoral party mix, not a single party or the government alone.

Standardization subtracts the country-year sample mean and divides by the sample standard deviation within the relevant series' available analysis sample, before binning or smoothing. WGI scores are then multiplied by minus one.

Figure 2 presents original-unit bin means in panel (a) and means of standardized country-year outcomes in panel (b). Figure A7 uses original-unit rhetoric measures. Each series is binned independently, with a separate zero-exposure bin and up to 12 quantile bins for positive exposure. Bin coordinates are the series-specific mean exposure and mean outcome. Figure A6 uses LOWESS with bandwidth 1. Each institutional-quality figure uses 15 quantile bins of its governance measure.

Descriptive statistics Table A4 reports unbinned statistics for the analysis samples.

Table A4: Descriptive statistics for the empirical analysis samples.

Variable	N	Countries	Mean	SD	Min.	Max.
Log(1 + incidents per million)	5,005	183	0.361	0.685	0.000	4.740
Military spending (% GDP)	4,266	162	2.198	2.835	0.000	117.350
Security rhetoric (net)	457	63	1.339	1.746	-1.690	16.109
Left–right ideology	457	63	-3.340	11.945	-36.620	39.861
Pro-military rhetoric	457	63	1.629	1.770	0.000	16.611
Peace rhetoric	457	63	0.493	0.881	0.000	9.714
Government inefficiency	3,929	181	0.000	1.000	-2.363	2.348
Weak control of corruption	3,939	181	0.000	1.000	-2.512	1.771
Political instability	3,944	183	0.000	1.000	-1.872	2.950

Notes: N counts country-year observations before binning or smoothing; units are countries/territories. The study period is 1990–2020, excluding 1993. WGI observations cover 1996–2020 with gaps in 1997, 1999 and 2001. Governance measures are reversed standardized scores; ideology and rhetoric are reported in original units. Military spending is a percentage of GDP. The terrorism row uses the union of plotted-outcome samples; other rows use series-specific samples.

A5 Descriptive evidence

A5.1 Institutional quality and terrorism

This section collects descriptive evidence on institutional quality and terrorism exposure using three World Governance Indicators (WGI) measures: government effectiveness, control of corruption, and political stability. We reverse the standardized indices so that higher values indicate greater government inefficiency, weaker control of corruption, or greater political instability. These measures are not direct empirical counterparts of ω. In the model, ω specifically captures the relative bureaucratic inefficiency of security provision compared with other government activities. The figures instead provide descriptive evidence on the broader institutional environments in which security-sector inefficiency may operate. These patterns motivate examining how institutional weaknesses that hinder policy implementation may also limit effective security provision.

Government inefficiency Figure A3 relates terrorism exposure to reversed standardized government effectiveness. Terrorism exposure is relatively low across much of the distribution but rises among countries with high government inefficiency. The association between government inefficiency and terrorism exposure motivates our examination of how inefficient implementation can weaken security provision.

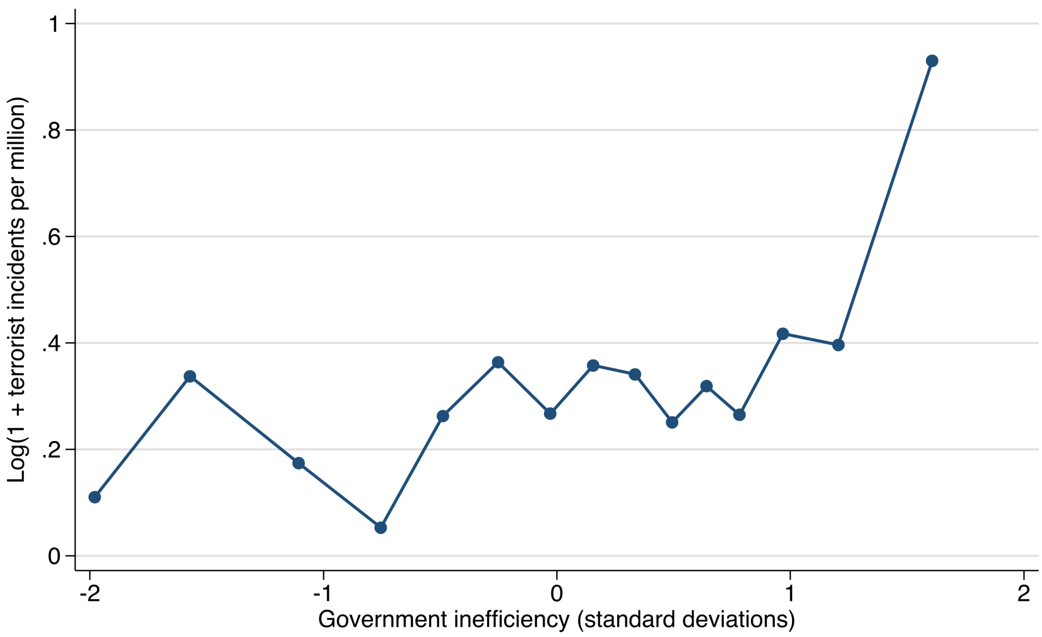

Figure A3: **Terrorism exposure and government inefficiency.** The figure reports means of $\log(1 + \text{terrorist incidents per million})$ across 15 quantile bins of reversed standardized WGI Government Effectiveness. Higher values indicate greater government inefficiency. The sample contains 3,929 country-years from 181 countries/territories over 1996–2020, with no observations in 1997, 1999 or 2001. Sources: GTD and WGI.

Control of corruption and political stability Figure A4 presents the relationship between weak control of corruption and terrorism exposure. Figure A5 presents the corresponding relationship for political instability.

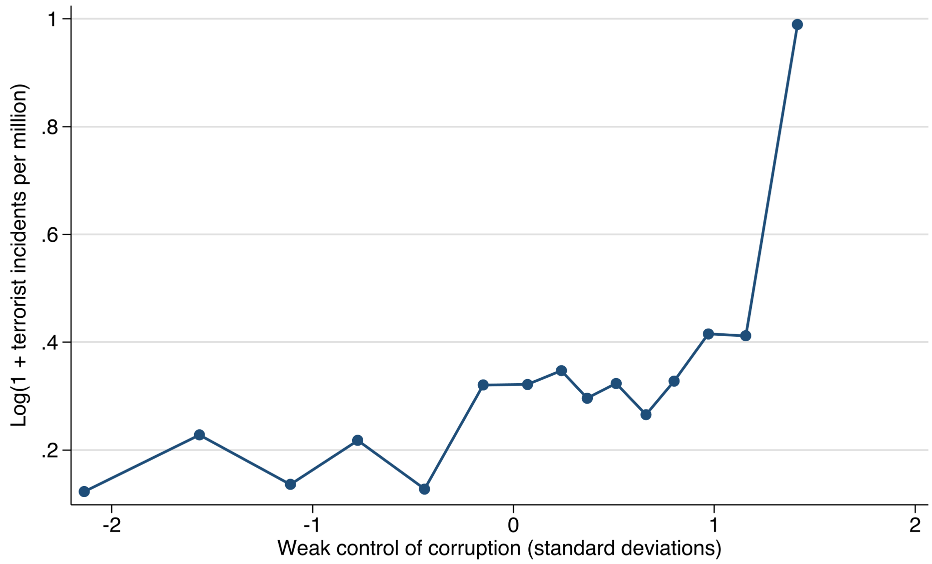

Figure A4: **Weak control of corruption and terrorism.** The figure reports means of log(1 + terrorist incidents per million) across 15 quantile bins of reversed standardized WGI Control of Corruption. Higher values indicate weaker control of corruption. The sample contains 3,939 country-years from 181 countries/territories over 1996–2020, excluding 1997, 1999 and 2001. Sources: GTD and WGI. The relationship is descriptive.

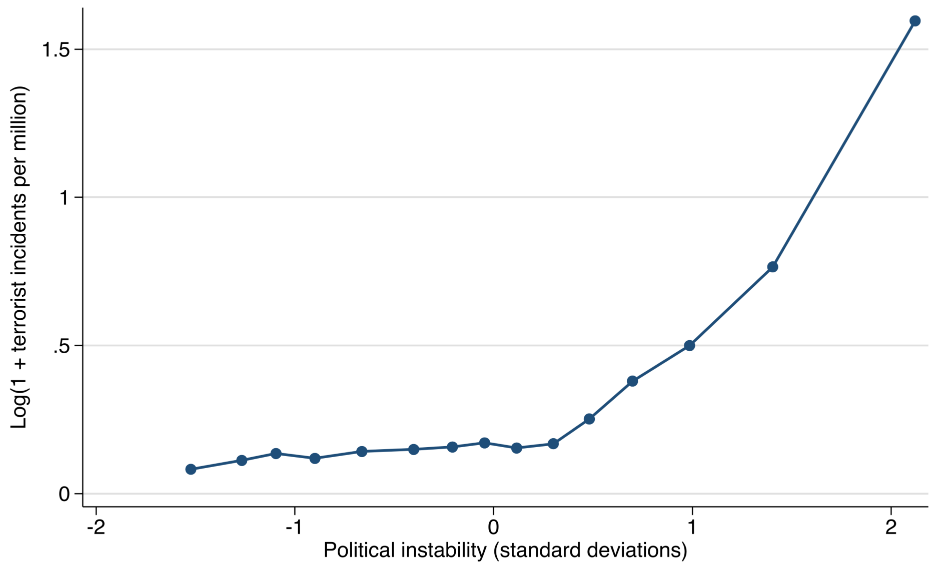

Figure A5: **Political instability and terrorism.** The figure reports means of log(1 + terrorist incidents per million) across 15 quantile bins of reversed standardized WGI Political Stability and Absence of Violence/Terrorism. Higher values indicate greater instability. The sample contains 3,944 country-years from 183 countries/territories over 1996–2020, excluding 1997, 1999 and 2001. Sources: GTD and WGI.

A5.2 Ideology and security rhetoric

The following descriptive comparisons supplement Figure 2. They motivate the interpretation of security preferences but are not intended to identify the model parameters or provide a causal test of the political mechanism.

Ideology comparison Figure A6 compares broad left–right ideology with net security rhetoric. In the sample, both standardized measures increase at high terrorism exposure, while their patterns differ at lower exposure. The figure therefore does not rule out an association between terrorism exposure and ideological position. The model's focus on a shift in the weight placed on security, from δ_L to δ_H, is a modeling choice rather than a conclusion established by this comparison.

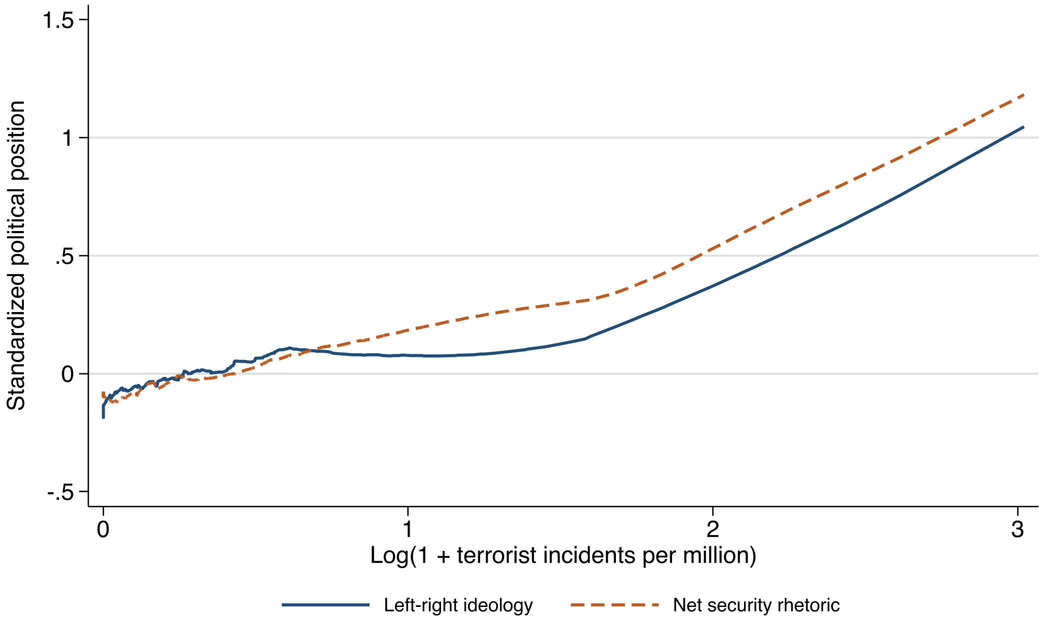

Figure A6: **Terrorism, ideology, and security rhetoric.** LOWESS curves (bandwidth 1) relate standardized left-right ideology and net security rhetoric to log(1 + terrorist incidents per million). Higher ideology values indicate a more right-wing position. Each series uses 457 country-years from 63 countries over 1990-2020, excluding 1993, with no exposure-tail trimming. Outcomes are standardized before smoothing within their available samples. Sources: GTD and MPD.

Composition of security rhetoric Figure A7 provides descriptive context for security preferences. Net security rhetoric and pro-military rhetoric are generally higher at greater terrorism exposure, although the relationships are not monotonic. Peace rhetoric also rises toward the upper end of the exposure distribution, but remains lower in level. Net rhetoric subtracts military-negative content from military-positive content; peace rhetoric is a separate category and is not its negative component. The figure motivates considering security-related political communication and the preference shift $\delta_H > \delta_L$.

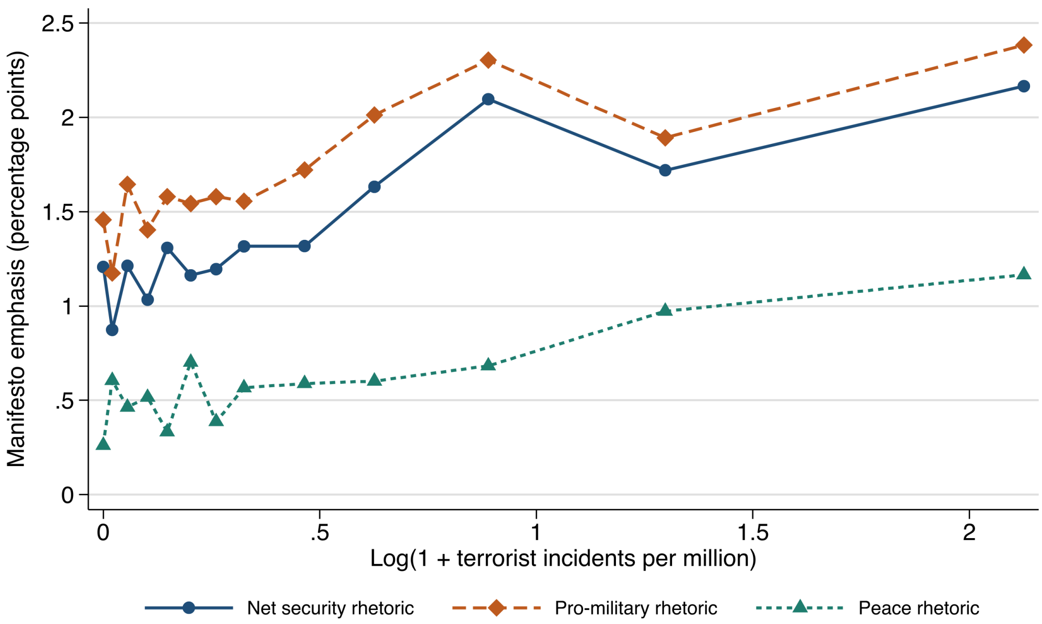

Figure A7: **Terrorism and the composition of security rhetoric.** The figure reports bin means of net security rhetoric (military-positive minus military-negative manifesto content), pro-military rhetoric and peace rhetoric against mean log(1 + terrorist incidents per million). Each series uses 457 country-years from 63 countries over 1990–2020, excluding 1993. A separate zero-exposure bin and up to 12 positive-exposure quantile bins are constructed for each series without tail trimming. Rhetoric measures are vote-weighted manifesto percentages (percentage-point differences for the net measure). Sources: GTD and MPD. See Appendix A4 for sample construction.

A6 Model notation

Table A5 collects the notation used in the model. Lower-case expenditure variables denote budget shares; upper-case expenditure variables denote amounts. Party or role identifiers appear as superscripts and election indices as subscripts in the main policy notation. The incumbent's s_2^I denotes an appropriation, whereas the challenger's s_2^{-I} denotes a platform. Benchmark parameter values are reported in Section 3.

Table A5: Model notation and interpretation.

Symbol	Meaning
Agents, indices and states	
$D, R;\ i$	Political parties; $i \in \{D, R\}$ indexes a party.
$I;\ -I;\ C$	Incumbent, opposing party, and challenger, respectively.
t	Election index, $t \in \{1, 2\}$.
$j \in \{L, H\}$	Low and high security-preference states, respectively.
T, NT	Terrorist attack and no terrorist attack, respectively.
SA	Security Agency.
Policy choices, implementation and agency outcomes	
GB	Total government budget.
$S;\ S_I$	Security expenditure in the voter's utility function, and the amount appropriated to the agency, respectively; $S_I = s_2^I GB$.
$S_e;\ c$	Agency effort and cost per unit of effort; cS_e is implemented security expenditure.
s_1^i	Party i's first-election security spending promise, expressed as a share of the government budget.
s_2^I	Incumbent's appropriation to the agency, expressed as a share of the government budget.
s_2^{-I}	Challenger's second-election spending platform, chosen after observing whether an attack occurs.
$\mathbf{s}^i;\ s_1, s_2$	Party strategy vector $\mathbf{s}^i = (s_1^i, s_2^i)$; in a symmetric equilibrium, s_1 and s_2 denote the common promise and incumbent appropriation, respectively.
Preferences and institutional parameters	
$U;\ CL$	Voter utility and civil liberties, respectively.
β	Preference weight on government expenditure other than security and counter-terrorism.
$\gamma_j;\ \phi$	Components of the security-preference weight: γ_j concerns terrorism risk reduction and ϕ captures civil-liberty concerns.
δ_j	Security-preference weight, $\delta_j = \gamma_j/\phi$.
δ_L, δ_H	Baseline and post-attack security-preference weights, with $\delta_H > \delta_L$.

Continued on next page

Table A5 (continued)

Symbol	Meaning
ε	Increase in security preference following an attack: $\delta_H = \delta_L + \varepsilon$.
ω	Bureaucratic inefficiency in security provision relative to other government activities.
ν	Coefficient linking implemented security spending to the modeled probability of no attack.
a	Positive weight on the credibility penalty for a discrepancy between promised and implemented expenditure.
$expc$	Exponent on the absolute promise–implementation discrepancy in the credibility penalty.
θ	Exogenous incumbency advantage entering second-election probabilities.
Ψ	Positive parameter governing electoral sensitivity to policy differences.
χ_t	Independent popularity shock in election t, uniform on $(-1/(2\Psi), 1/(2\Psi))$.
B	Value of a second-election victory relative to a first-election victory.
Probabilities, payoffs and equilibrium	
$P(T)$	Probability of a terrorist attack in the general voter-utility representation.
$P(NT)$	Modeled probability of no terrorist attack.
$P_1(i)$	Probability that party i wins the first election.
$P_2^{NT}(i;\cdot),\ P_2^{T}(i;\cdot)$	State-weighted second-election winning probabilities, conditional on the incumbent/challenger role indicated after the semicolon. Each includes the probability of its attack state.
EU_i	Party i's expected payoff across the two elections.
$U^i;\ BR^i$	Party i's payoff function and best-response function in the existence discussion.
$BR;\ B(\cdot)$	Common best-response function on symmetric profiles and joint best-response map, respectively. The latter is distinct from the scalar election weight B.
$*$; NE	An optimum or equilibrium value, according to context; Nash equilibrium.