Polit Behav (2015) 37:723–748
DOI 10.1007/s11109-014-9289-x

ORIGINAL PAPER

# Citizen Information, Electoral Incentives, and Provision of Counter-Terrorism: An Experimental Approach

**Andrew W. Bausch · Thomas Zeitzoff**

Received: 8 January 2014 / Accepted: 9 September 2014 / Published online: 23 September 2014


**Abstract** How does incomplete information about counter-terror provisions influence the strategic interaction between a government, terrorist groups, and the citizenry? We investigate this research question using a laboratory experiment and present two key findings. (1) Public counter-terror spending leads citizens to overly frequent "protected" targets such that it makes them easier targets for terrorists. (2) Additionally, we show that citizens over-estimate government counter-terror spending when they are unable to observe it. These findings suggest that asymmetric information and the small probability of a successful terrorist attack may lead to the inefficient provision of counter-terror. We also connect the findings to the larger literature on the principal-agent relationship between citizens and elected officials.

**Keywords** Terrorism · Principal-agent · Experiment

## Introduction

> "(Former U.S.) Homeland Security Chief Tom Ridge made that critical leap from 'be afraid' to 'be very afraid', raising the terrorist threat level to orange for financial sectors in New York, Washington, D.C., and northern New Jersey... Ridge's announcement comes amidst reports he will step down as

**Electronic supplementary material** The online version of this article (doi:10.1007/s11109-014-9289-x) contains supplementary material, which is available to authorized users.

A. W. Bausch (✉)
Department of Social and Decision Sciences, Carnegie Mellon University, Pittsburgh, USA
e-mail: bausch@nyu.edu

T. Zeitzoff
American University, Washington, DC, USA
e-mail: zeitzoff@american.edu

> head of Homeland Security after the election. Ridge himself has refused to comment on the story, though colleagues say he has often expressed a desire to spend more time at home, scaring his family."-John Stewart (August 2004)

The above quote is meant to playfully mock the now defunct U.S. Department of Homeland Security (DHS) color-coded Advisory System. Yet critics have argued that there is an element of truth hidden in Stewart's lampoon (Shapiro and Cohen 2007). They contend that the Transportation Security Administration (TSA) is "searching grandmothers" while ignoring terrorists (Lipton 2010). Other analysts argue that politics and the appearance of security take precedent over actual security (de Rugy 2006). This is not surprising given that citizens can observe certain kinds of security measures (airport screenings) and not others (intelligence gathering).

Counter-terror provision is just one of the numerous instances in which voters must evaluate politicians based on incomplete information about what politicians do and to what extent their policies influence outcomes. Many studies of retrospective voting have found that voters use a heuristic based on outcomes to sanction or reward politicians, even for events that they have little control over (Achen and Bartels 2004; Healy et al. 2010). A separate literature on counter-terrorism shows: (1) politicians may sub-optimally provide counter-terrorism by pandering to voters and over-investing in visible protection (Bueno de Mesquita 2007), and (2) terrorists behave strategically by moving away from targeting heavily defended sites, instead attacking weaker sites (Enders and Sandler 1993; Powell 2007a). The microfoundations of this research is the rational, strategic response of terrorists, politicians, and voters. However, the retrospective voting literature suggests that citizens use heuristics (i.e., are not fully strategic) to support or sanction politicians. A key question then is how do the strategic behavior of terrorists and politicians interact with the bounded rationality of citizen voters?

We examine how the partially observable nature of counter-terror influences strategic behavior between a representative government, citizen, and terrorist using a laboratory experiment.[1] A laboratory experiment is an ideal way to directly manipulate the level of incomplete information and observe the dynamics between the three main actors: citizen, terrorist, and government. These three actors are represented in our experiment. The government allocates a budget toward hardening potential terrorist targets. This hardening may or may not be observed by the citizen and terrorist. The citizen then tries to avoid a terrorist attack, and the terrorist tries to successfully complete an attack. After observing whether or not the terrorist attack was successful, the citizen can choose to reward the government.

We have two principal findings. (1) Public counter-terror spending makes citizens sub-optimally pool on publicly hardened targets, making citizens 'easier' targets for the terrorists. (2) Citizens behave as if the government invests more heavily in counter-terror when they do not have information about government investment (relative to when they do), and reward the government with higher

[1] From a methodological standpoint, it would be ideal to directly observe and/or manipulate citizens' and terrorists' level of incomplete information (from 100 % incomplete to 100 % complete) and observe subsequent counter-terror dynamics. Yet from an ethical standpoint (thankfully), this is not feasible, nor practical. Thus, we accept the artificiality of a laboratory.

bonuses. These findings paradoxically suggest that citizens desire to monitor politicians' counter-terror provisions (by making hardening public), may actually make them less safe. Additionally, politicians have a strong incentive to obfuscate security spending, by claiming credit for terrorism-free periods even if they have done little.

The paper is structured as follows. In the section titled "Information, Elections, and Strategic Interaction," we discuss the extant theory and literature on counter-terrorism and information asymmetries between citizens and politicians. The "Experiment" section outlines the experiment while the "Predictions" section generates our experimental predictions. The "Results" section discusses the results before the "Conclusion" section concludes and presents policy recommendations.

## Information, Elections, and Strategic Interaction

Our experiment examines two main issues related to counter-terrorism: (1) the strategic behavior of counter-terror provisions and (2) the principal-agent relationship between citizens, who have incomplete information about the provision of counter-terror, and a government that provides the counter-terror. We begin by addressing how counter-terrorist activities affect the strategic interaction between the citizens the government seeks to protect, and terrorists. Several game-theoretical papers have addressed the related dynamics between a government attempting to protect sites and the decision of a terrorist about when, where, and with how much effort to attack.[2] Many of these studies argue that there are two types of counter-terrorism: defensive policies that protect specific targets by reducing the likelihood of successful attacks and proactive, offensive measures, that go after the terrorists directly through attacks to kill or capture group members, to seize resources, or eliminate safe havens (Sandler and Siqueira 2009). This paper focuses on the former.

A number of empirical and formal studies show that terrorists will switch tactics or targets in response to counter-terrorism. While greater governmental efforts to deter a specific type of attack (or an attack at a given site) lowers the likelihood of attack with that tactic or in that place, it increases the likelihood of attacks with other tactics or in other locations (Landes 1978; Enders and Sandler 1993, 2002; Sandler and Arce 2003). Thus, counter-terrorism does not completely deter attacks, but changes the terrorists' calculus of where or by what means an attack is most likely to succeed. In an international context, several models show that individual governments are predisposed to overspend on defensive measures to avoid becoming the most vulnerable target because proactive expenditures produce an incentive for allies to free ride (Arce and Sandler 2005; Sandler and Siqueira 2006; Cadigan and Schmitt 2010). These studies show how states inadvertently deflect terrorism to their allies through defensive counter-terrorism. In a domestic context, Powell (2007b) shows the government uses its defense budget optimally against a

[2] For a complete review of the game-theoretical approaches to terrorism, see Sandler and Arce (2003) and Sandler and Siqueira (2009).

fully strategic actor when minimizing the terrorists' maximum payoff until the terrorist's payoff is equal across all defended sites.

However, if the defending government has private information about the relative vulnerabilities of the sites they are trying to protect, this creates a trade-off for the government. The government may desire to invest in defending the most vulnerable site, but doing so may signal to the terrorist that the site is, in fact, vulnerable, thus increasing the probability of an attack there. Powell (2007a) finds that secrecy concerns outweigh vulnerability concerns when more vulnerable sites are also harder to protect on the margin. In such cases, the government is better off allocating defense expenditures so as to avoid signaling its vulnerability to the attacker, despite seemingly doing too little to secure the vulnerable site.

While these models address the interplay between a government's defense strategy and the actions of the terrorist, they neglect the actions and choices of citizens and how they may be affected by governmental action.[3] Jindapon and Neilson (2009) find that if the median voter is risk-neutral, terrorists set the severity of attacks without regard for governmental activities while the government invests in the highest level of counter-terrorism. Meanwhile, risk-aversion on behalf of the citizenry induces terrorists to attempt more severe but less frequent attacks. Getmansky (2011a) shows formally that geographic areas that support the government become more attractive to terrorists, thus receiving more counter-terror investment, which in turn decreases the likelihood of attack there. Empirically, Getmansky (2011b) shows that areas where voters are most likely to switch to the opposition party are more likely to receive protection, while strong supporters are less likely to be protected because a successful attack is unlikely to change their vote.

What these models fail to address is the interplay between a citizen's reaction to counter-terrorism and a terrorist's updated strategy. For example, a citizen may be making a choice about a mode of transportation to work. She may ordinarily drive her car to her destination, though the route would lead her over a lightly protected bridge. She could also take public transportation such as a subway to her destination. After hearing rumors that an attack may be imminent and that terrorists may be targeting bridges or subways, the citizen may be unsure how to react. These rumors do not provide enough information for the citizens to decide between taking her car or taking the subway. We may assume that she takes her car, as usual, or that the mere threat of attack persuades her to stay home, despite such a decision being individually and socially suboptimal.[4]

Now imagine the same citizen hears rumors about a potential attack on a bridge or subway and observes the government increase counter-terror measures in and around subways. On one hand, she may believe that government is adding protection to the subways because an attack is most likely there and, thus, chooses to

[3] Looking more closely at partisanship, Berrebi and Klor (2006) finds that terrorists react strategically by attacking left-wing governments more than right-wing. Both Berrebi and Klor (2008) and Getmansky and Zeitzoff (2014) show that terrorism shifts voters to the right.

[4] With regards to the plausibility of citizens switching transportation methods, Blalock et al. (2009) finds an increase in driving fatalities after 9/11 that can be attributed to travelers substituting road trips in place of air travel. Furthermore, Cox et al. (2011) find a reduction in subway and bus trips after the 2005 London bombings.

avoid the subway. On the other hand, she may believe that counter-terrorism is effective and will prevent a subway attack or shift it to another location. For instance, during the Second Intifadah, following several high profile bombings against public buses, Israelis shifted their transportation habits to take buses (Kahneman 2011, p. 322). Assuming the terrorist can switch relatively easily between attacking the bridge and the subway, the formal models discussed above suggest that increased counter-terror activities directed at subways would prompt the terrorist to target the bridge.[5] However, if the citizen follows this logic, she will surely ride the subway to avoid the attack on the bridge.

Thus, if terrorists select targets based on publicly-observed counter-terrorism efforts, focusing attacks on less protected targets, citizens, observing increased protection on a site should favor those sites over more exposed targets. Depending on the goals of the terrorist, particularly if the terrorist places a higher value on human casualties than destroying infrastructure, the terrorist may now be willing to attack the more protected site. In effect, the lower probability of success for the terrorist is mitigated by the higher payoff of a successful attack. In contrast to Powell (2007a), where publicly observable counter-terrorism can signal weakness, in this situation publicly observable counter-terrorism activities increases the chance of an attack because of an increase in the value of the target to the terrorist as more citizens make use of the protected target. If citizens behave heuristically, and always go to the most protected targets, then a terrorist can exploit their feeling of safety by attacking the these targets—even if the protected target has a lower probability of success. Public counter-terrorism, in this case, serves as a coordinating device between the terrorist and the citizens. It allows the terrorist to predict how the value of the target changes in response to defensive measures.[6] A fully rational citizenry, knowing that both sites retain a chance of successful attack, would rely on both sites, with only a slight bias towards the more heavily defended one.

However, citizens may act heuristically and retain a bias towards the protected target if they do not grasp the strategic reasoning of the terrorist. The concept of "level-k" reasoning explains why this process may be difficult for the citizen (Camerer et al. 2004). At *level-0*, the citizens, responding simply to the government's action without considering the action of the terrorist, increasingly frequent the protected site. Terrorists, anticipating the citizen's action, use *level-1* reasoning, and attack the protected site. In turn, to increase their chance of avoiding the terrorist, citizens must correctly predict the terrorist's reaction to the citizen's own initial reaction to hardening.[7] The *level-2* reasoning required of citizens to

[5] Further evidence of this dynamic can be found in Iraq, when American forces pulled back in Iraq behind the heavily fortified 'Green Zone', and insurgents stepped up their attacks–even though the area was fortified http://www.nytimes.com/2010/09/30/world/middleeast/30iraq.html.

[6] While empirical work on how protecting a site affects the probability of attack there is difficult to come by, Berman and Laitin (2005) suggest that terrorists change tactics in response to site hardening, with more hardened sites being more likely to attract a suicide attack.

[7] This logic is similar to that of bounded rationality commonly exhibited in the beauty contest game. In the game, a pool of subjects choose an integer [0,100]. An individual who guesses closest to $\frac{2}{3}$ the average of the responses receives a monetary prize.If subjects completely backwardly induct, then they should guess 0 (the Nash Equilibrium). However, observed responses vary between 15 and 40, suggesting that

avoid the terrorists stretches the limit of the average person's reasoning capacity.[8] Paradoxically, the government's attempts to protect certain targets may in fact make individuals more vulnerable.

The first question this papers addresses is do how citizens and terrorists make decisions after observing public, defensive counter-terror measures? Do they act fully strategically and correctly randomize across targets? Or do they pool on publicly hardened targets?

The second question this papers examines is how citizens sanction or reward the government in the face of asymmetric information about government counter-terror activities. Bueno de Mesquita (2007) deviates from the above distinction between defensive and offensive counter-terrorism, dividing counter-terrorism into observable and tactic-specific efforts, such as airport security, and unobservable and general efforts, such as intelligence gathering. He argues that if the government has an incentive to engage in rent-seeking behavior, it creates a principle-agent problem between the government and the citizens. This inability to observe certain counter-terror provisions such as intelligence gathering places electoral pressure on the government to over-invest in publicly observable, tactic-specific counter-terror. Furthermore, if terrorists change tactics in response to defensive measures, tactic-specific counter-terror is suboptimal from a security standpoint, especially if the terrorists have a variety of tactics available (Bueno de Mesquita 2007; Faria 2006).

Governments over-invest in inefficient counter-terrorism strategies in response to the electorate's need to observe the government's actions. This inefficiency is further exacerbated by the fact that terrorist attacks are rare events, making it difficult for the citizen to evaluate whether it is actually government policy that is preventing terror attacks, or a variety of other factors, including luck. Empirical research suggests citizens blame the government for successful terrorist attacks, as terrorist attacks tend to hurt the incumbent parties' chance for reelection (Gassebner et al. 2008; Kibris 2011). What remains unclear though, is if citizens reward the government for a lack of attacks, even if they have no way to verify governmental action actually prevented an attack.

Several recent studies find that citizens use very coarse heuristics to reward incumbent governments with electoral support retrospectively. For instance, incumbents benefit electorally from events out of their control, including the outcomes of college athletic contests and lotteries (Healy et al. 2010; Bagues and Esteve-Volart 2011). Experimentally, Huber et al. (2012) show that lotteries explicitly stated to be unrelated to a simulated incumbent's ability affect subjects' willingness to retain that incumbent, with subjects keeping the incumbent more often after winning the lottery. This research suggests that a personal sense of well-being affects voters' choices in favor of incumbents, even if that sense of well-being is derived from non-political events. Likewise, voters punish the government for

Footnote 7 continued
humans have, or believe others have, a limited ability ("level-k") to reason backward and act fully strategic (Nagel 1995; Stahl and Wilson 1995).

[8] Camerer et al. (2004) estimate that an average person can reason through about one and a half steps of strategic interaction.

negative events outside of its control, including shark attacks, droughts, and other forms of extreme weather, though the punishment can be mitigated by an appropriate governmental response (Achen and Bartels 2004; Gasper and Reeves 2011; Cole et al. 2012). Wolfers (2002) shows that oil price drops decrease the probability of re-election for governors in oil-producing states, despite oil prices being unrelated to gubernatorial actions. In the context of terrorism, these results suggest that when voters know there is some chance of terrorism and have an imperfect ability to monitor government actions, an absence of successful attacks is likely to benefit the government. This benefit will accrue even if the government took no actions or its actions were ineffective.

We develop a laboratory experiment to test our claims about the strategic interaction of citizens and terrorists after observing public counter-terrorism and whether citizens reward or punish the government after the results of the interaction. While laboratory experiments invariably bring in a level of artificiality, they have several advantages over other methods (such as case studies and large-N studies). Since we are principally interested in the behavioral and strategic interaction between citizens, terrorists, and the government, the laboratory affords us the opportunity to isolate and control key variables of interest (e.g., risk from terrorism and degree of visibility of government counter-terror) to determine the effect on the behavior of the actors (McDermott 2011). Moreover, as Arce et al. (2011) argue, the interaction between terrorist and government can be complicated by multiple factors, making isolating believable explanations difficult. An experiment allows us to isolate these strategic interactions and avoid threats to causal inference. Finally, experiments may be especially useful to test competing hypotheses, such as when game theoretic models predict fully strategic behavior, and psychological theories suggest a more heuristic style of reasoning (Camerer 2003, p. 7). The next section details the set-up of our experiment.

## Experiment

We explore the interaction between a representative Government, Citizen, and Terrorist in the provision of counter-terrorism using a laboratory experiment.[9] The experiment sought to capture three important dynamics in counter-terrorism. (1) Both citizens and terrorists have incomplete information with respect to the vulnerabilities of targets.[10] (2) Citizens reward or sanction leaders based on incomplete information about the leader's provision of counter-terrorism. (3) The success of an attempted

[9] We denote Government, Citizen, and Terrorist with capital letters in the experiment. Lowercase government, terrorist, and citizen should be thought of as a more general argument.

[10] In reality, citizens may be less informed about target vulnerability than terrorists—who actively conduct surveillance and planning for attacks http://www.nytimes.com/2012/09/11/opinion/the-bush-white-house-was-deaf-to-9-11-warnings.html. Yet, we show in our experiment that even when the Terrorist and the Citizen have the same incomplete information, the Citizen still does not react as strategically as the Terrorist. This suggests that the dynamic of Citizens behaving less strategic than the Terrorists in our experiment is likely to be even stronger in real life, where terrorists have an informational advantage.

terrorist attack is not guaranteed, but rather is dependent on the level of defense of a target and the strategic behavior of the citizen and terrorist.[11]

We recruited subjects via New York University's Center for Experimental Social Science's undergraduate recruitment pool. 96 subjects participated in four experimental sessions in a computer lab. The experiment lasted approximately an hour and 15 min. Subjects were told at the outset of the experiment that they would be assigned to a group of three and to one of three roles.[12] The groups and roles remained the same throughout the experiment. Subjects were given a $10 show-up fee and were also paid according to the sum of points they earned in five randomly chosen rounds of the experiment (Morton and Williams 2010). Including the show-up fee, subjects earned an average of $19.61 ($1 = 12 points).

The experiment lasted for 4 blocks of 11 periods each (44 total periods). The experimental protocol proceeded as follows. In each period, the Government was given a budget of 10 points, which represented potential funds they could spend on counter-terrorism, and moved first. The Government then decided how many of the 10 points to keep and add to her own period payoff, and how many to invest in "hardening" options *A* or *B* against the Terrorist.[13] Every 1 point invested in hardening choice *A* or *B* adjusted the baseline probabilities in favor of (against) the Citizen (the Terrorist) by two points.

Next, the Citizen and the Terrorist decided simultaneously and independently of each other whether to choose option *A*, *B*, or *C*. The Citizen and the Terrorist had divergent preferences. The Citizen "won" with certainty if the Citizen and the Terrorist chose different options. If the Citizen and the Terrorist chose the same option, then with some probability ($p$) the Citizen won and with some probability ($1 - p$) the Terrorist won. This set-up, similar to a "matching pennies game," models the dynamics between citizens trying to avoid likely places where a terrorist could successfully strike, and a terrorist trying to strike a target successfully.[14] Both the Citizen and the Terrorist always knew the payoffs and the baseline probabilities associated with *A* or *B*, but not necessarily how much the Government spent on hardening them. The incomplete information about the Government's decision on the allocation of funds to harden targets represents citizens' and terrorists' uncertainty about how governments spend (or do not spend) on defense. Option

[11] For instance, see the failed plots of the so-called "Shoe Bomber" and the "Underwear Bomber" http://articles.cnn.com/2009-12-25/justice/richard.reid.shoe.bomber_1_terror-attacks-american-airlines-flight-qaeda?_s=PM:CRIME.

[12] In the experiment these were referred to as Players *1*, *2*, or *3*. See Online Appendix for a full list of instructions.

[13] On one hand, the assumption the Government keeps the points not spent on counter-terrorism spending may serve as a model of corruption. However, we find it more realistic that the Government shifts the counter-terrorism funds to another area of policy that helps it stay in power, rather than that the Government is personally appropriating the funds. In this way, the Government faces a trade-off between spending on counter-terrorism and on other policies. For example, see Chapter 3 of the 9/11 Commission Report on the low priority of counter-terrorism spending in the mid-1990s, even after the first World Trade Center Bombing (Kean 2011).

[14] We are agnostic as to what a terrorist "success" represents. It could be the terrorist seeking to inflict maximum casualties against a citizenry. Or it could be the terrorist seeking to strike a high value target. All that is required for our game is that the Terrorist and Citizen have opposing preferences.

*C* was a safer choice for the Citizen and paid lower in expectation than *A* or *B* for both the Citizen and the Terrorist. *C* models the choice of Citizens to change their daily routine and minimize the risk of suffering a Terrorist attack, even if it is inefficient. Meanwhile, *C* for the Terrorist represents the small cost of planning and then not following through with an attack.[15]

The baseline probabilities and payoffs for *A* and *B* were the same in expectation for both players. The baseline probabilities, displayed in Table 1, were randomly drawn for each treatment block and remained the same throughout all 11 periods of each block. After all three players learned whether the Citizen or the Terrorist won, the Citizen had the option to give a bonus of between 0 and 25 points to the Government. This bonus was costless to the Citizen. This bonus can be thought of as an election where the Citizen can reward the Government based on the information available to them in the game and the outcome (did they experience a terrorist attack or not).

The payoffs to the Citizen and the Terrorist are found in Table 1. The Government's payoff was $Y_G = 10 - I_G + B_Z$, where $I_G$ represents how much the Government invested in hardening in that round, and $B_Z$ is how much the Citizen gave as a bonus.

Our chief question of interest is how does varying information about the provision of counter-terror spending influence strategic behavior and subsequent electoral approval of the elected official? In our experiment we had four treatments that varied the level of information both the Citizen and the Terrorist observed about the Government's investment in hardening options *A* and *B*. The four treatments, which appeared to the subjects in random order by group, are listed below:

1. **Private Harden**: The Government's investment in hardening remains private.
2. **Public Harden One**: The Government's investment in hardening is made public for only option *A*.
3. **Public Harden Both**: The Government's investment in hardening is public for both options *A* and *B*.
4. **Government Choice**: Each round, the Government chooses whether to make their investment in hardening public for both options *A* and *B* or private for both options. The Citizen and the Terrorist are aware of the decision.

By experimentally varying the information available to a representative Citizen and the Terrorist, we can directly measure how incomplete information about terrorism influences strategic behavior between a terrorist group and citizenry. In particular, the *Public Harden One* treatment allows us to determine if the Citizen will respond to public hardening heuristically and choose the option where she knows hardening occurred, and if, in turn, the Terrorist will anticipate and attack there. Furthermore, we can determine how the principal-agent relationship between the government and the electorate changes with the level of information available to the electorate. The next section generates predictions related to this experimental set-up.

[15] The small probability of a successful Terrorist attack when both the Citizen and the Terrorist choose *C* can be thought of as an event targeted at a particular citizen, such as a kidnapping, whose victims tend to be middle class locals (Forest 2012). Santifort and Sandler (2013) find that, conditional on a successful kidnapping, in only 27.5 hostage incidents do terrorist achieve even some of their goals.

**Table 1** Baseline probabilities for the Citizen and the Terrorist for Lotteries between options A and B (if the citizen and the terrorist choose the same option)

| Probabilities | | | | Payoffs (in points) | | | | Expected payoff (in points) | |
|---|---|---|---|---|---|---|---|---|---|
| Terrorist win | Terrorist loses | Citizen win | Citizen loses | Terrorist win | Terrorist loses | Citizen win | Citizen loses | Terrorist | Citizen |
| 0.5 | 0.5 | 0.5 | 0.5 | 140 | −20 | 80 | −48 | 20 | 16 |
| 0.6 | 0.4 | 0.4 | 0.6 | 90 | −10 | 76 | −24 | 20 | 16 |
| 0.7 | 0.3 | 0.3 | 0.7 | 98 | −22 | 100 | −20 | 20 | 16 |
| 0.8 | 0.2 | 0.2 | 0.8 | 98 | −32 | 100 | −5 | 20 | 16 |
| *0.1* | *0.9* | *0.9* | *0.1* | *−60* | *−1* | *4* | *−10* | *2.55* | *2.6* |

(Note that probabilities for the option *C* (the safe option) and associated payoffs are highlighted in italics in the bottom row of the table. Also note that the expected value is skewed in favor of the Terrorist. This is to account for the fact that the Government had the ability to manipulate these baseline probabilities in favor of the Citizen.)

## Predictions

The game between the Citizen and the Terrorist most closely resembles a matching pennies game with a probabilistic element when the Terrorist successfully guesses the Citizen's choice. One key question is how does public investment in hardening by the Government affect the strategic behavior of the Citizen and the Terrorist. Theoretically, the hardening available to the Government here should have only a modest impact on the strategic interaction of the Citizen and the Terrorist. For instance, if the Government publicly hardens option *A*, and the Citizen increasingly chooses *A*, then the Terrorist will respond by also increasingly choosing *A*. The Citizen should then respond by increasingly choosing option *B*. These best-response dynamics continue until both the Citizen and the Terrorist are indifferent between options *A* and *B*. In other words, public signals of investment should not necessarily sway the Citizen to choose the hardened option overwhelmingly more often, especially not to the point where they would be better off choosing the unhardened option, because the Terrorist observes the signal.

We derive a prediction of how Citizens should respond to public hardening by measuring the increase in the frequency with which the Citizen plays that option in mixed strategy equilibrium. Solving for a mixed strategy equilibrium in this game is slightly more complicated than in a standard matching pennies game because even when the Citizen and the Terrorist make the same choice, the Citizen still can win. Thus, for the Terrorist to make the Citizen indifferent, the Terrorist solves the following equation for $t^a$:

$$t^a(P^a_{cw} \times Y^a_{cw} + (1 - P^a_{cw}) \times Y^a_{cl}) + (1 - t^a) \times Y^a_{cw} = \\ (1 - t^a) \times (P^b_{cw} \times Y^b_{cw} + (1 - P^b_{cw}) \times Y^b_{cl}) + t^a \times Y^b_{cw}$$

where $t^a$ represents the probability the Terrorist chooses option *A*, $P^i_{cw}$ represents the probability the Citizen wins when both the Terrorist and the Citizen chose option *i*, $Y^i_{cw}$ represents the payoff to the Citizen for winning when the Citizen chose option *i*, and $Y^i_{cl}$ represents the payoff to the Citizen for losing when the Citizen chose option *i*. The Citizen uses a similar equation to make the Terrorist indifferent, except the Terrorist loses with certainty when the Citizen and the Terrorism make different choices (the last term on both sides of the equation). Comparative statics for the Terrorist's and Citizen's equations are presented in the Online Appendix. The key results are that as the Government hardens an option, both the Citizen and the Terrorist should choose that option more frequently as part of a mixed strategy equilibrium. However, the maximum a Citizen should shift towards a hardened option is 1.27 percent per point of hardening.

In a standard matching pennies set-up, Goeree and Holt (2001) show that subjects employ a mixed strategy that closely resembled the Nash Equilibrium when those subjects' payoffs were equal for their two choices. However, when one of the two choices offered a higher payoff, subjects exhibited "own-payoff effects," picking the choice with the higher payoff much more frequently than the Nash Equilibrium prediction (Ochs 1995; Goeree and Holt 2001; Goeree et al. 2003). This tendency makes it easier for the subjects' opponents to correctly predict the subjects' choices. In

the present experiment, the Government is given the opportunity to change the expected value of the choices in favor of the Citizen by hardening an option. When the Government hardens an option, previous research suggests that Citizens may shift their choices to the hardened option too often, allowing the Terrorist to predict where the Citizen will be more often than in equilibrium. Moreover, the k-level reasoning model discussed above posits an extra step of reasoning required by the Citizen in response to Government hardening and the Terrorist's anticipated reaction. Accounting for the combination of own-payoff effects and the cognitive demands on the Citizen, it is likely that the Terrorist may benefit from public hardening.

Meanwhile, the interaction between the Government and the Citizen is a principal-agent problem. The Citizen wants the Government to invest in counter-terrorism to help ensure her safety. Should the Citizen avoid the terrorist attack due to the Government's effort, the Citizen can reward the government with a high bonus. In fact, previous work has shown that subjects vote retrospectively, which induces office-motivated politicians to act in the citizen's best interest and can increase the provision of public goods (Woon 2012; Hamman et al. 2011). However, in some treatments of our experiment, the Citizen has no way to monitor the Government and determine if counter-terrorism helped prevent an attack or not. In these cases, Citizens may incorrectly attribute the lack of a successful attack to Governmental action. Here, the Government may anticipate that the Citizen is unlikely to experience a successful attack and reward the Government regardless of the actions they took (or did not take). When this belief is held, the Government has the incentive to ignore the Citizen's desire for counter-terrorism and, instead, under invest in hardening targets.

We develop three hypotheses derived from the bounded rationality literature and our formal model.

**Hypothesis 1** When government counter-terror hardening is made public, citizens will overly frequent protected targets, making them easier targets for terrorists.

**Hypothesis 2** Anticipating **Hypothesis 1**, public hardening will cause terrorists to attack the hardened target at a greater frequency.

**Hypothesis 3** Given the literature on retrospective literature, in the absence of information about government counter-terror provision, citizens will reward governments for the absence of terrorism, and sanction them for its presence.

In the next section, we present the results of our experiment and provide interpretation.

## Results

### Strategic Interaction

We empirically examine whether our theoretical predictions hold in our experiment, beginning with a summary of the choices made by the Citizen and the Terrorist. Figures 1 and 2 present these choices broken down by the treatment.

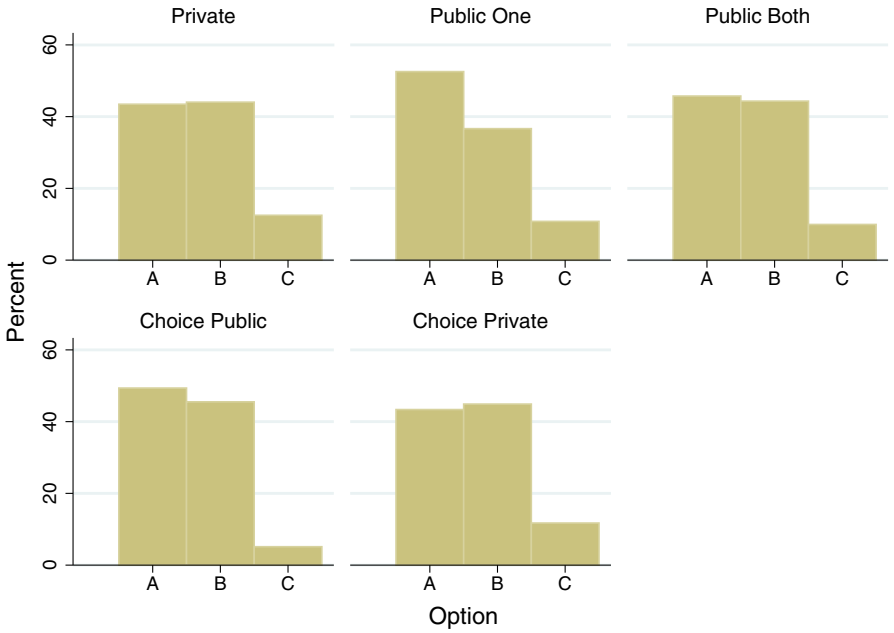


**Fig. 1** Histograms summarizing the citizens' choices by treatment

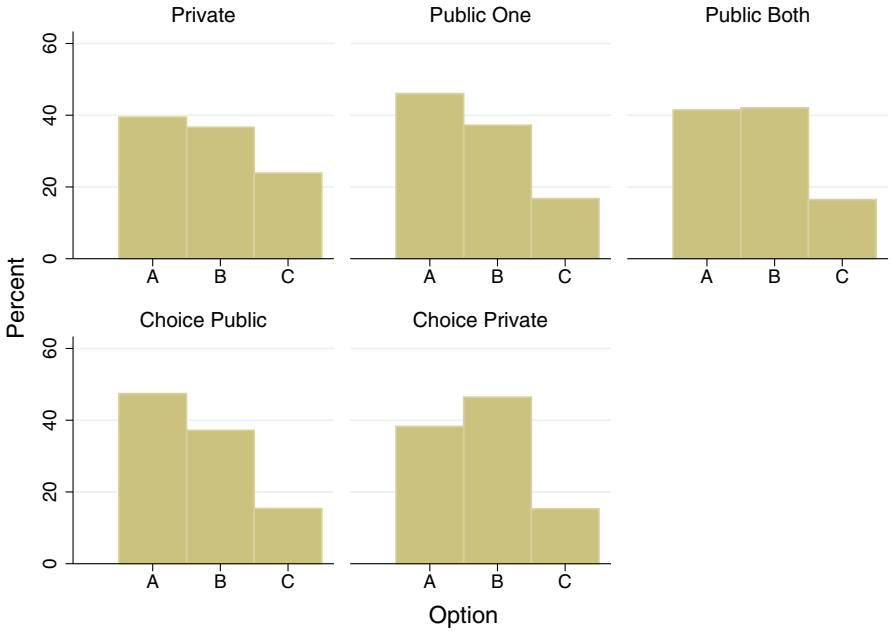


**Fig. 2** Histograms summarizing the terrorists' choices by treatment

Across all treatments, with the exception of the Government Choice treatment when the Government made its spending public, the Citizen chose option *C*, the safe option, about 10 percent of the time. That percentage was approximately halved when the Government actively decided to publicize its spending, most likely because this lead to an increase in hardening, as discussed below. Also of note is that, overall, the Citizen mixed fairly evenly between option *A* and *B*. The exception is when only the hardening for Option *A* was made public, in which case the Citizen gravitated towards that option. The repercussions of this tendency to choose sites with public hardening for the Citizen is addressed next. First, however, it is worth noting that the Terrorist selected option *C* more often than the Citizen. In some treatments, the Terrorist selected this safe option twice as often as the Citizen.

Table 2 addresses the effect of public hardening by the Government on the choices of the Citizen and the Terrorist. We include treatments where the hardening of option *A* was public.[16] The first two columns in Table 2 examine the marginal effect of a 1 point increase in public hardening of *A* on the probability that the Citizen (Model 1) and the Terrorist (Model 2) choose option *A*.[17] These models suggest that public investment in the hardening of *A* serves as a coordinating signal to choose *A*. A public 1 point increase in *Govt. Harden A* makes the Citizen approximately 2.6 % and the Terrorist approximately 2.8 % more likely to chose *A*, respectively. Government hardening leads Citizens to increase their selection of option *A* more than twice as much as the theoretical maximum (see Online Appendix). Public hardening, then, serves as a coordination device for the Terrorists, to their advantage, and they increase their frequency of choosing that option as well. These results confirm Hypotheses 1 and 2.

With both the Citizen and the Terrorist increasingly choosing *A* in response to public hardening on it, we address whether, conditional on choosing *A*, increases in the hardening of *A* leads the Citizen to win more. In effect, did public hardening on a site make the Citizen safer there? Model 3 shows it did not.[18] Rather *Govt. Harden A* had no discernible effect on the probability that the Citizen won if it chose *A*. Moreover, Models 4 and 5 show that the Citizen would have won more often by choosing *B* as the Government increasingly hardened *A*, and, conversely, the Terrorist would have lost more often. These results suggest that while the Terrorists were responding strategically to increases in public hardening of *A*, Citizens were not. It appears that Citizens did not engage in the extra step of reasoning required to avoid pooling on option *A*. This inability was problematic for the Citizen given that Terrorist responses did tend to correctly predict the Citizen's initial reaction to

[16] These treatment blocks include **Public Harden One**, **Public Harden Both**, and **Government Choice** in which the Government chose to make hardening public. The results do not change substantially when we exclude the **Government Choice** treatment.

[17] Throughout the paper we use random effects panel regression. Given that we randomly assigned individuals to groups and treatments, we assume that any baseline differences are orthogonal to the treatment (Greene 2008)—thus random effects is preferable. However, there may be a concern that in some models the use of a lagged dependent variable biases the estimates and random effects. The Online Appendix presents the results from Table 4 and Fig. 5 (Model 3) using fixed effects. The results there are nearly identical to those in Table 4 and Fig. 5 (Model 3), alleviating any concerns about the use of random effects.

[18] A discussion of learning during the experiment and public hardening appears in the Online Appendix.

**Table 2** Marginal effects of public Government hardening of option *A* on strategic behavior

| Dependent variable (Conditions) | Citizen chooses A | Terrorist chooses A | Citizen wins (If she chose *A*) | Citizen wins (If she chose *B*) | Terrorist wins (If she chose *B*) |
|---|---|---|---|---|---|
| | *Model 1* | *Model 2* | *Model 3* | *Model 4* | *Model 5* |
| *Govt. Harden A* | 0.0257*** | 0.0281*** | −0.00138 | 0.0249* | −0.0235* |
| | (0.00675) | (0.0071) | (0.0071) | (0.0121) | (0.0106) |
| *N* | 860 | 860 | 423 | 356 | 337 |
| Number of clusters | 32 | 32 | 32 | 32 | 32 |

Estimated from a random effects panel logit

* $p<0.05$ ** $p<0.01$ *** $p<0.001$

hardening. The Citizen's over-reliance on hardened sites points to a limitation in the ability of individuals to strategically respond to public information (coordination bias) and derive optimal best responses. In other words, the Citizens' bias towards the hardened target made it easier for the Terrorists to predict which option they would choose, and this predictable bias cost the Citizen points.

## Experimental Political Behavior and Counter-terrorism

We now address the principal-agent problem between the Citizen and the Government. Figure 3 presents the treatment effects for the different levels of information on the Government's investment in hardening. From the graph, three things become apparent. (1) In all the treatments, the Government kept most of the initial endowment for itself (none of the graphs are above 5 on average). (2) Increasing levels of information provided to the Citizen (and the Terrorist) about investment in hardening corresponds to higher levels of investment in hardening. (3) Finally, in the *Government Choice* treatment, Governments that make their investment choices public (*Choice Public*) invest the most in hardening, while those who chose to keep their choices private (*Choice Private*) invest the least. This suggests that announcing Governmental intentions to conceal the investment in counter-terror might serve as a signal to the Citizens of low expenditures.

Figure 4 presents the treatment effects for the different levels of information on the Citizens' bonus to the Government. The higher bonuses here do not seem to correlate with higher investment in Fig 3, suggesting that Citizens were not rewarding the Government for high investments on their behalf. Three effects can be directly observed from this graph. (1) On average, Citizens gave more than half of the allotted 25 points to the Government across all treatments. (2) Contrary to Fig 3, there is not a linearly increasing relationship in information about the

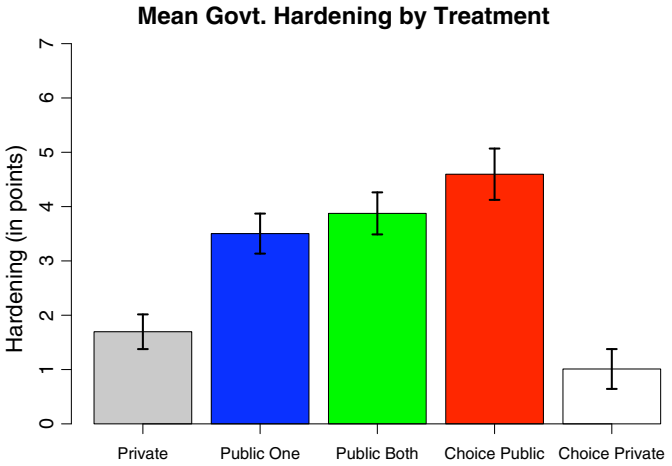


**Fig. 3** Treatment effects on Government investment hardening of targets (95 % CI bars. *Y-axes* represents point allocations)

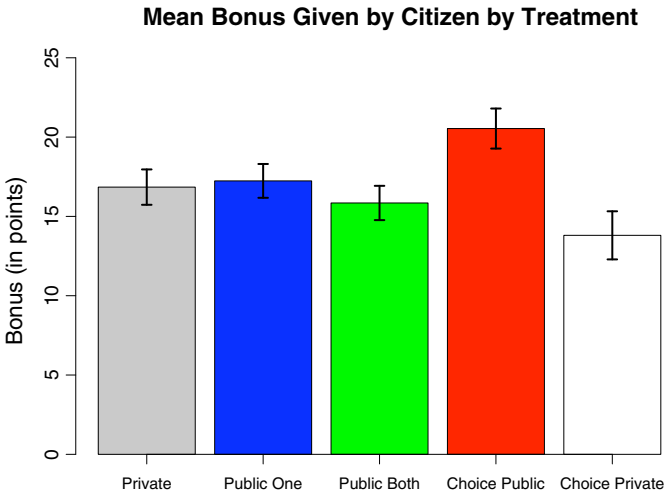


**Fig. 4** Treatment effects on Citizen bonus. (95 % CI bars. *Y-axes* represents point allocations)

Government's investment in hardening and bonus levels. In fact, Citizens gave slightly lower bonuses in *Public Both*, than *Private* and *Public One*. It appears that subjects are overestimating the amount that the Government invests when there is incomplete information. (3) Finally, Citizens accurately interpret the Government concealing their choices in *Government Choice* (*Private Choice*) as an indication the Government was likely to invest less in hardening. Taking points (1) and (2) together, whether or not the Citizen was the victim of a successful attack by the Terrorist likely influences the size of the bonus as much as the actual levels of Government spending. Because it was relatively easy for the Citizen to receive a positive payoff regardless of the Government's level of investment, Citizens gave high bonuses quite often regardless of whether they could attribute the successful outcome to the Government or not. As point (3) suggests, only actively concealing the Government's investment lowers this bonus on average, and even then somewhat moderately. We discuss the Citizen's willingness to reward the Government in more depth below.

Given the censored nature of the Government's investment in hardening ($0 \leq I_G \leq 10$), we use a Tobit model in Table 3 to estimate the effect of the treatments. In almost 45 % of the periods the Government did not invest anything in hardening (*N* Left cens.). Model 1 represents the core specification with dummy variables for the treatment effects (relative to *Private* treatment), the amount of bonus given by the Citizen in the previous period ($Bonus_{(t-1)}$), and a dummy variable for whether the Government chose to make its hardening choices public (*Public*) interacted with *Govt. Choice* (*Public* X *Govt.Choice*). Model 2 adds the lagged dependent variable ($Govt.Harden_{(t-1)}$) and Model 3 includes dummy variables to account for time trends within and across blocks.[19] Governments

[19] These are to account for learning in subjects' behavior both within treatment blocks and across time.

**Table 3** Random effects panel tobit model of Government hardening by treatment group (relative to private hardening treatment)

| | Dep. Var.: Govt. Harden | | |
|---|---|---|---|
| | *Model 1* | *Model 2* | *Model 3* |
| *Public Harden One* | 4.432*** | 3.569*** | 3.609*** |
| | (0.444) | (0.439) | (0.438) |
| *Public Harden Both* | 5.503*** | 4.490*** | 4.199*** |
| | (0.452) | (0.449) | (0.444) |
| *Govt. Choice* | −2.819*** | −2.984*** | −3.230*** |
| | (0.644) | (0.625) | (0.621) |
| *Public X Govt. Choice* | 8.906*** | 8.456*** | 8.813*** |
| | (0.726) | (0.700) | (0.702) |
| $Bonus_{(t-1)}$ | 0.116*** | 0.0819*** | 0.0976*** |
| | (0.0212) | (0.0207) | (0.0205) |
| $Govt.Harden_{(t-1)}$ | – | 0.427*** | 0.396*** |
| | – | (0.0526) | (0.0519) |
| *Constant* | −4.093*** | −4.102*** | −2.939*** |
| | (0.976) | (0.825) | (0.854) |
| *Within block trends* | | | ✓ |
| *Across block trends* | | | ✓ |
| *N* | 1376 | 1376 | 1376 |
| *N* Left cens. ($Y \leq 0$) | 615 | 615 | 615 |
| *N* Right cens. ($Y \leq 10$) | 165 | 165 | 165 |

Robust standard errors clustered in parentheses at the group level

* $p<0.05$ ** $p<0.01$ *** $p<0.001$

invested more on average when the Citizen gave a higher bonus in the previous period (positive coefficient on $Bonus_{(t-1)}$). Moreover, largely confirming the results from Fig. 3, increased levels of information relative to the *Private* treatment increase *Govt. Harden* except for *Govt. Choice*. There is an interactive effective between *Govt. Choice* and *Public*, with Governments that chose to go *Public* investing more (positive sign on the interaction term *Public X Govt. Choice*), and those who did not investing less (negative coefficient on *Govt. Choice*) relative to the *Private* treatment. Taken together, these results indicate that increased monitoring ability on behalf of the Citizen leads to increased counter-terror expenditures.

In Table 4, we examine how the information available to Citizens influences their decision to award a bonus to the Government. As in Table 3, we use a Tobit model to account for censoring of the Government bonus ($0 \leq B_Z \leq 25$). In contrast to Government spending on investment decisions, there is a large amount of right censoring. In over 50 % of the periods, Citizens gave the Government the full bonus (25 points). For each specification we include dummy variables for each treatment effect (relative to *Private* treatment), the amount that the Citizen was able to see the

**Table 4** Random effects panel tobit model of Citizen's bonus by treatment group (relative to private hardening treatment)

| | Dep. Var.: Bonus | | |
|---|---|---|---|
| | *Model 1* | *Model 2* | *Model 3* |
| *Visible Harden* | 2.033*** | 1.422*** | 1.505*** |
| | (0.246) | (0.224) | (0.224) |
| *Public Harden One* | −5.811** | −4.970** | −4.717** |
| | (1.880) | (1.728) | (1.751) |
| *Public Harden Both* | −12.67*** | −9.203*** | −8.286*** |
| | (1.972) | (1.830) | (1.849) |
| *Govt. Choice* | −5.073** | −4.265** | −3.060 |
| | (1.774) | (1.628) | (1.630) |
| *Citizen wins* | 11.87*** | 12.72*** | 13.13*** |
| | (1.400) | (1.289) | (1.281) |
| $Bonus_{(t-1)}$ | – | 0.874*** | 0.826*** |
| | – | (0.0726) | (0.0716) |
| *Constant* | 17.19*** | 2.658 | 1.055 |
| | (4.998) | (3.977) | (4.152) |
| *Within block trends* | | | ✓ |
| *Across block trends* | | | ✓ |
| *N* | 1408 | 1376 | 1376 |
| *N* Left cens. ($Y \leq 0$) | 285 | 275 | 275 |
| *N* Right cens. ($Y \leq 25$) | 731 | 723 | 723 |

Robust standard errors clustered in parentheses at the group level

* $p<0.05$ ** $p<0.01$ *** $p<0.001$

Government invest in hardening (*Visible Harden*),[20] and whether the Citizen won in that period (*Citizen win*). Not surprisingly, Citizens that won in a given period award a higher bonus that period across treatments. Furthermore, a visible increase in the hardening of targets by the Government leads to an increase in bonuses awarded by the Citizen. Interestingly, *Public Both* has a negative sign on it. This coefficient should be interpreted as an interaction term with *Visible Harden*, and even at average levels of investment in the *Public Both* treatment there is a slightly negative effect.[21] This deserves to be highlighted: as Citizens become more aware of how much the Government hardened, they actually reward them *less* even thought they are hardening *more* relative to the *Private hardening treatment*.

Since the effect of *Visible Harden* depends on the treatment and is not directly interpretable from Table 4, we plot the predicted level of bonus given by the Citizen

[20] For instance in *Private* treatment, *Visible Harden* could only take on zero, whereas in *Public Both* treatment *Visible Harden* would take on the value $I_G$. In the *Public One* treatment, only the investment in option *A* counts toward the *Visible Harden* variable

[21] Governments invested on average 3.875 points, so from Model 3 the main effect on the latent $Bonus = 1.505 \times 3.875 + -8.286 = -2.454125$

in each treatment based on the average level of hardening in each treatment.[22] The horizontal dashed lines represent the mean level of bonus given in the *Private* (blue) and *Choice Private* (red) lines, where no public hardening can occur, as a reference. The plots confirm that rather than rewarding the increased hardening that accompanies increased information with higher bonuses, the Citizen appear to largely to give similar bonuses across treatments (and actually give less in the *Public Both*).[23]

Our results overall suggest that Citizens respond heuristically in both their strategic interaction with the Terrorist and the Government. Citizens chose options with public hardening to a fault. When observing hardening, the average Citizen seems to consider only the increased the safety of an option without also considering the Terrorist's strategic response. In general, Citizens tended to fail in taking the step to level-2 reasoning and account for both the increased safety of a site and the Terrorist's prediction about their own behavior with respect to that increased safety.

Furthermore, while the Government in our experiment spent relatively little on counter-terrorism, as indicated by the large number of left censoring ($I_G = 0$), Citizens rewarded them quite highly with the bonus, as indicated by the large amount of right censoring ($B_Z = 25$).[24] A concern may be that our experimental set up artificially induced conservatism in Government counter-terror provision because the government hardening does not "protect" the Citizen that much, and because in some of the treatments, the Citizen lacked the ability to monitor Government spending. However, we argue that this is false for two reasons. (1) The hardening itself is fairly significant—if a government invests the full amount (10 points), they can harden a one option's baseline probability by 20 % points. Admittedly, in equilibrium, government hardening has a small effect on the actual choice of the Citizen (a maximum of 1.27 % per point). However, in a purely rational world, this spending makes both sites safer as the Terrorist adjusts his mixed strategy towards the hardened site. Therefore, while hardening induces only a small change in behavior in equilibrium, it should have a significant effect on the payoff of the Citizen.[25] Nevertheless, as the experimental results show, the Citizen choose sites with publicly observable counter-terrorism more frequently than they should have, allowing the Terrorist to match the choice of the Citizen more easily. The net result is that the Citizen was no safer after public hardening. However, this finding is due to the poor heuristic of the Citizens rather than because hardening in the experiment is inherently ineffective.[26] (2) Many counter-terror provisions are inherently

[22] The Government invested an average of 3.50, 3.875, and 4.60 points in hardening in the *Public One*, *Public Both*, and *Choice Public* treatments, respectively.

[23] Bonuses in the *Public Both* are on average 1.4 points lower than in the *Public One* treatment (two-tailed *p value*=0.07). It should be emphasized that treatment order was randomized across groups.

[24] In fact, when we reformulate the dependent variables $I_G$ and $B_Z$ as binary variables—1 right censored for $B_Z$, 0 otherwise and 1 left censored for $I_G = 0$, 0 otherwise, the results are even stronger (see Online Appendix).

[25] See the Online Appendix for a discussion of the expected increase in the Citizen's payoff per point of hardening, which is calculated to be between .5 and .9.

[26] Powell (2007b) makes a similar point in arguing why strategic defense does not necessarily make targets uniformly safer due to strategic behavior on the part of terrorists.

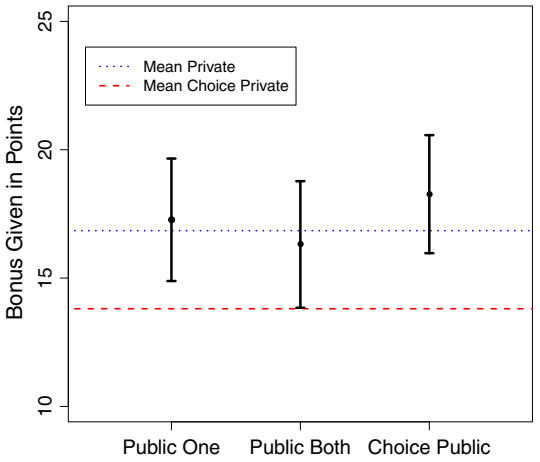


**Fig. 5** Predicted bonus based on average level of hardening across treatments. 95 % CI bars around the point estimate from the delta method. *Y-axes* represents point allocations. Point estimates are based on a Tobit model of average level of government hardening in each of the three treatments in which subjects could observe levels of hardening (*Public One, Public Both, Choice Public*) from Table 4, Model 3. For reference, the higher dashed horizontal line is the mean level of bonus in *Private* treatment, and the lower dashed horizontal line is the mean level of hardening in the *Choice Private*

secretive in nature, and given the low probability of terrorist attacks, this may exacerbate a lack of spending on unobservable, but effective, counter-terror (Bueno de Mesquita 2007). Furthermore, we do not mean to say that governments spend nothing on counter-terror, but rather that citizens may think that governments are allocating money to counter-terror, but because of a lack of oversight/monitoring, they may be in fact using it for other purposes.

The experiment suggests that Citizens responded with a heuristic decision rule that placed a large emphasis on whether they avoided a successful Terrorist attack, and placed less emphasis on the actual provision of counter-terrorism. In short, Citizens rewarded the Government for the absence of terrorism, especially when they were unable to monitor the Government's effort. This result confirms Hypothesis 3 and is consistent with the larger literature on retrospective voting that finds voters may reward or punish politicians based on outcomes—even those in which the politician had no hand in creating (Healy et al. 2010; Bagues and Esteve-Volart 2011; Achen and Bartels 2004; Gasper and Reeves 2011; Wolfers 2002). Our findings suggest that Citizens with incomplete information have limited ability to effectively monitor Government behavior. However, when Citizens could monitor the Government, they decreased the bonus (the negative sign on *Public Both* in Table 4) even though the Government actually spent more on hardening in this treatment. This final point indicates the Government did consider their likely bonus when setting their level of hardening—when that level could be observed. When hardening could not be observed, the low probability of a successful attack meant the Government felt keeping points (representing spending on other policies) to be a better strategy than hardening. We interpret this as a trading off of counter-terrorism spending for investment in other policies.

## Conclusion

We used an experiment to investigate the interplay between counter-terror provisions and the strategic interaction between Citizens and a Terrorist. We find that after observing counter-terror activities, Citizens responded heuristically by tending to choose a target that had been the recipient of counter-terror hardening. Terrorists, anticipating the Citizens' actions, increasingly directed attacks against this "more secure" option. Government hardening serves as a coordinating signal to both Terrorists and Citizens, at a detriment to the latter. Moving outside the experimental context, if terrorists are flexible with respect to which locations they target, and citizens believe governmental counter-terrorism is effective, citizens are likely to frequent locations with a high level of observable counter-terrorism. Thus, terrorists wishing to inflict maximum causalities should be willing to attack protected locations despite the lower probability of success. This effect may be even stronger factoring in the increased psychological impact of a successful attack on a seemingly protected target.[27] The "psychological bonus" terrorists receive from successfully attacking a seemingly protected site[28] would further strengthen our results that citizens relying on simple heuristics and reacting predictably to counter-terrorism can raise the value of protected sites for terrorists.

Furthermore, we show that when citizens cannot perfectly monitor governmental expenditures, they use heuristics that reward the government based on observed outcomes. In the absence of a successful terrorist attack *and* information on government hardening, the Citizens in the experiment rewarded the Government *as if* its policy was responsible for preventing the attack. However, the results indicate that the Governments in the experiment actually did very little to lower the probability of a successful attack—especially when its actions were not observable. The principal-agent problem between the citizen and the government is particularly acute with counter-terrorism for two reasons. First, many counter-terror operations can only be effective if done covertly (Bueno de Mesquita 2007). Second, terrorist attacks are probabilistic events. It is difficult to attribute the absence of terrorism to government policy definitively or know the scale of damage that was prevented.[29] Therefore, citizens are left with little information other than whether or not an attack occurred. Our experimental results show that citizens frequently rewarded the government for the lack of successful attacks, attributing effective policy to the government despite being unable to monitor this policy.

The key question following a stylized laboratory experiment is how externally valid are the conclusions and what are the limitations of these conclusions? We randomly varied the ability of a representative citizen to monitor government investment in counter-terror, but kept the ratio of spending to counter-terror

[27] We argue that a terrorist defeating the government at its strongest may instill fear in the citizenry, making attacks at protected sites even more valuable to the terrorist.

[28] This could also be a high value target, whereby hardening the target signals that it is high value to terrorists (Powell 2007b).

[29] For instance, both Getmansky and Sinmazdemir (2013) (terror attacks) and Johnston (2012) (drone strikes) use the quasi-random success of terror attacks and drone strikes, comparing successes to failures, to estimate their causal effect.

improvement fixed.[30] Yet, it might be fruitful in future experiments to vary this ratio, or allow the government to use more offensive measures (prevent the terrorist from being able to attack). If this ratio were increased, it would allow the government to increase a site's safety in such a way that it approached 100 %. As sites approach complete safety, the nature of the game presented here permits strategies that approach pure strategy equilibria. Thus, as sites become nearly 100 % safe, the cost associated with the citizens' first-order strategy of choosing sites with public hardening, which we observed in our experiment, would decrease dramatically and eventually approach zero.

Another assumption we make in the experiment is that citizens' ability to reward the government with support is independent of the success of a terrorist attack.[31] If we tied the payoff of the Government to the payoff of Citizen in the experiment, it would remove some of the tension of the principal-agent problem between the two by aligning their incentives. This principal-agent relationship is something we, along with other researchers (Bueno de Mesquita 2007; Healy et al. 2010), have argued is central to political behavior. Yet, it is possible that if we altered this principal-agent relationship in such a way, the Government could harden targets in a manner that changes or eliminates the use of poor heuristics on the part of the Citizens. Nevertheless, with respect to citizen support for the government, we show that citizens reward governments for the absence of terrorism, and, when they have the ability to monitor government counter-terror spending, for expending resources on counter-terror. Even if the payoff of the citizen and government are closely linked, citizens are still likely to reward governments for the absence of terrorism as long as (1) citizens have incomplete information about government counter-terror spending and (2) the base rate of successful terror attacks remains low.

The broader implications of this work suggest that any time the citizens have incomplete information about the government's actions, they are likely to form their opinion based on observed outcomes.[32] This type of retrospective voting incentivizes the government to under-provide public goods that cannot be monitored. In fact, (Healy and Malhotra 2009) show empirically that even government spending on natural disaster preparedness, which may be observable, is not rewarded by citizens.

Together with our research, this finding suggests that under-provision of goods will be particularity problematic for low-probability events, such as terrorism and natural disasters, due to asymmetric information between voters and elected officials and voter re-election heuristics that emphasize outcomes (Achen and Bartels 2004). Imperfect monitoring of government actions and low-probability

[30] 1 point spent by the government reduced the probability of a successful terrorist attack by 2 percentage points if the terrorist guessed correctly.

[31] I.e., the size of the bonus the Citizen can give the Government does not depend on the success of a Terrorist attack.

[32] The recent uproar over the disclosure by former US intelligence contractor Eric Snowden of secret US surveillance programs run by the National Security Agency (NSA) on ordinary Americans suggests an interesting alternative http://www.theguardian.com/world/the-nsa-files. In our model, the government can shirk by simply not investing in private hardening. However, as the Snowden revelations suggest, they may use secrecy to as a cover for far more invasive surveillance. Modeling this would be an interesting extension.

events combine to hinder the ability of retrospective voting to sanction or reward elected officials. Furthermore, recent research suggests that voters do not uniformly evaluate parties–rather certain parties "own" national security and terrorism issues (Abramson et al. 2007; Wright 2012; Getmansky and Zeitzoff 2014). For instance, left-wing, or more dovish politicians may be blamed more or rewarded less than right-wing politicians for identical amounts of terror. Finally, this inability to sanction or reward governments accurately has implications for the fundamental democratic relationship between elected officials and voters. It also suggests that democratic accountability may distort incentives in strategic defense spending, which in turn influences the tactics and strategies of terrorists.

**Acknowledgments** Funding for this research was provided by NSF award number 0905044.

## References

Abramson, Paul R., Aldrich, John H., Rickershauser, Jill, & Rohde, David, W. (2007). Fear in the voting booth: The 2004 presidential election. *Political Behavior*, *29*(2), 197–220.

Achen, C. H., Bartels, & L. M. (2004). Blind retrospection. Electoral responses to drought, flu, and shark attacks. Estudios/Working Papers (Centro de Estudios Avanzados en Ciencias Sociales), (199), 1.

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

Arce, D. G., Croson, R. T. A., & Eckel, C. C. (2011). Terrorism experiments. *Journal of Peace Research*, *48*(3), 373–382.

Bagues, M., & Esteve-Volart, B. (2011). Politicians Luck of the Draw: Evidence from the Spanish Christmas Lottery. *Working Paper*.

Berman, E., & Laitin, D. (2005). Hard targets: Theory and evidence on suicide attacks. *NBER Working Paper*, 11740.

Berrebi, Claude, & Klor, Esteban F. (2006). On terrorism and electoral outcomes theory and evidence from the Israeli-Palestinian conflict. *Journal of Conflict Resolution*, *50*(6), 899–925.

Berrebi, Claude, & Klor, Esteban F. (2008). Are voters sensitive to terrorism? Direct evidence from the Israeli electorate. *American Political Science Review*, *102*(03), 279–301.

Blalock, Garrick, Kadiyali, Vrinda, & Simon, Daniel H. (2009). Driving fatalities after 9/11: A hidden cost of terrorism. *Applied Economics*, *41*(14), 1717–1729.

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

Cadigan, J., & Schmitt, P. M. (2010). Strategic entry deterrence and terrorism: Theory and experimental evidence. *Public Choice*, *143*(1), 3–22.

Camerer, Colin. (2003). *Behavioral game theory: Experiments in strategic interaction*. Princeton: Princeton University Press.

Camerer, Colin F., Ho, Teck-Hua, & Chong, Juin-Kuan. (2004). A cognitive hierarchy model of games. *The Quarterly Journal of Economics*, *119*(3), 861–898.

Cole, S., Healy, A., & Werker, E. (2012). Do voters demand responsive governments? Evidence from Indian disaster relief. *Journal of Development Economics*, *97*, 167–181.

Cox, Andrew, Prager, Fynnwin, & Rose, Adam. (2011). Transportation security and the role of resilience: A foundation for operational metrics. *Transport Policy*, *18*(2), 307–317.

de Rugy, V. (2006). Spending too much the wrong way on the wrong things. *CATO Unbound*.

Enders, W., & Sandler, T. (1993). The effectiveness of antiterrorism policies: A vector-autoregression-intervention analysis. *American Political Science Review*, *87*, 829–844.

Enders, W., & Sandler, T. (2002). Patterns of transnational terrorism, 1970–1999: Alternative time-series estimates. *International Studies Quarterly*, *46*(2), 145–165.

Faria, J. R. (2006). Terrorist innovations and anti-terrorist policies. *Terrorism and Political Violence*, *18*(1), 47–56.

Forest, James J. F. (2012). Global trends in kidnapping by terrorist groups. *Global Change, Peace and Security*, *24*(3), 311–330.
Gasper, J. T., & Reeves, A. (2011). Make It Rain? Retrospection and the attentive electorate in the context of natural disasters. *American Journal of Political Science*, *55*(2), 340–355.
Gassebner, M., Jong-A-Pin, R., & Mierau, J. O. (2008). Terrorism and electoral accountability: One strike, you're out!. *Economics Letters*, *100*(1), 126–129.
Getmansky, A. (2011a). Protection politics: The effect of Government's counter terrorism policy on terrorists' choice of targets. *Working Paper*.
Getmansky, A. (2011b). Whose life is worth more? Domestic politics, counterterrorism, and terrorists selection of targets. *Working Paper*.
Getmansky, A., & Zeitzoff, T. (2014). Terrorism and voting: The effect of rocket threat on voting in Israeli elections. *American Political Science Review*, Forthcoming.
Getmansky, A., & Sinmazdemir, T. (2013). Success breeds failure: The effect of terrorist attacks on land seizures in the Israeli-Palestinian conflict. *Working Paper*.
Goeree, J.K., & Holt, C.A. (2001). Ten little treasures of game theory and ten intuitive contradictions. *American Economic Review*, 1402–1422.
Goeree, J. K., Holt, C. A., & Palfrey, T. R. (2003). Risk averse behavior in generalized matching pennies games. *Games and Economic Behavior*, *45*(1), 97–113.
Greene, William H. (2008). *Econometric analysis* (6th ed.). Upper Saddle River: Prentice Hall.
Hamman, J. R., Weber, R. A., & Woon, J. (2011). An experimental investigation of electoral delegation and the provision of public goods. *American Journal of Political Science*, *55*(4), 738–752.
Healy, A., & Malhotra, N. (2009). Myopic voters and natural disaster policy. *American Political Science Review*, *103*(03), 387–406.
Healy, A. J., Malhotra, N., & Mo, C. H. (2010). Irrelevant events affect voters' evaluations of government performance. *Proceedings of the National Academy of Sciences*, *107*(29), 12804–12809.
Huber, A., Hill, Seth J., & Lenz, Gabriel S. (2012). Sources of bias in retrospective decision-making: Experimental evidence on voters limitations in controlling incumbents. *American Political Science Review*, *106*(04), 720–741.
Jindapon, P., & Neilson, W. S. (2009). The impact of societal risk attitudes on terrorism and counterterrorism. *Economics and Politics*, *21*(3), 433–451.
Johnston, Patrick B. (2012). Does decapitation work? Assessing the effectiveness of leadership targeting in counterinsurgency campaigns. *International Security*, *36*(4), 47–79.
Kahneman, Daniel. (2011). *Thinking, fast and slow*. New York: Macmillan.
Kean, Thomas. (2011). *The 9/11 commission report: Final report of the national commission on terrorist attacks upon the United States*. Washington, DC: Government Printing Office.
Kibris, A. (2011). Funerals and elections: The effects of terrorism on voting behavior in Turkey. *Journal of Conflict Resolution*, *55*(2), 220–247.
Landes, W. M. (1978). An economic study of US aircraft hijacking, 1961–1976. *Journal of Law and Economics*, *21*(1), 1–31.
Lipton, E. (2010). *Officials regret curbs on adding to terror watch list*. New York: The New York Times.
McDermott, R. (2011). New directions for experimental work in international relations. *International Studies Quarterly*, *55*, 503–520.
Morton, R. B., & Williams, K. C. (2010). *Experimental political science and the study of causality: From nature to the lab*. Cambridge: Cambridge University Press.
Nagel, R. (1995). Unraveling in guessing games: An experimental study. *The American Economic Review*, *85*, 1313–1326.
Ochs, J. (1995). Games with unique, mixed strategy equilibria: An experimental study. *Games and Economic Behavior*, *10*(1), 202–217.
Powell, R. (2007a). Allocating defensive resources with private information about vulnerability. *American Political Science Review*, *101*(04), 799–809.
Powell, R. (2007b). Defending against terrorist attacks with limited resources. *American Political Science Review*, *101*(03), 527–541.
Sandler, T., & Arce, D. G. (2003). Terrorism and game theory. *Simulation and Gaming*, *34*(3), 319–337.
Sandler, T., & Siqueira, K. (2006). Global terrorism: Deterrence versus pre-emption. *Canadian Journal of Economics*, *39*(4), 1370–1387.
Sandler, T., & Siqueira, K. (2009). Games and terrorism. *Simulation and Gaming*, *40*(2), 164–192.
Santifort, Charlinda, & Sandler, Todd. (2013). Terrorist success in hostage-taking missions: 1978–2010. *Public Choice*, *156*(1–2), 125–137.

Shapiro, Jacob N., & Cohen, Dara Kay. (2007). Color bind: Lessons from the failed homeland security advisory system. *International Security*, *32*(2), 121–154.

Stahl, Dale O., & Wilson, Paul W. (1995). On players models of other players: Theory and experimental evidence. *Games and Economic Behavior*, *10*(1), 218–254.

Wolfers, J. (2002). Are voters rational? Evidence from gubernatorial elections. *Working Paper*.

Woon, J. (2012). Democratic accountability and retrospective voting: A laboratory experiment. *American Journal of Political Science*, *56*, 913–930.

Wright, John R. (2012). Unemployment and the democratic electoral advantage. *American Political Science Review*, *106*(04), 685–702.