Useless Prevention vs. Costly Remediation*

Jean Guillaume Forand†

February 19, 2015

Abstract

I model the dynamic agency relationship underlying prevention. In each period, a principal sets a budget for an agent that has private information about a problem, which the agent can direct to solving the problem or divert into rents. Problems are persistent and rectifiable: they randomly generate observable disasters until enough resources have been committed to solving them. I characterise the principal's equilibrium trade-off between (a) preventing disasters while squandering transfers in informational rents to agents facing trivial problems and (b) limiting transfers and remediating costly disasters that eliminate agents informational advantage and prove the need for action.

JEL Classification: D72, H10, C73

Keywords: Agency; Prevention; Dynamic Moral Hazard; Adverse Selection

Had FEMA correctly foreseen the damages generated by Hurricane Katrina, could it have convinced its congressional overseers to finance the costly investments in New Orleans' levees required to avert them? If the levees are updated and withstand the hurricane, then FEMA's counterfactual claim that the old levees would have failed is hard to corroborate, which opens FEMA's managers, and the politicians that monitor

*I thank Aloysius Siow for help in the early stages of this project. I also thank Phil Curry, Johanna Goertz, Francisco Gonzalez, Tasos Kalandrakis, Martin Osborne, Alan Wiseman and audiences at the University of Toronto, the Wallis Institute for Political Economy, MPSA 2011, CPEG 2011 and Stony Brook ICGT 2012. Finally, the Editors and two anonymous referees provided excellent comments and suggestions.

†Department of Economics, University of Waterloo, Hagey Hall of Humanities, Waterloo, Ontario, Canada N2L 3G1. Email: jgforand@uwaterloo.ca. Website: http://arts.uwaterloo.ca/∼jgforand

them, to charges of having unnecessarily diverted public funds. If FEMA fails to take preventative action, then the costly crisis that ensues produces unquestionable evidence that the old levies were weak. However, in this case, FEMA's managers are open to retribution, in the form of career repercussions or public embarrassment, at the hands of their political monitors, angered that the former knew about the problems but remained inactive.

This paper studies agency problems in prevention, which, as highlighted by the example above, are pervasive in the relationships of political principals (e.g., executive office holders, legislative committees) and their agents in the public service (e.g., civil servants, law enforcement officials, central bankers). To exploit such an agent's expertise in identifying and solving problems that can eventually cause a crisis, a political overseer delegates investments in prevention, but retains power over their scale by controlling the resources put at the agent's disposal.[1] In turn, the principal answers to voters or interest groups both for her stewardship of public funds and for any unprevented damages. Because the public employees tasked with prevention belong to large bureaucracies whose preferences can favour organisational growth and prestige over the public interest,[2] this generates a complex incentive problem. On the one hand, if the agent's warnings trigger an allocation of funds by the principal, then by exaggerating trivial problems the agent can capture surplus resources, assured that the corresponding absence of damages will corroborate this successful, but spurious, prevention. If, on the other hand, the principal ignores the agent's advice, then neither useless nor useful prevention occurs, with potentially disastrous, but informative, consequences. Paradoxically, the most credible argument in favour of prevention investments is supplied by failing to undertake them in the first place.

Agency problems in the public provision of prevention are not restricted to natural disaster preparedness. Terrorist attacks have spurred debates on the trade-off between civil liberties and terrorism prevention.[3] However, although law enforcement agencies have priviledged information about the severity of terrorist threats, they can have preferences for stringent laws that encroach on civil liberties, and the resulting adverse

[1]E.g., congressional committees' 'power of the purse', as described by Fenno (1966).
[2]E.g., as described in Niskanen (1971).
[3]See Donohue (2008) and Posner (2005).

selection problem muddies public oversight of these agencies' activities.[4] Financial crises have raised the question of whether central banks should try to preemptively deflate asset price bubbles.[5] Supporters of this position argue that a central bank's role in promoting financial stability suggests preventing bubbles that could lead to crises, while opponents argue that bubbles are hard to identify and that the economic costs of mistaken interventions are high. A central bank may also shy away from taking corrective actions against potential bubbles to limit political fallout harmful to its independence: even if successful, such actions could lead to a recession for which the public and its representatives would blame the bank and not the bubble.

I model an infinite-horizon dynamic game between a political principal and an office-motivated agent in which, critically, the existence of an initial problem is known only to the agent. If a problem exists and it is left untended, it randomly generates costly collateral damages. In each period, the principal sets the budget available to the agent, which can be privately invested in prevention or diverted into rents. The initial problem is persistent, in that any leftover problem from past periods remains in a current period, and rectifiable, in that it can be entirely eliminated given enough resources.[6] Furthermore, disasters are more costly to remediate than to prevent (i.e., 'an ounce of prevention is worth a pound of cure'), so that the principal would devote resources to avoiding damages if convinced of both the existence of a problem and of the reliability of the agent's spending. I focus on the perfect Bayesian equilibria of the game, which limit the principal's ability to commit to future budgets. In particular, the principal would benefit from committing to transfer resources to the agent after a problem has been fully rectified, even though such transfers are sure to be wasted. The constraints on prevention incentives imposed by limited commitment are stringent, because an agent becomes redundant, and hence hard to reward, after successfully preventing a disaster.

An important question is how inefficiencies in prevention investments are linked to the informational frictions in the relationship between the principal and the agent.

[4]See Dragu (2011), Dragu and Polborn (2013) and Di Lonardo (2014).

[5]See Bernanke and Gertler (2001) and Bordo and Jeanne (2002).

[6]Banks and Sundaram (1993, 1998), who present a general model of dynamic moral hazard and adverse selection without commitment, and Biais, Mariotti, Rochet, and Villeneuve (2010) and Myerson (2012), who provide models of dynamic risk prevention with moral hazard, commitment and transferable utility, all consider history-invariant production technologies. See Board and Meyer-ter Vehn (2013), Fernandes and Phelan (2000) and Jarque (2010) for dynamic agency problems with persistence.

In the version of the model with no initial information asymmetry but with moral hazard in prevention, I show that the perfect Bayesian equilibrium which minimises the principal's initial costs involves full prevention in the initial period. In this equilibrium, the timing of prevention investments is efficient and the principal runs no risk of costly disasters. However, the resources devoted to rectifying the problem in the first period are inefficiently bloated, as the principal allows the agent to divert some of this budget into rents. Why does the principal not benefit from distributing its transfers across time in order to limit the agent's rent-seeking in early periods? With persistent problems and no commitment, promising future rents to the agent involves delaying prevention, which imposes additional risks of collateral damages that the principal finds costlier than the savings due to stronger incentives.

It follows that any delays in prevention investments, which are inefficient and generate the possibility of disasters, are tied to the agent's inability to credibly establish the existence of a problem. In particular, collateral damages are informative about the underlying problem, so that the principal can use their arrival as a tool to screen agents by limiting prevention prior to disasters. In the full model in which both problems and prevention are unobservable, I exploit the features of the cost-minimising equilibrium under confirmed initial problems to derive two natural candidates for equilibrium. In the *prevention equilibrium*, the principal limits the rent-seeking of agents facing legitimate problems by implementing the full-prevention assessment under observable initial problems and surrendering the full-prevention transfer to agents not facing any problems. In the *remediation equilibrium*, the principal screens agents' private information by not allocating any budget until collateral damages prove the need for investments. The first arrival of collateral damages puts an end to information asymmetries and the principal implements the full-prevention equilibrium under observable initial problems. Therefore, failing to task the agent with prevention and appearing to simply wait for a disaster need not be due to the principal's short-sightedness or inability to understand the risks of inaction. Rather, it may be a rational attempt to counter the agent's informational advantage.

How the agent is held accountable for the occurrence of damages varies depending on whether the equilibrium calls for prevention or for remediation. In the prevention equilibrium, the principal disregards justification for the spending but punishes the agent for any damages by cutting off access to discretionary office benefits. In the remediation equilibrium, agents are not trusted with any transfers prior to the arrival of damages, in return for which they bear no blame for them. The remediation equilibrium yields lower costs to the principal when problems are sufficiently unlikely, while she prefers the prevention equilibrium when problems are sufficiently likely. This suggests a rationale for the ex post occurrence of damages to be inversely related to their ex ante likelihood: when damages are more likely ex ante, privately informed agents can more credibly make the case for prevention investments that eliminate the risk of damages ex post.

Although a complete characterisation of cost-minimising perfect Bayesian equilibria is difficult, my final result provides some foundations for my focus on the prevention and remediation equilibria. In particular, I show that as the principal becomes arbitrarily farsighted, any cost-minimising equilibrium must yield costs that converge to those of either the prevention or the remediation equilibria. In other words, it is as though a farsighted principal first decides whether to engage in prevention or not. If so, she imposes the equilibrium yielding the least moral hazard rents to agents facing problems, ignoring adverse selection altogether. If not, then she remediates all confirmed problems by imposing the equilibrium offering the least informational rents to agents not facing any problems. Principal farsightedness complements problem persistence and rectifiability, since a farsighted principal knows that the problem must be dealt with eventually and fully internalises the costs of doing so. This leaves no role for partial prevention investments prior to the arrival of damages.

The literature on bureaucratic budgets addresses how political principals manage the informational advantage that agents derive from their expertise to counter unjustified allocations of resources (adverse selection) and/or low-quality policy outputs (moral hazard).[7] Banks (1989), Banks and Weingast (1992) and Bendor, Taylor, and Van Gaalen (1987) focus on the effects of costly monitoring on the principal's ability to curtail these asymmetries.[8] Although the ex post quality of prevention investments can conceivably be audited, their ex ante justification is much more difficult to verify (e.g., an audit can reveal that a security agency is properly applying strong anti-terrorism

[7]How public service institutions bridge this expertise gap also motivates the large literature on communication and delegation in bureaucracies, as surveyed by Gailmard and Patty (2012).

[8]The last assuming commitment, the first two without.

laws, but wether such laws actually thwart attacks can still elicit debate). With prevention, the principal learns the extent of the problem through disasters: in particular, this shields those agents that squander resources on trivial problems from audits. Also, while most models of agency funding are static, the principal's ability to delay prevention investments to acquire information highlights the importance of a dynamic environment. A notable exception is the repeated model of budget allocation with monitoring and moral hazard of Ting (2001).[9]

The growing literature on prevention in models of elections does not deal with bureaucratic incentives, but it highlights the standards of accountability that voters (principals) impose on politicians (agents). Fox and Van Weelden (2013) also investigate a setting in which costly damages may optimally, but inefficiently, go unprevented. In their model, disasters are not prevented because they are unlikely and information about their risk is symmetric, and the agent anticipates that the principal's retention decision is most likely taken when no problem exists. Ashworth and Bueno de Mesquita (2013) present a model in which disasters are exogenous and they generate more precise information about an agent's competence than in normal times, so that the principal optimally applies more stringent retention standards following crises for which the agent bears no blame. Bueno de Mesquita (2007) studies terrorism prevention in a multi-tasking model with moral hazard, focusing on how the principal optimally distorts her retention standard so that agents' under-invest in unverifiable tasks and restrain their rent-seeking. Below, I discuss an electoral extension of my model that preserves all my results.

Model

A principal and an agent interact over an infinite horizon $t = 1, 2, ...$ Their relationship is indexed by an *initial problem* $q_0 \in \{0, \overline{q}\}$, where $\overline{q} \in (0, 1)$. The initial problem is privately observed by the agent, and the principal's belief about the problem's type is given by $\lambda = Pr(q_0 = 0)$. At the start of some period t, let $q_t \leq q_0$ denote whatever part of the initial problem remains at t. Again, only the agent observes the problem q_t,

[9] See also Calvert, McCubbins, and Weingast (1989) for a folk theorem in a repeated budgeting game with complete information.

which randomly generates *collateral damages*, or a *disaster*, $d_t \in \{0, \overline{d}\}$ in each period t, where $\overline{d} > 1$ is the money cost of damages. The likelihood of collateral damages increases in the scale of the underlying problem (e.g., a strong terrorist threat leads to a higher risk of an attack, or poorly maintained levees are more easily breached when subjected to hurricanes), and for simplicity, I assume that $Pr(d_t = \overline{d}|q_t) = q_t$. The agent cannot credibly disclose the existence of an initial problem $q_0 = \overline{q}$, which the principal can only verify by observing collateral damages. Meanwhile, an agent who draws $q_0 = 0$ faces a trivial problem and never oversees any collateral damages.

The principal delegates the prevention of collateral damages to the agent. In each period t, the principal can transfer money resources $\tau_t \geq 0$ to the agent. The agent can divert any amount $\pi_t \in [0, \tau_t]$ of this budget into private rents, and he devotes the balance to prevention investments. Although budgets are observable, prevention expenditures are unobservable. Problems can be reduced, and even eliminated, if the agent invests enough of the resources transferred by the principal into prevention. Specifically, I assume that at the beginning of period $t+1$, the remaining problem is given by $q_{t+1} = \max\{q_t - (\tau_t - \pi_t), 0\}$. At the end of any period t, the principal can irreversibly dismiss the agent. For simplicity, I assume that the principal cannot replace the agent, so that following a dismissal at t, the problem q_t persists, and continues to periodically generate collateral damages, without the possibility of future prevention.

Given rents π_t, the agent's stage payoff in period t is

$$u(\pi_t) + b,$$

where u is a twice continuously differentiable and strictly concave function, and $b > 0$ is a benefit that the agent derives from being retained by the principal. I assume that $u(0) = 0$, $u'(0) = \infty$ and $\lim_{\pi \to \infty} u'(\pi) = 0$. Because the agent's marginal benefit from diverting a small share of any budget is arbitrarily high, the agency costs due to moral hazard are strictly positive: having the agent invest any amount in prevention requires delivering some rents. The principal is risk-neutral and bears all costs associated to transfers and the remediation of damages, but does not pay for the benefit b (or gain from discontinuing it by dismissing the agent). Specifically, given problem q_t, transfer τ_t and rent choice π_t, the principal's expected cost at t is

$$\tau_t + \max\{q_t - (\tau_t - \pi_t), 0\}\overline{d},$$

whether or not the principal dismisses the agent at the end of the period. The principal discounts future payoffs according to $\beta_P < 1$ and the agent according to $\beta_A < 1$.

Remarks

First, note that the marginal cost of fixing problems in period t is 1, and its marginal benefit, which consists of foregone damages in current and future periods, is $\bar{d}/_{1-\beta_P} > 1$. It follows that the costs of collateral damages are minimised when problem are fully resolved in the initial period: in the absence of agency costs, it is more costly to remediate the damages generated by problems than to prevent them.

Second, the temptation for both the principal and the agent to delay costly investments in prevention is derived from the infrequent arrival of collateral damages, so that the existence of a problem can be hidden. In this sense, the unbounded horizon of the game reflects both the flexibility in the timing of prevention investments that a principal may value in her relationship with the agent, as well as the difficulty in motivating the latter to take timely action. However, unresolved problems are persistent, so that damages can only be avoided temporarily (e.g., terrorist groups continue plotting attacks unless intercepted by law-enforcement agencies, or insufficient levies ultimately face a storm severe enough to breach them).

Third, the contractual environment is limited. For one, the principal cannot commit to future budgets. In particular, it is never optimal for the principal to transfer resources to an agent who has fully resolved a problem. Also, the threat of dismissal is a coarse tool to provide incentives. The benefit b represents all advantages that the agent derives from the relationship with the principal that are independent of his ability to divert rents. Many interpretations of what constitutes the dismissal of the agent, with the corresponding loss of access to the stream of benefits, are possible. If the agent is viewed as a bureau, then this can include its loss of prestige, responsibilities, or political capital that leads to a reduced role in policy-making. If the agent is viewed as an individual bureaucrat, then this can include his transfer to another job, demotion, or loss of reputational capital through public embarrassment at the hands of his political principal.

Fourth, the assumption that a dismissed agent cannot be replaced is justified if problems have relationship-specific components and the costs of damages once the agent has been dismissed represent the loss of knowledge or ability in that particular relationship that cannot be compensated by the arrival of a new agent. The assumption is also consistent with various forms of civil service protection. For example, tenure protects individual civil servants and some agencies of the state have various degrees of formal independence from political pressures (e.g., central banks). Given such constraints in the public service, the outside option to the principal and the agent is the breakdown in their working relationship, which, in the model, is what dismissal achieves. Finally, in the Conclusion, I discuss how my results can be extended to a model in which dismissed agents can be replaced.

Strategies and Equilibrium

I focus on (pure strategy) perfect Bayesian equilibria. In many of my results, I focus on the properties of those equilibria which maximise the principal's payoffs (among perfect Bayesian equilibria), which, in this model, is equivalent to minimising the expected discounted cost of dealing with problems.. A *private history at t*, which is known only to the agent, is

$$h_A^t = (q_0, (\tau_1, \pi_1, d_1), ..., (\tau_{t-1}, \pi_{t-1}, d_{t-1})),$$

while a *history at t*, which contains only those components that are publicly observed, namely transfers and realised damages, is

$$h^t = ((\tau_1, d_1), ..., (\tau_{t-1}, d_{t-1})).$$

Given any history h^t, the principal's *retention strategy* $\delta(h^t) \in \{0, 1\}$ indicates whether the principal retains the agent ($\delta(h^t) = 1$) or not at the beginning of period t, and, conditional on retention, the principal's *transfer strategy*, $\tau(h^t) \geq 0$, specifies the resources transferred to the agent. Given any private history h_A^t and any transfer τ_t, the agent's *rent diversion strategy*, $\pi(h_A^t, \tau_t) \in [0, \tau_t]$, specifies the share of the transfer that the agent consumes as rents at t. Finally, given any history h^t, let $\rho(h^t) \in \Delta([0, \overline{q}])$ denote the principal's *belief* about the current level of the problem. An *assessment* $\sigma = (\pi, \tau, \delta, \rho)$ collects strategies and beliefs.

Given a private history h_A^t, its corresponding history h^t, an assessment σ and a transfer τ_t, denote the expected discounted sum of costs to the principal by $V(\sigma; h^t)$ and the expected discounted sum of payoffs to the agent by $U(\sigma, \tau_t; h_A^t)$. Further, for any assessment σ, define $V(\sigma) \equiv V(\sigma; \emptyset)$ as the initial cost to the principal. Abusing notation, let $V(\sigma; q)$ for $q \in [0, \overline{q}]$ denote the principal's costs under assessment σ conditional on the initial problem being q. It follows that

$$V(\sigma) = \lambda V(\sigma; 0) + (1 - \lambda) V(\sigma; \overline{q}),$$

which highlights the principal's trade-off when facing agents that are privately informed about initial problems. The first term, $V(\sigma; 0)$, captures the informational rents paid out to 0-type agents. The second term, $V(\sigma; \overline{q})$, captures (a) the costs of prevention, including the moral-hazard rents paid out to $\overline{q}$-type agents, as well as (b) the costs associated to remediation of damages. These two sources of costs are weighed by the principal's prior λ.

A *perfect Bayesian equilibrium* is an assessment $\sigma^* = (\delta^*, \tau^*, \pi^*, \rho^*)$ such that (a) the principal's dismissal and transfer strategies (δ^*, τ^*) are optimal, i.e., for all histories h^t and all strategies (δ, τ),

$$V(\sigma^*; h^t) \leq V((\delta, \tau, \pi^*, \rho^*); h^t), \tag{1}$$

(b) the agent's rent strategy π^* is optimal, i.e., for all private histories h_A^t, all transfers τ_t and all strategies π,

$$U(\sigma^*, \tau_t; h_A^t) \geq U((\delta^*, \tau^*, \pi, \rho^*), \tau_t; h_A^t), \tag{2}$$

and (c) belief ρ^* is derived from strategies $(\delta^*, \tau^*, \pi^*)$ by Bayes' rule where possible. Now let σ^N be a no-prevention assessment such that, given any private history h_A^t and corresponding history h^t and for all $s \geq t$ and all τ_s, $\delta^N(h^s) = 0$, $\tau^N(h^s) = 0$ and $\pi^N(h_A^s, \tau_s) = \tau_s$. It can be verified that σ^N is a perfect Bayesian equilibrium, and furthermore

$$V(\sigma^N; h^t) = \frac{\mathbb{E}[d_t | \rho(h^t)]}{1 - \beta_P}.$$

Given an equilibrium σ^*, it is useful to consider the following condition: for all histories h^t,

$$V(\sigma^*; h^t) \leq V(\sigma^N; h^t). \tag{3}$$

Note that no perfect Bayesian equilibrium yields higher costs to the principal than the no-prevention profile σ^N, because the principal can unilaterally achieve cost $V(\sigma^N; h^t)$ following any history by dismissing the agent. Therefore, condition 3 is necessarily satisfied in all perfect Bayesian equilibria. Furthermore, to characterise equilibrium payoffs, it is without loss of generality to restrict attention to the set of assessments σ^* satisfying conditions 2 and 3. This follows from standard arguments in repeated games, because any perfect Bayesian equilibrium path can be supported by an assessment that punishes deviations by the principal by reverting to the no-prevention equilibrium.[10] No such simple characterisation can be obtained for the agent's optimality condition 2, since his choice of prevention investment is private.

Observable Initial Problems

How should the principal delegate prevention if initial problems are observable? Because prevention investments remain unobservable, this focuses attention on how the principal confronts moral hazard due to rent diversion when the existence of a problem is not subject to adverse selection. My main result in this section shows that in this case, a perfect Bayesian equilibrium in which all problems are resolved in the first period minimises the principal's costs. This does not mean that the principal bears no agency costs from delegating prevention investments: the agent extracts rents from his ability to misallocate resources, but only in the initial period. That the principal cannot gain from having prevention investments occur over multiple periods may be surprising, since a central insight from models of dynamic moral hazard (with or without commitment) is that the principal can sharpen the incentives she provides to the agent in current periods by paying out rents in the future. However, when problems are persistent and the principal cannot commit to making transfers to the agent after full prevention has occurred, promising future rents necessarily comes at the cost of delaying prevention. I show below that the principal does not benefit from such delayed investments, because any reduction in the agent's rents induces too much additional risk of collateral damages.

This result links the occurrence and efficiency of prevention investments to the informational advantage held by the agent. As long as those subject to possible disasters

[10] E.g., see Section 2.6 of Mailath and Samuelson (2006).

and those they delegate to prevent them have conflicting preferences over the allocation of investments, this spending is wasteful. However, this conflict of interest alone does imply that preventable disasters occur, since it is consistent with complete, although inefficient, prevention. Furthermore, the scale of the underlying problem, when fully understood by both the principal and the agent, is no barrier to prevention investments, because the principal understands that more important problems generate more disaster risk. Rather, as I show below, principals rationally risk the occurrence of costly disasters only when agents' claims to be facing important problems cannot be verified.

A useful first step is to focus on the optimal rent diversion problem of an agent that is tasked by the principal to invest in prevention for the last time. Specifically, consider a history h^t following which (a) the current level of the problem is q_t, (b) a final transfer τ_t occurs at time t, and (c) the principal dismisses the agent following damages occurring in any period $t' \geq t$. The agent's final choice of rents solves the problem

$$\max_{\pi \in [0,\tau_t]} u(\pi) + \frac{b}{1 - \beta_A(1 - \max\{q_t - (\tau_t - \pi), 0\})}. \tag{4}$$

The first term of the objective function captures the agent's stage benefit from diverting rents π from the final transfer τ_t received from the principal and the second term captures his expected retention benefits in current and future periods. Because no prevention occurs in any period following t, $q_{t+1} = \max\{q_t - (\tau_t - \pi), 0\}$ is the final and permanent level of the problem. The following assumption, which is a second-order condition ensuring that the solution to problem 4 is unique, is maintained in the rest of the paper.

Assumption 1. *For all $\pi \geq 0$,*

$$u''(\pi) + \frac{2b\beta_A^2}{[1 - \beta_A]^3} < 0.$$

Note that the agent's expected benefits in the problem from 4 depend on his choice of rents only through the final level of the problem q_{t+1}. Furthermore, the agent's payoff starting from t depends on the problem q_t and the transfer τ_t only through their difference $q_t - \tau_t$: if π^* is the solution to the problem from 4 given q_t and τ_t, then if instead the principal faced problem $q_t + \Delta$, and allocated transfer $\tau_t + \Delta$, the agent's optimal choice of rents would still be π^*, leaving the same final problem level q_{t+1}. In the proof of Proposition 1 in the Appendix, I use this observation to derive a function $\underline{\pi}$ such that, given a targeted final problem $q_{t+1} \in [0, \overline{q}]$, $\underline{\pi}(q_{t+1})$ denotes the rents that must be transferred to the agent to induce the level of prevention that generates this final problem. In other words, given any problem q_t, the minimum cost to the principal at t of inducing the final problem q_{t+1} is the transfer $q_{t+1} - q_t + \underline{\pi}(q_{t+1})$.

An important property of rents $\underline{\pi}(q_{t+1})$, which is key for my results, is that they are increasing in q_{t+1}. In words, agents that leave smaller problems behind also extract less rents.[11] This relies on the expected benefits to the agent in the problem from 4 being convex in the final level of damage q_{t+1}. Since the marginal cost in terms of foregone benefits of diverting transfers into rents is higher for lower targeted final levels of damage, those agents that better internalise the risk of damages imposed by rent-seeking are those agents that leave small problems behind.

The following assumption, which ensures that the principal is willing to supply rents $\underline{\pi}(0)$ to the agent in order to eliminate problem $\overline{q}$, is maintained in the rest of the paper.

Assumption 2.

$$\begin{aligned}\overline{q} + \underline{\pi}(0) &\leq \frac{\overline{q}\overline{d}}{1-\beta_P} \\ &= V(\sigma^N; \overline{q}).\end{aligned}$$

The following result shows that, under Assumption 2, an equilibrium in which all problems are prevented in the initial period minimises the principal's initial costs. This implies that when Assumption 2 fails, σ^N is the only perfect Bayesian equilibrium of the game, and no prevention can be achieved.

Proposition 1. *When initial problems are observable, there exists a perfect Bayesian equilibrium σ^{PR} in which*

i. Problems are fully prevented in the initial period.

ii. The agent receives rents $\underline{\pi}(0)$ in the initial period, and no transfers are made in subsequent periods.

[11]This property simplifies my results, but what is critical is that if the problem q_{t+1} increases, then these rents cannot decline too much relative to the costs imposed by this growing problem.

iii. The principal dismisses the agent if and only if damages occur.

Furthermore, if σ^* is a perfect Bayesian equilibrium, then

$$V(\sigma^*;\overline{q}) \geq V(\sigma^{PR};\overline{q}).$$

The central difficulty in the proof of Proposition 1, detailed in the Appendix, is to show that no equilibrium yields lower initial costs to the principal than σ^{PR}. A first remark is that in all equilibria transfers cease in finite time. To see this, note that in any equilibrium, the sequence of problems $\{q_t\}$ is weakly decreasing and converges to some limit $\hat{q} \in [0,\overline{q}]$, so that in the limit, the principal's costs include a stream of damages that converges to $V(\sigma^N;\hat{q})$. Furthermore, because $u'(0) = \infty$, the agent diverts a positive amount of rents irrespective of how much prevention he is tasked with. Therefore, if she persists in asking the agent to invest in prevention in the limit, the principal's optimality condition 3 must fail. Second, given any targeted final problem level q_{t+1}, the agent's incentives to devote transfers to prevention are sharpest when the principal dismisses the agent following all damages. This implies that, in any cost-minimising equilibrium, continuation play following a final investment in prevention corresponds to that in the reduced continuation game above. That is, in any final period of prevention in period t, rents are minimised (at $\underline{\pi}(0)$) when the problem is fully resolved and the principal dismisses the agent if and only if damages occur. Third, once the resting points of cost-minimising equilibrium prevention are determined, I show that delay is never optimal and the principal bears strictly lower costs if any final period of prevention is brought forward to the initial period.

The following corollary collects comparative statics results for the principal's initial cost under equilibrium σ^{PR}.

Corollary 1. *$V^{PR}(\sigma^{PR};\overline{q})$ is increasing in $\overline{q}$, decreasing in β_A, and is independent of β_P. Furthermore, $\frac{\partial}{\partial \overline{q}} V^{PR}(\sigma^{PR};\overline{q}) = 1$.*

To constrain rent-seeking, the agent's accountability persists after prevention investments have ceased, so that the moral hazard rents $\underline{\pi}(0)$ extracted by the agent are decreasing in its discount factor β_A. These rents are independent of the initial severity $\overline{q}$ of the problem, as they depend only on the continuation problem faced by the agent following prevention, which in turn depends on the final level of the problem anticipated by the principal, which in equilibrium is 0. Under σ^{PR}, all problems are prevented in the initial period, so that the principal's costs do not depend on her farsightedness β_P.

Choosing Between Prevention and Remediation

If the agent's claims about the severity of problems cannot be verified, the principal's dilemma is that the occurrence of a disaster constitutes the only credible proof that prevention is warranted. In this section, I exploit the result that σ^{PR} minimises the principal's costs when the existence of problems is publicly known to construct two candidate equilibrium assessments if problems are privately observed by agents. First, consider an assessment that has full prevention of any problem present in the initial period. By Proposition 1, the cost-minimising equilibrium in this class requires $\overline{q}$-type agents to follow σ^{PR}, the *prevention assessment*. The cost to the principal under σ^{PR} is given by

$$V(\sigma^{PR}) = V(\sigma^{PR}; \overline{q}),$$

and is independent of λ. However, with probability λ, the initial transfer $\tau_1^{PR} = \overline{q} + \underline{\pi}(0)$ is diverted as information rents by 0-type agents. Second, consider an assessment σ^{RE}, the *remediation assessment*, in which there are no transfers (and no prevention) until collateral damages are observed. The agent cannot be held liable for the first occurrence of damages, which screens agent types and proves the need for prevention investments, which are then undertaken according to the cost-minimising equilibrium σ^{PR}. The cost to the principal under σ^{RE} is given by

$$V(\sigma^{RE}) = (1 - \lambda) \frac{\overline{q}[\overline{d} + \beta_P V(\sigma^{PR}; \overline{q})]}{1 - \beta_P(1 - \overline{q})}.$$

Note that because $V(\sigma^{PR}; \overline{q})$ does not depend on the prior λ, $V(\sigma^{RE})$ depends on λ only through the likelihood of collateral damages. Under σ^{RE}, 0-type agents extract no information rents, since transfers are conditioned on verified problems.

Under σ^{RE}, the principal sacrifices prevention to have the arrival of damages eliminate the agent's informational advantage. Under σ^{PR}, the principal acts as though the agent had credibly demonstrated the need for investments to ensure that any existing problem is prevented at the lowest cost. These assessments provide insight into how to resolve the trade-off between prevention and remediation, and the following result details some of their properties.

Proposition 2. *Given any prior* λ

i. σ^{RE} *is a perfect Bayesian equilibrium.*

ii. There exists a prior $\underline{\lambda} \in [0, 1)$ *such that* σ^{PR} *is a perfect Bayesian equilibrium if and only if* $\lambda \leq \underline{\lambda}$.

iii. There exists a prior $\lambda^* \in [0, \underline{\lambda}]$ *such that*

$$V(\sigma^{RE}) \begin{cases} \leq V(\sigma^{PR}) & \textit{whenever } \lambda \geq \lambda^*, \\ \geq V(\sigma^{PR}) & \textit{whenever } \lambda \leq \lambda^*. \end{cases}$$

That σ^{RE} is an equilibrium under all priors λ (part *i*), follows from Proposition 1. In particular, note that under σ^{RE} the principal receives the stage payoff to σ^N until damages occur, followed by the payoff to σ^{PR}, so that condition 3 is satisfied in all histories prior to the arrival of damages. Contrary to σ^{RE}, σ^{PR} is not a perfect Bayesian equilibrium under all priors λ (part *ii*). Under the prevention assessment, condition 3 can be written as

$$\begin{aligned} &\lambda[V(\sigma^{PR}; 0) - V(\sigma^N; 0)] + (1 - \lambda)[V(\sigma^{PR}; \overline{q}) - V(\sigma^N; \overline{q})] \\ &\qquad = V(\sigma^{PR}; \overline{q}) - V(\sigma^N; \overline{q}) + \lambda[V(\sigma^N; \overline{q}) - V(\sigma^N; 0)] \geq 0, \end{aligned}$$

which (a) fails if $\lambda = 1$ (since $V(\sigma^{PR}; 0) - V(\sigma^N; 0) > 0$), ($b$) is satisfied if $\lambda = 0$ (by Assumption 2, strictly if the assumption also holds strictly) and (c) has a lefthand side that is increasing in λ (since $V(\sigma^N; \overline{q}) - V(\sigma^N; 0) > 0$). Therefore, σ^{PR} is an equilibrium if and only if the likelihood that a problem exists is sufficiently high.

Resolving the trade-off between prevention and remediation implies a negative relationship between the ex ante and ex post likelihood of damages (part *iii*). Remediation of damages yields lower initial costs to the principal if the prior λ is high: damages are not likely ex ante and they are not prevented. Intuitively, problems that are hard to verify, such as the impact of climate change or the importance of terrorist threats, are hard to prevent. The principal may prefer to run the risk of catastrophic damages, using these as a call to action and not as evidence of poor behaviour by the agent. On the other hand, if the prior λ is low prevention yields lower costs: damages are likely ex ante but not observed ex post. For problems whose importance is less contentious, prevention is possible, although the principal risks wasteful investments. Note that no such trade-off exists for the agent, because both types prefer prevention to remediation.

The prior λ^* at which the prevention and remediation assessments yield the same initial costs is the unique solution to

$$(1-\lambda^*)\frac{\overline{q}[\overline{d}+\beta_P V(\sigma^{PR};\overline{q})]}{1-\beta_P(1-\overline{q})} = V(\sigma^{PR}). \tag{5}$$

Note that whenever $\lambda \leq \lambda^*$, we have that

$$\begin{aligned} V(\sigma^{PR}) &\geq V(\sigma^{RE}) \\ &\geq V(\sigma^{N}), \end{aligned}$$

and hence σ^{PR} is an equilibrium under λ (yielding $\lambda^* \leq \underline{\lambda}$).

Corollary 2. *The cutoff prior λ^* has the following properties*

i. *λ^* is increasing in $\overline{d}$, b and β_A, and β_P.*

ii. *If β_P is close to 1, then λ^* is decreasing in $\overline{q}$, while if β_P is close to 0, then λ^* is increasing in $\overline{q}$.*

First, the principal risks damages only when opting for remediation, so that λ^* is increasing in $\overline{d}$ (part *i*). Second, by Corollary 1, the principal's costs to prevention are lower whenever the agent has less incentives to divert transfers into rents, so that both $V(\sigma^{PR})$ and $V(\sigma^{RE})$ are decreasing in b and β_A. However, under remediation these cost savings occur later (in expectation), leading the principal to favour immediate prevention. Third, by Corollary 1, the principal's costs do not depend on β_P under prevention, whereas under remediation any eventual damages are more costly to a farsighted principal. Therefore, farsighted principals favour prevention and impatient principals, who have a higher opportunity cost of prevention, favour remediation.

An increase in the initial severity $\overline{q}$ of the problem (part *ii*) induces two countervailing effects. On the one hand, severe problems entail higher risks of damages, giving the principal more incentives for prevention (damage effect). On the other hand, severe problems are more costly to prevent and lead to higher rents for agents facing trivial problems, giving the principal more incentives for remediation (rent effect). Interestingly, although part *i* shows that farsighted principals prefer prevention and impatient principals prefer remediation, they respond differently to increases in the initial severity of the problem. Farsighted principals become more likely to opt for remediation, and impatient principals become more likely to opt for prevention. The key to this result is that whether the damage effect or the rents effect contribute more to the principal's costs following an increase in the potential severity of a problem depends on the extent to which the principal already favours prevention or remediation.

A farsighted principal has a higher willingness to pay out rents to avoid damages and at the margin between prevention and remediation, she opts for prevention even if the existence of a problem is unlikely. In that case, the marginal cost of additional prevention is high as it consists mostly of informational rents. In other words, those principals who invest the most in prevention are most sensitive to increases in agents' private information and are most likely to treat claims that larger investments are required with skepticism. Formally, if β_P is sufficiently close to 1, then differentiating both sides of identity 5 with respect to $\overline{q}$ and evaluating at $\lambda = \lambda^*$ yields that

$$\begin{aligned}\frac{\partial}{\partial \overline{q}}\left[(1-\lambda)\frac{\overline{q}[\overline{d} + \beta_P V(\sigma^{PR};\overline{q})]}{1-\beta_P(1-\overline{q})}\right]\Bigg|_{\lambda=\lambda^*,\ \beta_P=1} &= \frac{V(\sigma^{PR};\overline{q})}{\overline{d} + V(\sigma^{PR};\overline{q})} \\ &< 1 \\ &= \frac{\partial}{\partial \overline{q}} V(\sigma^{PR}),\end{aligned}$$

where the final equality follows from Corollary 1. If the principal is farsighted, delay in the arrival of collateral damages brings no benefits and the costs associated to prevention do not depend the period in which they are incurred. However, under remediation, these costs are paid only with probability $(1-\lambda)$.

On the other hand, an impatient principal has a lower willingness to pay out rents towards prevention, so that at the margin between prevention and remediation, she opts for remediation even if the existence of a problem is likely in the hopes of avoiding damages in the short run. In that case, the marginal cost of additional remediation is high. A similar argument to that above shows that when the principal is impatient, the damage effect dominates, as

$$\begin{aligned} \frac{\partial}{\partial \overline{q}} \left[(1-\lambda) \frac{\overline{q}[\overline{d} + \beta_P V(\sigma^{PR}; \overline{q})]}{1 - \beta_P(1-\overline{q})} \right] \Bigg|_{\lambda=\lambda^*,\ \beta_P=0} &= \frac{V(\sigma^{PR}; \overline{q})}{\overline{q}} \\ &> 1 \\ &= \frac{\partial}{\partial \overline{q}} V(\sigma^{PR}). \end{aligned}$$

An impatient principal trades off the risk of damages in the initial periods under σ^{RE} against the costs of prevention under σ^{PR}. Under cutoff prior λ^*, we have that $V(\sigma^{RE})|_{\beta_P=0} = (1-\lambda^*)\overline{d}\overline{q}$. Because $V(\sigma^{PR}; \overline{q}) = \overline{q} + \underline{\pi}(0)$, the expected level of damages $(1-\lambda^*)\overline{d}$ must be high, in particular strictly greater than 1, for the principal to be indifferent between σ^{RE} and σ^{PR}. In that case, an increase in the severity of the problem leads to a steep increase in the costs of remediation.[12]

Cost-Minimising Equilibria

The prevention and remediation assessments are not, in general, the only candidates for equilibrium. They do, however, impose constraints on the set of cost-minimising equilibria. First, these assessments bound the set of cost-minimising equilibrium payoffs. Specifically, given any cost-minimising equilibrium σ^*, (a) the informational rents of 0-type agents must be less than those they would obtain under the prevention assessment (i.e., $V(\sigma^*; 0) \leq V(\sigma^{PR}; 0)$), and (b) the expected costs imposed by $\overline{q}$-type agents must be less than those they generate in the remediation assessment (i.e., $V(\sigma^*; \overline{q}) \leq V(\sigma^{RE}; \overline{q})$). Indeed, by Proposition 1, (a) follows because the prevention assessment has full prevention and minimises the moral hazard rents of $\overline{q}$-types, and (b) follows because the remediation assessment minimises the informational rents of 0-types. Second, and more importantly, if the principal is sufficiently farsighted, then the prevention and remediation assessments fully characterise the principal's payoffs under cost-minimising equilibria. Specifically, the next result shows that any assessment other

[12]The partial derivative $\frac{\partial \lambda^*}{\partial \overline{q}}$, need not be monotone in β_P, which is why I present only limiting results for sufficiently farsighted or impatient principals.

than the prevention or remediation assessments yields higher costs to a principal that is sufficiently farsighted.

Proposition 3. *Fix any prior* λ *and suppose that, given any* β_P*,* $\sigma^*_{\beta_P}$ *is a perfect Bayesian equilibrium that minimises the principal's initial costs. Then*

$$\lim_{\beta_P \to 1} V(\sigma^*_{\beta_P}) \in \left\{ \lim_{\beta_P \to 1} V(\sigma^{RE}), \lim_{\beta_P \to 1} V(\sigma^{PR}) \right\}.$$

If the principal discounts future costs, then her uncertainty about the existence of a problem can introduce incentives for partial prevention investments. The key intuition underlying Proposition 3, whose proof is in the Appendix, is that a farsighted principal does not benefit from shifting transfers and/or the risk of damages from earlier to later periods, and her costs depend on the total amount of prevention (or the final level of the problem), not its distribution in time. In other words, because a farsighted principal fully internalises the costs of unresolved problems, a high discount factor complements problem persistence and rectifiability. Specifically, for any discount factor the remediation assessment is cost-minimising among those equilibria in which no prevention occurs prior to the arrival of damages, and, conditional on investing in prevention prior to the arrival of damages, a farsighted principal prefers to fully prevent the problem immediately. This last fact relies on two remarks. First, a farsighted principal prefers the remediation assessment to any assessment in which a problem remains unresolved. Intuitively, to a sufficiently farsighted principal, the costs of the damages associated to any level of leftover problem that persists forever dwarf the costs of preventing that problem once it manifests itself through a disaster. Second, a farsighted principal prefers the prevention assessment to any assessment in which problems are fully prevented eventually. Intuitively, a sufficiently farsighted principal anticipates eventually disbursing rents $\underline{\pi}(0)$ (approximately) in full, so that all rent payments that must support earlier, partial, prevention investments are wasted.

Unfortunately, a complete description of all perfect Bayesian equilibria given any discount factor is difficult. As in the proof of Proposition 1, the main difficulty in the proof of Proposition 3 is to characterise arbitrary cost-minimising perfect Bayesian equilibria. Critically, the condition $\beta_P \to 1$ ensures that assessments in which partial prevention occurs prior to the arrival of damages cannot be cost-minimising, and I have not been able to show that this holds in general. As opposed to the case in which the existence of a problem is observable, an impatient principal need not benefit from bringing forward any planned prevention expenditures, because even in the absence of rent considerations, the return to an additional unit of transfer is weighed by the probability $(1-\lambda)$ that these funds are put to good use. Also, assessments with partial prevention bring up difficult issues that my previous results sidestep: future prevention investments provide incentives for the agent's current rent choices, which depends on the details of the particular assessment; the principal learns about the existence of the problem on the equilibrium path and updates her beliefs accordingly, which, because of the absence of commitment, leads to an endogenous sequence of optimality conditions 3; more broadly, because the dynamic game between the principal and the agent is non-stationary and the agent can take private actions that have persistent effects, the usefulness of standard recursive approaches is limited. This complexity frustrates the search for arguments that range over all possible equilibria.[13]

Agent-Appropriated Budgets

Consider the variant of my model in which the agent has the authority to appropriate his budget from the principal. Specifically, fix $\bar{\tau} > 0$ and suppose that, at the beginning of each period t in which he has not been dismissed, the agent can levy transfer $\tau_t \in [0, \bar{\tau}]$ from the principal. Given a private history h_A^t, let $\tau(h_A^t)$ denote the agent's *taxation strategy* (abusing notation), and given an assessment $\sigma = (\delta, \tau, \pi, \rho)$, let $\tilde{U}(\sigma; h_A^t)$ denote the agent's expected discounted sum of payoffs evaluated before his budget choice at t. Furthermore, define the no-prevention equilibrium $\tilde{\sigma}^N$ for this model such that, given any private history h_A^t and corresponding history h^t and for all $s \geq t$ and all τ_s, $\tilde{\delta}^N(h^s) = 0$, $\tilde{\tau}^N(h_A^t) = \bar{\tau}$ and $\tilde{\pi}^N(h_A^t, \tau_s) = \tau_s$. Also, note that any equilibrium σ^* must satisfy

$$\begin{aligned} \tilde{U}(\sigma^*; h_A^t) &\geq \tilde{U}(\tilde{\sigma}^N; h_A^t) \\ &= u(\bar{\tau}) + b, \end{aligned} \tag{6}$$

[13]That being said, I have not succeeded in constructing equilibria that dominate both the prevention and remediation assessments for low discount factors, so that the question of cost-minimising equilibria is open.

because the agent can always obtain a payoff of at least $u(\bar{\tau}) + b$ by appropriating the maximal budget $\bar{\tau}$. Because the agent's budget levy is observable, an argument as detailed above ensures that it is without loss of generality for equilibrium payoffs to restrict attention to the set of assessments σ^* satisfying conditions 2, 3 (with $\tilde{\sigma}^N$ substituted for σ^N) and 6. Finally, note that in the two equilibria σ^{PR} and σ^{RE} considered above, the lowest (equilibrium path) payoff to any agent following any history is $b/1-\beta_A$, which is, for example, obtained by a $\overline{q}$-type agent following successful prevention and by a 0-type agent under σ^{RE} following all histories. Therefore, the assumption that

$$\frac{\beta_A}{1-\beta_A} b \geq u(\bar{\tau}),$$

which ensures that these agents are not willing to sacrifice their future stream of office benefits by a one-time diversion of the maximal budget $\bar{\tau}$, is sufficient to guarantee that all my results extend to the model in which the agent appropriates his budget.

This extension allows the model to address electoral accountability for prevention investments. Indeed, while a budget-setting principal is appropriate for modelling the relationship between political overseers and the bureaucracy, a budget-setting agent is appropriate for modelling the relationship between voters and incumbent politicians, in which retention alone provides incentives for appropriate taxation and prevention investments. By highlighting how the agent's accountability for ex ante prevention and ex post disasters evolves with the tasks that the principal expects him to carry out in equilibrium, my results shed light on seemingly conflicting evidence regarding how voters treat politicians that invest in prevention relative to those that oversee disasters.[14] On the one hand, Healy and Malhotra (2009) find that voters in the U.S. reward the incumbent presidential party for disaster relief spending, but not for disaster preparedness spending. However, Garrett and Sobel (2003) find that nearly half of the funds distributed by FEMA are driven by politics as opposed to need, so that voters have reason to be wary of such investments.[15] On the other hand, voters do punish politicians

[14]See Fox and Shotts (2009), in which politicians are differentiated in both skill and ideology and their incentives to act as delegates (implement policies pandering to voters' tastes) or trustees (implement policies they think best) vary with the importance voters' reelection decisions place on screening incompetent vs. ideologue politicians.

[15]See also Gasper and Reeves (2011) on voter responses to natural disasters, and Coats, Karahan, and Tollison (2006) for evidence of political manipulation of Department of Homeland Security grants for the prevention of terrorism following 2001.

following unfavourable events they are delegated to avoid. For example, Kayser and Peress (2012) show that voters decompose their countries' economic performance into worldwide and national components and hold incumbent politicians accountable for the latter, and Powell and Whitten (1993) show that voters' behaviour is more closely tied to economic conditions when political institutions allow for more precise allocation of responsibility to politicians.

My results also contribute to a recent discussion about whether empirical results of Healy and Malhotra (2009), or those in Achen and Bartels (2004), which show that voters punish politicians for unfavourable events over which politicians have no obvious control, should be interpreted as evidence of voters failing to behave rationally.[16] Ashworth and Bueno de Mesquita (2014) stress that voters' decisions should not to be analysed in isolation, but understood as a response to politicians' anticipated behaviour, showing, in particular, that voters can benefit from committing to a retention standard that is not sequentially rational. In the remediation equilibrium, a voter is justified in failing to reward an incumbent for prevention, because if a disaster occurs, the incumbent is not punished. In contrast, punishing an incumbent for devoting tax dollars to prevention is not appropriate in the prevention equilibrium, because this would destroy the incumbent's ex ante incentives to invest in prevention.

Conclusion

In this paper, I focus on an uninformed principal's trade-off between the useless prevention of trivial problems and the costly remediation of unprevented disasters. I show that moral hazard, due to agents' ability to privately divert the principal's funds, does not imply delays in prevention investments. It is the combination of moral hazard and adverse selection, which is due to agents' not being able to credibly communicate the scale of a problem, that generates the possibility that costly observable damages can precede prevention investments. If the principal is farsighted, the trade-off between prevention and remediation is stark and the principal either opts for full prevention or waits for unprevented problems to generate disasters that prove the need for action.

A number of questions warrant further study. An example is the effect of the arrival

[16] See the surveys by Ashworth (2012) and Healy and Malhotra (2013).

of new problems. On the one hand, the possible existence of numerous problems could make it easier for agents to extract rents by claiming a need for resources to support investments in prevention. However, if in equilibrium the transfers that a principal allocates to agents are rationed, more problems could conceivably help discipline the spurious demands of agents who wish to limit the arrival of damages to remain in good standing.

An important assumption of my model is that the principal cannot replace dismissed agents. A first remark is that all equilibrium outcomes of my model can be supported in a model with agent replacement. If the continuation equilibrium with any new agent is the no-transfer equilibrium σ^N, then the principal is always indifferent between retaining her current agent without allocating any further transfers or opting for a replacement. In such equilibria, all prevention investments, if they occur, must be made by the initial agent, as in my model. Second, even allowing for more general continuation equilibria following replacement, the principal's inability to commit would imply that any equilibrium with replacement would feature a final round of prevention by a final agent, and this agent's incentives would be the same as in my model. It would then follow that my results on cost-minimising equilibria, which rely on prevention investments ending in finite time, would still hold. For example, the arguments establishing that σ^{PR} is cost-minimising when the existence of a problem is known rely on delays in prevention being overly costly for persistent and rectifiable problems. Replacement of agents does not help, as it only introduces a further source of delay.

While my results on cost-minimising equilibria persist with agent replacement, an interesting question concerns the role of competition between agents on the principal's incentives to direct agents towards prevention or remediation in different classes of equilibria. The value of opting for a fresh agent would depend on the continuation equilibrium following replacement, hence the standards of accountability that the principal applies to incumbent agents would be tied to future agents' expected performance.

References

Achen, C. H. and L. M. Bartels (2004). Blind Retrospection: Electoral Responses to Drought, Flu, and Shark Attacks. *Working paper*.

Ashworth, S. (2012). Electoral Accountability: Recent Theoretical and Empirical Work. *Annual Review of Political Science 15*, 183–201.

Ashworth, S. and E. Bueno de Mesquita (2013). Disasters and Incumbent Electoral Fortunes: No Implications for Democratic Competence. *Working paper*.

Ashworth, S. and E. Bueno de Mesquita (2014). Is Voter Competence Good for Voters?: Information, Rationality, and Democratic Performance. *American Political Science Review 108*(03), 565–587.

Banks, J. and R. Sundaram (1993). Adverse Selection and Moral Hazard in a Repeated Elections Model. In *Political Economy: Institutions, Competition, and Representation: Proceedings of the Seventh International Symposium in Economic Theory and Econometrics*, pp. 295. Cambridge University Press.

Banks, J. and R. Sundaram (1998). Optimal Retention in Agency Problems. *Journal of Economic Theory 82*(2), 293–323.

Banks, J. S. (1989). Agency Budgets, Cost Information, and Auditing. *American Journal of Political Science*, 670–699.

Banks, J. S. and B. R. Weingast (1992). The Political Control of Bureaucracies Under Asymmetric Information. *American Journal of Political Science*, 509–524.

Bendor, J., S. Taylor, and R. Van Gaalen (1987). Politicians, Bureaucrats, and Asymmetric Information. *American Journal of Political Science*, 796–828.

Bernanke, B. and M. Gertler (2001). Should Central Banks Respond to Movements in Asset Prices? *American Economic Review 91*(2), 253–257.

Biais, B., T. Mariotti, J. Rochet, and S. Villeneuve (2010). Large risks, Limited Liability, and Dynamic Moral Hazard. *Econometrica 78*(1), 73–118.

Board, S. and M. Meyer-ter Vehn (2013). Reputation for Quality. *Econometrica 81*(6), 2381–2462.

Bordo, M. D. and O. Jeanne (2002). Monetary Policy and Asset Prices: Does 'Benign Neglect' Make Sense? *International Finance 5*(2), 139–164.

Bueno de Mesquita, E. (2007). Politics and the Suboptimal Provision of Counterterror. *International Organization 61*(01), 9–36.

Calvert, R., M. D. McCubbins, and B. R. Weingast (1989). A Theory of Political Control and Agency Discretion. *American journal of political science 33*(3).

Coats, R., G. Karahan, and R. Tollison (2006). Terrorism and Pork-Barrel Spending.

Public Choice 128(1), 275–287.

Di Lonardo, L. (2014). On Electoral Incentives, Terrorist Threats and Counterterrorism Policies. *Working paper*.

Donohue, L. (2008). *The Cost of Counterterrorism: Power, Politics, and Liberty.* Cambridge Univ Press.

Dragu, T. (2011). Is There a Trade-off between Security and Liberty? Executive Bias, Privacy Protections, and Terrorism Prevention. *American Political Science Review 105*(01), 64–78.

Dragu, T. and M. Polborn (2013). The Rule of Law in the Fight Against Terrorism. *American Journal of Political Science 58*(2), 511–525.

Fenno, R. F. (1966). *The Power of the Purse: Appropriations Politics in Congress.* Little, Brown.

Fernandes, A. and C. Phelan (2000). A Recursive Formulation for Repeated Agency with History Dependence. *Journal of Economic Theory 91*(2), 223–247.

Fox, J. and K. Shotts (2009). Delegates or Trustees? A Theory of Political Accountability. *Journal of Politics 71*(4), 1225–37.

Fox, J. and R. Van Weelden (2013). Hoping for the Best, Unprepared for the Worst. *Working paper*.

Gailmard, S. and J. W. Patty (2012). Formal Models of Bureaucracy. *Annual Review of Political Science 15*, 353–377.

Garrett, T. and R. Sobel (2003). The Political Economy of FEMA Disaster Payments. *Economic Inquiry 41*(3), 496–509.

Gasper, J. and A. Reeves (2011). Make It Rain? Retrospection and the Attentive Electorate in the Context of Natural Disasters. *American Journal of Political Science*.

Healy, A. and N. Malhotra (2009). Myopic Voters and Natural Disaster Policy. *American Political Science Review 103*(03), 387–406.

Healy, A. and N. Malhotra (2013). Retrospective Voting Reconsidered. *Annual Review of Political Science 16*, 285–306.

Jarque, A. (2010). Repeated Moral Hazard with Effort Persistence. *Journal of Economic Theory*.

Kayser, M. A. and M. Peress (2012). Benchmarking Across Borders: Electoral Accountability and the Necessity of Comparison. *American Political Science Review 106*(03), 661–684.

Mailath, G. and L. Samuelson (2006). *Repeated Games and Reputations: Long-Run Relationships*. Oxford University Press.

Myerson, R. (2012). Moral Hazard in High Office and the Dynamics of Aristocracy. *Working paper*.

Niskanen, W. A. (1971). *Bureaucracy and Representative Government*. Aldine Press.

Posner, R. (2005). *Preventing Surprise Attacks: Intelligence Reform in the Wake of 9/11*. Rowman & Littlefield.

Powell, G. and G. Whitten (1993). A Cross-National Analysis of Economic Voting: Taking Account of the Political Context. *American Journal of Political Science 37*(2), 391–414.

Ting, M. M. (2001). The "Power of the Purse" and its Implications for Bureaucratic Policy-Making. *Public Choice 106*(3-4), 243–274.

Appendix

Proof of Proposition 1. The first step is to establish the claims made in the text about the outcomes of the reduced continuation game following a final transfer by the principal. Specifically, fix a perfect Bayesian equilibrium σ and consider a history h^t following which (a) the current level of the problem is q_t, (b) a final transfer $\tau_t > 0$ occurs at time t and (c) the agent is dismissed whenever damages occur at time $t' \geq t$. Letting $q_{t+1} = \max\{0, q_t - (\tau_t - \pi)\}$, note that any optimal choice of rents π^* in the problem 4 satisfies $q_{t+1} = q_t - (\tau_t - \pi^*)$. Indeed, if instead $q_t - (\tau_t - \pi^*) < 0$, diverting rents $\pi^* + \epsilon$, where $\epsilon \in (0, \tau_t - q_t]$, would yield strictly higher payoffs to the agent (and the same final problem q_{t+1}), a contradiction. Problem 4 can then be rewritten as

$$\max_{\pi \in [\max\{0, \tau_t - q_t\}, \tau_t]} u(\pi) + \frac{b}{1 - \beta_A(1 - [q_t - (\tau_t - \pi)])}. \tag{7}$$

Denote the objective function of the problem (7) by $f(\pi)$. We have that

$$f''(\pi) = u''(\pi) + \frac{2b\beta_A^2}{[1-\beta_A(1-[q_t-(\tau_t-\pi)])]^3} \tag{8}$$
$$\leq u''(\pi) + \frac{2b\beta_A^2}{[1-\beta_A]^3}$$
$$< 0,$$

where the final inequality follows from Assumption 1. Hence, $f(\pi)$ is strictly concave in π, implying that the problem 7 has a unique solution.

My next claim is that any solution π^* to problem 7 must satisfy $f'(\pi^*) = 0$, or

$$u'(\pi^*) = \frac{b\beta_A}{[1-\beta_A(1-[q_t-(\tau_t-\pi^*)]]^2}. \tag{9}$$

To show this, first note that it must be that $f'(\pi^*) \leq 0$. Suppose instead, towards a contradiction, that $f'(\pi^*) > 0$. Then it must be that $\pi^* = \tau_t$. In that case, $V(\sigma; h^t) = V(\sigma^N; q^t) + \tau_t$, so that by condition 3 transferring $\tau_t > 0$ cannot be optimal for the principal, a contradiction. Second, it must be that $f'(\pi^*) \geq 0$. Suppose instead, towards a contradiction, that $f'(\pi^*) < 0$. Then it must be that $\pi^* = \max\{0, \tau_t - q_t\}$. However, because $u'(0) = \infty$, it must be that $\pi^* > 0$, so that we have $\pi^* = \tau_t - q_t > 0$. Fix $\bar{\epsilon} > 0$ and, for any $\epsilon \leq \bar{\epsilon}$, let π_ϵ^* denote the solution to problem 7 when the final transfer is $\tau_t - \epsilon$. Note that, for all $\epsilon \leq \bar{\epsilon}$, we must have $\pi_\epsilon^* > \tau_t - \epsilon - q_t$. If instead for some ϵ we had $\pi_\epsilon^* = \tau_t - \epsilon - q_t$, then, by transferring $\tau_t - \epsilon$ instead of τ_t following history h^t, the principal's cost would be $V(\sigma; h^t) - \epsilon$, a contradiction. But then, for all $\epsilon \leq \bar{\epsilon}$, we have that $f'(\pi_\epsilon^*) = 0$. However, because u' is continuous, $f'(\pi^*) = 0$, yielding the desired contradiction.

Fix q_t and $q_{t+1} < q_t$, and define $\underline{\pi}(q_{t+1})$ as the unique solution to

$$u'(\underline{\pi}(q_{t+1}, t-t')) = \frac{b\beta_A}{[1-\beta_A(1-q_{t+1})]^2}. \tag{10}$$

If the principal transfers $\tau_t = q_t - q_{t+1} + \underline{\pi}(q_{t+1})$, it follows from equation 9 that the agent's optimal rent choice at t is $\pi^* = \underline{\pi}(q_{t+1})$, and that the final problem is q_{t+1}. Because $u'' < 0$ and the righthand side of equation 9 is decreasing in q_{t+1}, it can be verified that $\underline{\pi}(q_{t+1})$ is increasing in q_{t+1}.

With the above results in hand, I establish the remaining claims of Proposition 1. Define the assessment σ^{PR} such that (a) $\tau_1^{PR} = \overline{q} + \underline{\pi}(0)$ and $\pi_1^{PR} = \underline{\pi}(0)$, (b) given any history h^t such that $\tau_1 = \underline{\pi}(0)$, $\tau_s = 0$ for all $1 \leq s \leq t-1$ and $d_s = 0$ for all $2 \leq s \leq t-1$, we have that $\tau^{PR}(h^t) = 0$, $\delta^{PR}(h^t, 0, 0) = 1$ and $\delta^{PR}(h^t, 0, \overline{d}) = 0$, and (c) for all other histories, $\sigma^{PR} = \sigma^N$. To verify that σ^{PR} is an equilibrium, note that, because σ^N is an equilibrium, σ^{PR} satisfies conditions 2 and 3 following any history in which σ^{PR} calls for continuation play according to σ^N. Second, note that at any history h^t such that $\tau_1 = \overline{q} + \underline{\pi}(0)$, $\tau_s = 0$ for all $1 \leq s \leq t-1$ and $d_s = 0$ for all $2 \leq s \leq t-1$, the principal believes that $q_t = 0$, so that

$$
\begin{aligned}
V(\sigma^{PR}; h^t) &= 0 \\
&= V(\sigma^N; h^t),
\end{aligned}
$$

and condition 3 is satisfied. Also, because profile σ^{PR} calls for continuation play according to σ^N following $(h^t, 0, \overline{d})$ whether the principal dismisses the agent or not, dismissing the agent is optimal for the principal. Third, under σ^{PR}, the principal bears initial cost

$$
\begin{aligned}
V(\sigma^{PR}; \overline{q}) &= \overline{q} + \underline{\pi}(0) \\
&\leq V(\sigma^N; \overline{q}),
\end{aligned}
$$

where the inequality follows from Assumption 2, so that condition 3 is satisfied. Finally, for the agent, the arguments from the reduced continuation game above establish that setting $\pi_1 = \underline{\pi}(0)$ is optimal following the initial (and final) transfer of $\tau_1 = \overline{q} + \underline{\pi}(0)$ by the principal.

To show that no equilibrium yields lower initial costs than σ^{PR}, fix a perfect Bayesian equilibrium σ. First, I claim that all terminal histories involve only a finite number of periods of prevention. Suppose, towards a contradiction, that there exists some terminal history h which involves no final period of prevention. The sequence of problems $\{q_t\}$ associated with this history is monotone and converges to some $\hat{q} \in [0, \overline{q})$. If we let $h^{\hat{t}}$ denote superhistories of h in which prevention occurs, then from condition 3 we have that

$$\begin{aligned}\lim_{\hat{t}\to\infty} \mathbb{E}\Big[\sum_{t\geq\hat{t}} \beta_P^{t-\hat{t}}\tau(h^t)\Big] + V(\sigma^N;\hat{q}) &= \lim_{\hat{t}\to\infty} V(\sigma;h^{\hat{t}}) \\ &\leq \lim_{\hat{t}\to\infty} V(\sigma^N;q_{\hat{t}}) \\ &= V(\sigma^N;\hat{q}),\end{aligned}$$

which holds only if $\lim_{\hat{t}\to\infty}\mathbb{E}\big[\sum_{t\geq\hat{t}}\beta_P^{t-\hat{t}}\tau(h^t)\big]=0$, so that $\lim_{\hat{t}\to\infty}\tau(h^{\hat{t}})=0$. It follows that, for $\hat{t}$ large, the agent's choice of rents $\pi^{\hat{t}}$ solves a problem that is approximately that in problem 4 of choosing a final level of prevention, with the exception that the principal's dismissal strategy δ need not dismiss the agent following some damages occurring in periods $t'\geq\hat{t}$. It follows from standard arguments that the rents that the agent diverts at $\hat{t}$ are lowest when the principal dismisses the agent after damages in all periods $t'\geq\hat{t}$, so that

$$\begin{aligned}\lim_{\hat{t}\to\infty}\pi^{\hat{t}} &\geq \lim_{\hat{t}\to\infty}\underline{\pi}(\hat{q}) \\ &> 0.\end{aligned}$$

But then, for $\hat{t}$ high enough, $\tau(h^{\hat{t}})<\pi^{\hat{t}}$, a contradiction.

The previous claim ensures that to any terminal history h corresponds a superhistory h^f of h along with final period $f\geq 1$ of prevention such that

i. $q_f > q_{f+1} = q_{f+t}$ for all $t\geq 1$, and

ii. $\tau(h^{f+t})=0$ for all $t\geq 1$.

That is, the principal makes a final transfer at time f and the equilibrium calls for no further prevention investments. Let F be the set of final prevention periods for all such histories and fix history h^f with $f\in F$. Following this history, the distribution of disasters is determined by q_{f+1}.

Given any $\epsilon>0$, let $f', f''\in F$ be such that $q_{f'}-\inf\{q_f : f\in F\}<\epsilon$ and $q_{f''+1}-\inf\{q_{f+1} : f\in F\}<\epsilon$. Since ϵ is arbitrary, assume that, for all $f\in F$, $\pi(h_A^f)\geq\underline{\pi}(q_{f+1})\geq\underline{\pi}(q_{f''+1})$. Furthermore, it follows from condition 3 that

$$\begin{aligned}V(\sigma;h^f) &= \tau(h^f)+V(\sigma^N;q_{f+1}) \\ &\leq V(\sigma^N;q_f).\end{aligned}$$

In particular,

$$\begin{aligned}\underline{\pi}(q_{f''+1}) + [q_f - q_{f''+1}] + V(\sigma^N; q_{f''+1}) &\leq \underline{\pi}(q_{f+1}) + [q_f - q_{f+1}] + V(\sigma^N; q_{f+1}) \\ &\leq \pi(h_A^f) + [q_f - q_{f+1}] + V(\sigma^N; q_{f+1}) \\ &= V(\sigma; h^f), \qquad (11)\end{aligned}$$

where the first inequality follows because $[q_f - q_{f+1}] + V(\sigma^N; h^{f+1})$ is strictly increasing in q_{f+1} since $\overline{d} > 1$. In words, the equations in 11 state that at problem level q_f, holding prevention prior to h^f fixed, the principal gains by according transfer $\underline{\pi}(q_{f''+1}) + [q_f - q_{f''+1}]$ to decrease the problem to minimal $q_{f''+1}$, since in equilibrium it is willing to allocate more transfers in order to achieve a higher permanent problem level. By inequality 11, the principal's initial discounted cost is such that

$$\begin{aligned}V(\sigma; \overline{q}) \geq \overline{q} - q_{f'} + \mathbb{E}_{f \in F}\Bigg[&\frac{(1-\beta_P)}{1-\beta_P^{f+1}} V(\sigma^N; q_{f'}) \\ &+ \beta_P^{f+1}\left[\underline{\pi}(q_{f''+1}) + [q_{f'} - q_{f''+1}] + V(\sigma^N; q_{f''+1})\right]\Bigg] \\ \geq [\underline{\pi}(q_{f''+1}) + [\overline{q} - q_{f''+1}]] &+ V(\sigma^N; q_{f''+1}),\end{aligned}$$

In words, the first inequality states that the principal's initial discounted cost $V(\sigma; \overline{q})$ is at least the cost to paying $\overline{q} - q_{f'}$ in the initial period to attain the stage payoff to equilibrium σ^N at problem level $q_{f'}$ for some (random) number of periods followed by a final move to problem $q_{f''+1}$ at cost $\underline{\pi}(q_{f''+1}) + [q_{f'} - q_{f''+1}]$, and the second inequality states that the principal gains relative to σ by moving all final prevention forward to the first period. It remains to determine which permanent problem achieves the lowest cost for the principal. For this, note that $\underline{\pi}(q) + [\overline{q} - q] + V(\sigma^N; q)$ is minimised at $q = 0$, since both $\underline{\pi}(q)$ and $[\overline{q} - q] + V(\sigma^N; q)$ are increasing in q, so that setting the final damage to 0 in the initial period minimises principal's costs.

□

Proof of Corollary 1. This result depends on the comparative statics of $\underline{\pi}(0)$, which, following equation 10, is determined by

$$u'(\underline{\pi}(0)) = \frac{b\beta_A}{[1-\beta_A]^2}.$$

It can be verified that $\underline{\pi}(0)$ is decreasing in b and β_A. □

Proof of Proposition 3. The result follows from this claim: *Let σ be a perfect Bayesian equilibrium. If β_P is sufficiently close to 1, then $V(\sigma) \geq \min\{V(\sigma^{PR}), V(\sigma^{RE})\}$.*

To show this, let σ be a perfect Bayesian equilibrium. If under σ no transfers (and no prevention) occur following any history in which no damages have been observed, then, by Proposition 1, we have that $V(\sigma) \geq V(\sigma^{RE})$.

Now suppose that there exists some history h^T in which no collateral damages have been observed and for which $\tau(h^T) > 0$ and, for all $t > T$, $\tau(h^t) = 0$ for all subhistories h^t of h^T in which no damages have been observed. That is, history h^T captures the final period of prevention prior to a first observation of damages called for under profile σ. First, suppose that under σ, we have that $q_{T+1} > 0$ for agents of type $\overline{q}$, and assume that β_P is large enough that $V(\sigma^{PR}; q_{T+1}) \geq V(\sigma^N; q_{T+1})$. Then

$$\begin{aligned}
V(\sigma) \geq{} & \lambda \sum_{t=1}^{T} \beta_P^{t-1} \tau(h^t) + (1-\lambda) \Bigg[\frac{1-\beta_P^{T-1}(1-q_{T+1})^{T-1} q_{T+1}}{1-\beta_P(1-q_{T+1})} \left[\overline{d} + \beta_P V(\sigma^{PR}; \overline{q})\right] \\
& + \beta_P^{T-1}(1-q_{T+1})^{T-1} \Bigg[[\tau(h^T) + [\overline{q} - q_{T+1}] \\
& + \frac{q_{T+1}}{1-\beta_P(1-q_{T+1})} \left[\overline{d} + \beta_P V(\sigma^{PR}; q_{T+1})\right] \Bigg] \Bigg] \\
\geq{} & \lambda \beta_P^{T-1} \tau(h^T) + (1-\lambda) \left[\frac{q_{T+1}}{1-\beta_P(1-q_{T+1})} \left[\overline{d} + \beta_P V(\sigma^{PR}; \overline{q})\right] \right] \\
\to_{\beta_p \to 1}{} & \lambda \tau(h^T) + V(\sigma^{RE}) \\
>{} & V(\sigma^{RE}).
\end{aligned}$$

The second inequality follows from arguments in the proof of Proposition 1, which imply that

$$\begin{aligned}
\tau(h^T) + [\overline{q} - q_{T+1}] + V(\sigma^{PR}; q_{T+1}) &\geq \underline{\pi}(0) + [\overline{q} - q_{T+1}] + V(\sigma^{PR}; q_{T+1}) \\
&= 2\underline{\pi}(0) + \overline{q} \\
&> V(\sigma^{PR}; \overline{q}).
\end{aligned}$$

In words, if under σ the problem is not fully resolved before the arrival of collateral damages, then a farsighted principal is better off under σ^{RE}.

Second, suppose that under σ, we have that $q_{T+1} = 0$ for agents of type $\overline{q}$. Then

$$\begin{aligned}
V(\sigma) \geq \lambda \sum_{t=1}^{T} \beta_P^{t-1} \tau(h^t) + (1-\lambda)\Bigg[& \frac{1-\beta_P^{T-1}(1-q_{T+1})^{T-1} q_{T+1}}{1-\beta_P(1-q_{T+1})} \left[\overline{d} + \beta_P V(\sigma^{PR}; \overline{q})\right] \\
& + \beta_P^{T-1}(1-q_{T+1})^{T-1} V(\sigma^{PR}; \overline{q}) \Bigg] \\
\geq \lambda \beta_P^{T-1} V(\sigma^{PR}; \overline{q}) + (1-\lambda)\Bigg[& \frac{1-\beta_P^{T-1}(1-q_{T+1})^{T-1} q_{T+1}}{1-\beta_P(1-q_{T+1})} \left[\overline{d} + \beta_P V(\sigma^{PR}; \overline{q})\right] \\
& + \beta_P^{T-1}(1-q_{T+1})^{T} V(\sigma^{PR}; \overline{q}) \Bigg] \\
\to_{\beta_p \to 1} V(\sigma^{PR}) & + (1-\lambda)[1-(1-q_{T+1})^{T-1}]\overline{d} \\
> V(\sigma^{PR}). &
\end{aligned}$$

The second inequality follows because, given that $q_{T+1} = 0$, we have that

$$\begin{aligned}
\sum_{t=1}^{T} \beta_P^{t-1} \tau(h^t) &\geq \beta_P^{T-1}[\overline{q} - q_{T+1} + \tau(h^T)] \\
&\geq \beta_P^{T-1}[\overline{q} + \underline{\pi}(0)] \\
&= \beta_P^{T-1} V(\sigma^{PR}; \overline{q}).
\end{aligned}$$

In words, if under σ the problem can be fully resolved before the arrival of collateral damages, then a farsighted principal is better off under σ^{PR}.

□