Propping Up Allies With Uncertain Prospects

Kristopher W. Ramsay* Joseph J. Ruggiero†

September 1, 2026

Abstract

We study a model where a belligerent depends on continued support from an ally to fight a war. Both parties have aligned preferences in war, but the donor prefers to withdraw support if the prospects of winning are too slim. To regulate their ally's beliefs about their ability to win, the recipient may deviate from optimal battlefield strategies in equilibrium. This can occur even when the recipient is strong and victory is likely. In fact, greater fighting ability can make it more difficult for the strong belligerent to fight optimally because setbacks are even more likely to imply weakness. Allowing the donor to occasionally observe conditions on the ground can make both strong and weak recipients worse off. Limited commitments also do not resolve the problem created by the conditionality of support. For any finite commitment length, a sufficiently long war between patient states induces the same behavior as a game with no commitment.

Keywords: alliance politics, formal theory, delegation, reputation, adverse selection, repeated games

*William P. Boswell Professor, Department of Politics, 303 Fisher Hall, Princeton University, Princeton, NJ 08544. (kramsay@princeton.edu)

†Assistant Professor, Department of Politics, University of Virginia, P.O. Box 40078, Charlottesville, VA 22902. (jruggiero@virginia.edu)

1 Introduction

In 2021, the continuation of United States support to the Afghan government grew increasingly uncertain. At the same time, Afghan commanders were sending their most capable forces, the Commandos, to reinforce threatened positions and recapture territory lost to the Taliban.[1] These operations came at exceptional cost, with commanders regularly ignoring systems in place that were designed to provide troops with rest and recovery.[2] But in addition to control over territory, Afghan leaders were fighting for external aid from the US. Accepting setbacks could signal weakness to the US donors and jeopardize the continuation of support. On the other hand, pressing forward could signal strength even when the best long-term military choice would be to avoid extreme losses, preserve forces, and concede ground.

The tradeoff facing Afghan commanders reflects a general problem in wars where the belligerents depend on external support from an ally. States often have interests in foreign wars, and when such interests are compelling, states may offer support, such as weapons, equipment, and financial aid, to their preferred side in hope of influencing the outcome. An aligned state may be willing to extend support as long as their perception of the war is favorable, backing out if they determine that success is too unlikely to justify the ongoing costs. However, potential donors often lack information on the ground, such as whether territory gained is worth the losses incurred to recover it, and need to depend on the recipient's tactical choices to form expectations. As a result, when deciding how to fight, belligerents receiving external funds consider how the donor will view their actions. If the optimal tactics suggest bad news and cause donors to second-guess their involvement, recipients must choose between fighting effectively today and securing support for tomorrow.

In this paper, we propose a model to understand the incentives in the case where a belligerent's ability to fight depends on the support of an external ally. Both parties have aligned preferences in the war outcome, but the donor prefers to withhold financing if the odds of victory are too slim. Moreover, the belligerent has private knowledge of the front lines that informs their tactics. When circumstances are not conducive to successful attacks (when it *rains*), it is optimal to maintain position with defense. On the other hand, when

[1]Refer to SIGAR (2023).

[2]Refer to United States Department of Defense (2020).

an opportunity arises (it *shines*), it is better to press forward with offensive attacks. We allow for this *weather* to be endogenous to the quality of the recipient: a strong belligerent is more likely to be able to advance their offensive goals on any given day. Consequently, battlefield actions not only affect outcomes, but also donor beliefs about the potential to win and hence whether they receive future external support. When the donor cannot perfectly monitor conditions on the ground, this can generate an incentive to deviate from optimal tactics, even though it is against both donor's and recipient's war aims to do so.

We find that both strong and weak recipient belligerent deviate from optimal military tactics in equilibrium when their continued support is sufficiently uncertain (Theorem 1). When the donor is confident that the recipient can win, the recipient has freedom to accept occasional setbacks on the battlefield without fear of losing resources going forward. On the other hand, if funding is insecure, one big setback may be the last straw. As a result, belligerents that are highly dependent on external support may need to avoid setbacks at all costs, even if it places them in a worse position to win the war (e.g., allowing troops to take on too much damage in refusing to concede a small amount of territory). Counterintuitively, because an occasional setback becomes increasingly informative of weakness when strong recipients are expected to succeed more often, greater strength can make it harder for the recipient to follow optimal battlefield tactics (Proposition 1).

The recipient's incentive to engage in this distortionary behavior depends on the donor's ability to monitor battlefield conditions. Counterintuitively, we find that limited monitoring can make both weak and strong recipients worse off (Proposition 2). With no monitoring, the recipient can always prevent a degradation of the donor's beliefs about their prospects by attacking regardless of the day's weather. On the other hand, monitoring allows the donor to occasionally observe unfavorable conditions on the ground despite these obfuscating recipient strategies. Even a strong recipient can, through bad luck, experience the occasional rain and, if the donor happens to observe them, it can cause them to withdraw their support. Limited monitoring can thus increase the risk of losing support without mitigating the incentive to distort tactics. In contrast, when conditions on the ground can be perfectly monitored and bad news cannot be concealed, hence the incentive to distort disappears.

Additionally, our analysis shows that limited commitments do not mitigate the problem caused by the conditionality of support (Proposition 3). While such a commitment can prevent the donor from withdrawing support in the short term, the recipient's actions continue to affect the donor's beliefs about their prospect of winning. When the commitment period ends, the donor bases their decision to continue supporting the recipient or not on those beliefs. As a result, the recipient retains the incentive to conceal bad news from the donor even when the commitment is active and funding is secure. In fact, we show that commitments of any finite length fail to change recipient behavior in long enough wars involving patient enough states (Proposition 4).

These results have empirical implications on how we understand strategic interactions surrounding external support in war. First, the main mechanism of distorted battlefield tactics is illustrated in the case of US support in Afghanistan from 2015-2021. Insecurity about continued assistance met with a constrained US intelligence on the ground provided Afghan commanders an incentive to demonstrate strength even when doing so required overly costly offensive operations that were counterproductive to their overall aims in the war. Second, in light of our results on the role of monitoring, we examine lack of transparency between the US and Ukraine from 2022-2024. Our theory provides an alternative explanation for this behavior, focusing our attention on the beliefs of the donor as a causal determinant rather than the quality or strength of the recipient. Finally, our commitment results can shed light on why limited commitments of support to the Soviet Union during World War II could remain ineffective even while they were honored, as US assessments of Soviet prospects continued to evolve during the commitment period, ultimately shaping how decisions over subsequent support would be made.

1.1 Related literature

Our paper makes theoretical contributions to the literature on the dynamics of conflict, particularly that on the role of resources and information. Most importantly, we show how dependence on continued external support changes the recipient's battlefield tactics when those tactics affect what the donor learns about the war. Related work in this domain has studied belligerent learning while fighting (Powell, 2004), moral hazard in alliance contracts (Benson, Meirowitz, and Ramsay, 2014), and the role of markets in war (Krainin et al., 2022). Our focus on dependence on external resources also connects the paper to work on state capacity, weak states, and foreign financing (Dube and Naidu, 2015; Fearon and Laitin, 2003; Besley and Persson, 2010). A related literature studies the effect of third-party intervention on conflict duration, finding that intervention can reduce the duration of war and facilitate quicker stalemates (Beardsley, 2012), while external support can prolong civil wars (Cunningham, 2010; Regan, 2002), though it depends on the fungibility of the support (Sawyer, Cunningham, and Reed, 2017).

Our model is most closely related to two recent articles that also study asymmetric information and learning in the context of war fighting. First, Krainin et al. (2020) develops a bargaining model where a government faces uncertainty about rebel strength that deteriorates over time, showing how this can lead to rational quagmires where fighting continues despite repeated failures. Their model also features dynamic learning about military capabilities through battles, but differs in several key respects. Most importantly, the players' preferences are opposed and they do not incorporate third-party intervention that can affect the learning process. Second, Grillo and Nicolò (2025) shares our focus on learning and third party intervention; however, it differs in how it conceives of intervention. Their model features direct military aid and sanctions affecting the attacker's learning about their own strength. Our model examines when allies strategically withhold or provide support, and how the recipient's choices influence the donor's beliefs about winning prospects, even when the ally donor and recipient have aligned preferences in the conflict.

We focus on a conflict setting, but our results also speak to formal models of political delegation more broadly. In particular, models of American politics have studied how an agent's behavior today affects whether a principal can continue to delegate to them tomorrow. For example, Gailmard and Patty (2007) study a model in which the legislature repeatedly decides whether to allow a bureaucrat to make policy, and hence the bureaucrat considers the legislature's future decisions when making their policy choices. Additionally, Gailmard (2021) shows how an agent may make bargaining concessions to a third party to conceal weakness when observed policy outcomes affect the principal's retention decision. We instead focus on a principal and agent that agree on the best action after delegation, but distortion nonetheless arises because the agent's behavior affects the principal's willingness to continue delegating in the future.

Finally, the paper also relates to work in economic theory on dynamic principal-agent problems. Studies of repeated moral hazard have examined how an ongoing relationship can provide incentives when the principal cannot directly observe or control the agent's behavior (Spear and Srivastava, 1987; Padró i Miquel and Yared, 2012). Moreover, Lipnowski and Ramos (2020) model repeated delegation when to an agent that privately observes the payoff-relevant conditions but has preferences that systematically differ from the principal's. We study a model of repeated delegation in which we add a separate persistent dimension of private information, with the distribution of the payoff-relevant conditions the agent privately observes being governed by their private type, and the agent's choices therefore potentially revealing information about that type.

2 Model

Two countries, a donor and a recipient, play a repeated game. Time is discrete and indexed by $t = 1, 2, \ldots, T$, with $T > 1$ finite or infinite, and future payoffs are discounted by $\delta \in (0, 1)$. At the beginning of the game, the recipient's type $\theta \in \{0, 1\}$ is determined by Nature according to prior $\mu_1 = \Pr(\theta = 1) \in (0, 1)$, with 1 denoting strong and 0 weak. The recipient knows their type but the donor does not.

At the beginning of each period t, the donor chooses whether to fund the recipient. If the donor does not fund, denoted by $d_t = 0$, then the donor receives a series of flow payoffs of $c > 0$ for the remaining $T - t + 1$ periods and the recipient receives zero.

On the other hand, if the donor funds, $d_t = 1$, Nature then determines whether the recipient will have an opportunity for a successful attack or not, $\omega_t \in \{0, 1\}$, which we refer to as *rain* or *shine*. Specifically, the probability of shine s_θ depends on the recipient's private type, with the assumption that strong types generate more opportunities, $s_1 > s_0$. While s_0 and s_1 are known to both players, the realized *weather* ω_t is only known to the recipient. Having observed the weather, the recipient then chooses to defend or attack, $r_t \in \{0, 1\}$ where 1 indicates attack and 0 indicates a defense. For both players, attacking in the shine yields a flow payoff $a > 0$, defending in the rain yields $b > 0$, and players receive zero otherwise. We say that a recipient that maximizes their period t flow payoff, $r_t = \omega_t$, is playing *sincerely* in that period. At the end of each period, the flow payoff is publicly observed by the donor with probability $\lambda \in [0, 1]$. Players receive their flow payoffs whether or not the donor observes them, and thus monitoring affects only the donor's information. We refer to λ as the donor's monitoring capacity, and we are especially interested in strategic interactions with low monitoring.

We also make the following assumptions on the model primitives.

Assumption 1 (Payoffs)**.** *Given δ, s_0, s_1, the payoffs a, b, c satisfy*

$$s_1 a > c, \tag{1}$$

$$\frac{(s_1 a - c)(c - s_0 a)}{a(1 - s_0 s_1) - c(2 - s_0 - s_1)} > b \tag{2}$$

$$\delta(s_0 a + (1 - s_0)b) > b. \tag{3}$$

Inequality (1) of Assumption 1 guarantees both that donors want to fund strong types, and that donors would still prefer to fund strong types even if they never defend in the rain. Inequality (2) bounds the defense payoff b from above so that, from any belief at which the profitability of funding depends on the recipient's tactics, a single defense revealing the rain leaves the donor unwilling to fund even a sincere recipient. Finally, inequality (3) ensures that the recipient would prefer a payoff of zero today with funding tomorrow to a payoff of b today and the end of funding. Together, these conditions essentially guarantee that a is sufficiently large and b is sufficiently small to generate the strategic interaction of greatest interest, where the donor does not want to fund a weak type and the weak type the recipient does not want to forfeit the promise of tomorrow's funding for an immediate gain today.

Assumption 2 (Reputational Incentives)**.** $s_0 + 1/\delta > 2$.

In Assumption 2, we require that the weak type's probability of favorable shine conditions is sufficiently large given the common discount factor. Rearranged, the inequality reads $b > \delta(1 - s_0)b/(1 - \delta)$, i.e., the value of defending in the rain today is greater than the expected discounted value of defending in every future rainy period. We show that, given support continues, the largest possible difference in the recipient's future value across donor beliefs is bounded by the value of defending in future rainy periods (Lemma 7, Appendix C.1). The assumption therefore ensures that no reputational gain justifies sacrificing a defense today unless the donor's continued support is at stake. This serves to isolate the effect of the conditionality of support, as recipients in our setting never distort their tactics merely to appear strong, and any insincere play relates directly to the risk of losing the donor's support. If the condition does not hold, a better reputation could be worth more than a defense even when the donor's support is not at risk, and deviations from sincere play would no longer indicate that support is in doubt.

Timing

To summarize, the timing is as follows. The stage game is illustrated in Figure 1.

1. Nature draws the recipient's strength, $\theta \in \{0, 1\}$, observed by the recipient and for which the donor has a well-specified prior belief, μ_1.

2. The donor stays out or funds, $d_t \in \{0, 1\}$. If they stay out, the game ends.

3. Nature rains or shines $\omega_t \in \{0, 1\}$ with probs. $1 - s_\theta$ and s_θ, respectively, where $1 > s_1 > s_0 > 0$.

4. The recipient observes the weather, then plays defends or attacks, $r_t \in \{0, 1\}$.

5. The donor updates beliefs by Bayes' rule given the recipient's action and flow payoffs if they observe them (which occurs with prob. λ).

   If $t < T$, period $t + 1$ proceeds from step 2. If $t = T$, total payoffs are realized.

We restrict attention to strategies that are Markov in the donor's belief μ and, when $T < \infty$, the number of periods remaining. When $T = \infty$, strategies depend only on the current belief. For every belief at which funding is part of the donor's strategy, posteriors follow from Bayes' rule whenever possible. We also assume initial priors that do not immediately end the game (otherwise there is no interaction to study). We refer to any such Markov perfect Bayesian equilibrium simply as equilibrium and focus on equilibria with the following monotone properties.

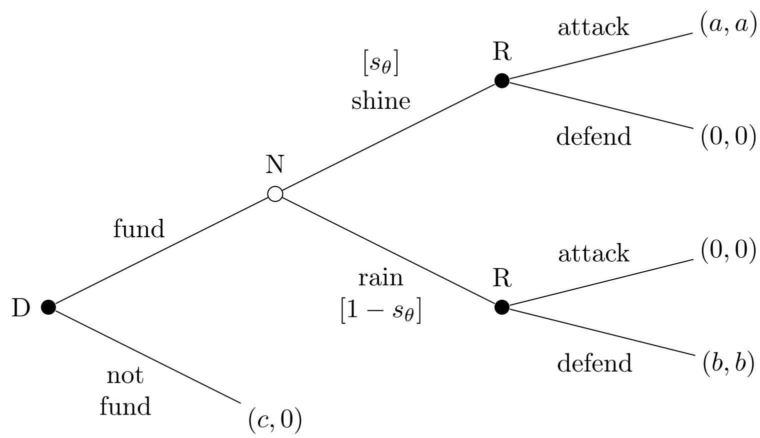

Figure 1: Stage game between a donor (D) and a recipient (R) of type θ. *Note:* Nature (N) chooses the weather. The payoffs list the donor's flow payoff first and the recipient's second.

Definition 1 (Monotonicity)**.** *An equilibrium σ is called monotone if*

(i) at every information set, the donor's posterior belief about recipient type is weakly increasing in the recipient's action for each monitoring realization,

(ii) for each type θ of the recipient, the probability of defending in the rain is weakly increasing across nondegenerate beliefs, and

(iii) the donor funds if and only if their current belief μ is sufficiently large, $\mu \geq \bar{\mu}_t \in [0, 1]$ for some threshold that may depend on time.

Definition 1 pins down the equilibria in which attacking looks strong, avoiding ones where strong types can signal strength by defending, for example. The donor reads an attack as good news about strength by part (i), and funds those recipients they sufficiently believe in by part (iii). Additionally, by part (ii), each recipient is more likely to distort their battlefield tactics as their reputation becomes less secure.

2.1 Comments on the model

The donor and the recipient want the same things in the war, and they disagree only in that the donor holds an outside option worth c each period. Every strategic distortion in our results thus arises from the funding relationship itself, not from divergent war aims. This is a deliberate choice that allows the model to isolate how the conditionality of support affects fighting behavior on the battlefield.

The flow payoffs a, b, and c admit a broader interpretation than fixed returns. In any period, an attack launched in the shine may resolve into a distribution over battlefield outcomes, such as a territorial advance, human casualties, or equipment destroyed. Then, a can be read as the expected value of that distribution, incorporating the probability that the offensive succeeds outright as well as the probability that it fails at great cost. The payoffs b and c admit a similar reading, with each period's flow payoff standing in for the expected draw from the distribution of possible events. Under non-zero monitoring, this interpretation requires that the donor's observation of a realized outcome still reveals the weather.

We think of the large, finite T horizon as the most practical version for the application. Wars of external support cannot go on forever: appropriations expire, stockpiles deplete, elections change leaders, and opponents lose the capacity to continue. However, analyzing every horizon disciplines the mechanism. The equilibrium takes the same threshold form whether the horizon is short, long, or infinite, and the incentive to distort battlefield tactics is present throughout.

Finally, the weather ω_t reflects conditions on the battlefield that the recipient observes directly but the donor is not guaranteed to learn, such as opportunities for attack, the state of the recipient's units, or general morale. A recipient whose funding depends on the donor's beliefs has an incentive to misrepresent these conditions. Any observer can watch the front line move, but gaining or losing territory is not the same as winning or losing the war. An offensive that gains ground while exhausting troops and equipment can leave the recipient worse off than no offensive at all, and a defense that concedes small losses today but preserves the force for tomorrow can be a tactical success. Judging these gains and losses requires information that the recipient alone has, and even well-resourced observers get wrong.[3] We model donors as being able to recover information that reveals these short-term gains and losses with monitoring probability λ. Because of our interest in the effect of support's conditionality on tactics, we are especially interested in situations with low to no monitoring.

[3]For example, analysts widely debated whether the 2023 Ukrainian counteroffensive was a success, despite agreement on how the front line had moved (Kofman and Lee, 2023).

3 Equilibrium

In this section, we characterize the monotone equilibria of interest. Sincere play returns both players the greatest flow payoff, with the donor preferring it conditional on funding today and the recipient preferring it conditional on funding in the future. However, as we will see, the donor's inability to commit to future funding creates an incentive for the recipient to deviate from sincere play.

Before the main result, we note that the donor's funding threshold in any monotone equilibrium weakly declines over time when monitoring is rare (Proposition 6, Appendix C.6). To see why, suppose the donor funded a belief below the next period's funding threshold. An attack that both types make carries no information, the belief either stays where it is or reveals the rain, and support ends regardless. With no reputational concerns, the recipient would rather defend in the rain and collect b. At higher funded beliefs, where an attack does preserve funding, the recipient attacks in the rain to keep it. Defending in the rain would then be more attractive at the lower belief than at the higher one, which reverses part (ii) of Definition 1. No monotone equilibrium can fund such a belief, and thus the funding threshold weakly declines over time.

Next, consider the recipient at beliefs near the funding threshold. When play is sincere, a defense reveals the rain, and the donor's belief falls to a posterior we denote $\mu^-(\mu)$. Then, a recipient who defends at a belief below $(\mu^-)^{-1}(\bar{\mu}_{t+1})$ drops the donor's posterior below the next funding threshold and forfeits all future support. Whether and how this danger shapes strategic behavior is captured by the following result.

Theorem 1. *Fix $T \leq \infty$. There exists a $\lambda^* \in (0,1)$ such that, for all $\lambda < \lambda^*$,*

(i) a monotone equilibrium exists;

(ii) in every monotone equilibrium, there exists a threshold $\bar{\mu}_t \in (0,1)$ such that the donor funds in period t if and only if $\mu \geq \bar{\mu}_t$;

(iii) in every monotone equilibrium, there exists a threshold $\mu_t^* \in [(\mu^-)^{-1}(\bar{\mu}_{t+1}), 1]$ for all $t < T$, with $(\mu^-)^{-1}(\bar{\mu}_{t+1}) > \bar{\mu}_t$, such that both types of recipient play sincerely for any $\mu > \mu_t^*$ and pool on attack for any $\mu \in [\bar{\mu}_t, \mu_t^*)$ in each period t.

In every monotone equilibrium, we see the same general pattern across periods: if the donor's current belief is sufficiently high, the recipient will fight sincerely on the battlefield, confident they will continue to receive support. If the donor's current belief is sufficiently low, they will simply pull funding. In between these two thresholds, the donor will continue to fund but the recipient will distort their battlefield tactics by attacking in every period, regardless of the weather.[4]

To see why deviation from sincere play arises, consider the recipient's problem at a belief just above the donor's funding threshold. Both types attack in the shine, and thus a defense reveals the rain whether or not the donor observes the flow payoff. Consider then a funded belief low enough that a revealed rain drops the donor's posterior below the next funding threshold. Defending there collects b and ends support. Attacking conceals the rain with probability $1 - \lambda$, and a recipient who keeps funding collects at least one more funded sincere period, which is worth more than b by Assumption 1. Additionally, the two types never separate at a funded belief, with one playing sincerely and the other attacks regardless of the weather. Because both types attack in the shine, separation would require one type to attack in the rain while the other defends, sacrificing the current payoff b in exchange for a more favorable reputation. However, Assumption 2 ensures that the resulting improvement in continuation value is never large enough to justify this sacrifice.

The result takes the donor's monitoring capacity to be sufficiently low. Attacking in the rain trades a certain payoff of b for a chance at continued funding, and that chance is the probability $1 - \lambda$ that the flow payoff goes unobserved. When the donor is unlikely to observe flow payoffs, concealment almost always works and attacking in the rain is worth it. If monitoring capacity were high, on the other hand, an attack in the rain is likely to be exposed, and the recipient then loses the defense payoff while the donor learns of the rain anyway. The bound λ^* guarantees that, at every belief where a defense would end funding, attacking today is likely enough to conceal the rainy weather that it is worth sacrificing the flow of b.

Figure 2 illustrates the equilibrium dynamics, outlining the trajectory of donor beliefs about a strong recipient over the course of war. The two panels differ in what the donor

[4] This intermediate region is nonempty in all periods except the terminal period T.

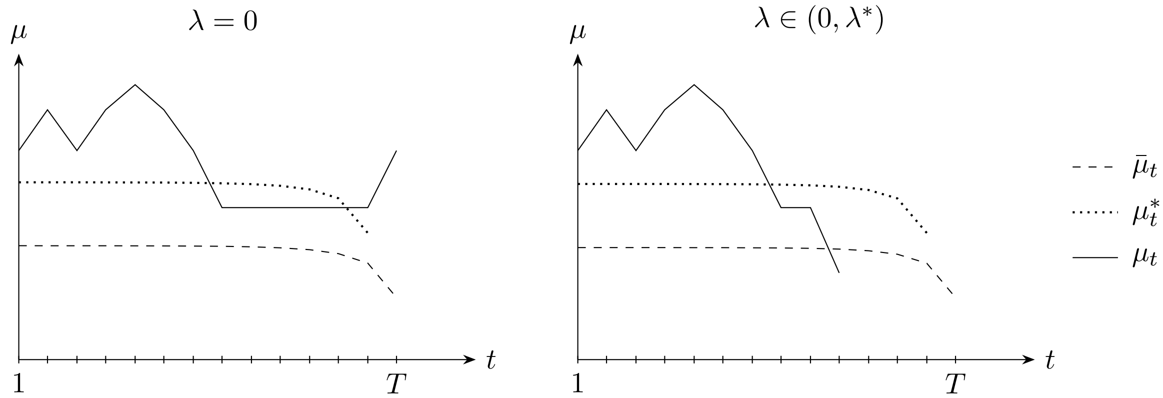

Figure 2: On-path belief dynamics μ_t for a strong recipient. *Note:* The solid line is a realized belief path. Both types play sincerely above the dotted threshold μ_t^* and always attack below it, and the donor funds above the dashed threshold $\bar{\mu}_t$. At $\lambda = 0$, funding in the first period guarantees funding through period T, whereas at $\lambda > 0$ even the strong type can lose funding. The parameter values used to generate the figure are $a = 1$, $b = 1/12$, $c = 1/2$, $\delta = 11/20$, $s_0 = 1/3$, $s_1 = 2/3$, $\mu_1 = 3/4$, $T = 14$, and for the right panel, $\lambda = 1/30$.

can learn. In the left panel, there is no monitoring $\lambda = 0$ and thus, because actions are the donor's only source of information, the belief does not move while the recipient pools on attack. Because the thresholds decline over time, a belief can escape the pooling region, in which case sincere play resumes. The donor learns nothing when funding is in doubt, but can learn again only when future funding is sufficiently secure. On the path of play, the belief never falls below the funding threshold, and funding that begins in the first period continues through the entire war.

On the other hand, the right panel of Figure 2 illustrates the trajectory of beliefs under low monitoring $\lambda \in (0, \lambda^*)$. The flow payoff can be observed, in which case a flow payoff of zero after an attack reveals rain, pushing the donor's belief about recipient quality below their funding threshold. As shown in the figure, the donor's support can end on the equilibrium path of play *even for the strong type* that the donor would prefer to fund perpetually. The weak type does run this risk more heavily, though, as it rains more frequently for them. In the limit as $T \to \infty$, strong types lose support with non-zero probability, while weak types lose support almost surely.

Theorem 1 pins down the shape of every monotone equilibrium but not its thresholds, and different monotone equilibria can pool on different sets of beliefs. Given we are interested in how the conditionality of support can create strategic distortions from sincere play, we are especially interested in equilibria at the boundary of the recipient's willingness to play sincerely. Thus, for the remainder of our analysis, we study what we call most-sincere equilibria.

3.1 Most-sincere equilibria

We now define the class of most-sincere equilibria. Our results concern distortion in battlefield tactics, so the natural equilibria to examine are those in which the recipient distorts by attacking in the rain the least. Then, the strategic distortions that survive reflect the environment, and are not artifacts of selecting a poorly behaved equilibrium.

Definition 2 (Most-sincere Equilibria)**.** *Define* $\underline{\mu}_t^* := \inf_{\sigma \in \Sigma} \mu_t^*(\sigma)$ *the lowest recipient's sincerity threshold, where* Σ *is the set of monotone equilibria. A monotone equilibrium is most-sincere if and only if* $\mu_t^* = \underline{\mu}_t^*$ *for every period* $t < T$*.*

Intuitively, a most-sincere equilibrium is one in which the recipient plays sincerely for the largest range of beliefs, among all monotone equilibria.

Lemma 1. *Suppose* $\lambda < \lambda^*$*. Then,*

(i) most-sincere equilibria exist;

(ii) every most-sincere equilibrium has the lowest donor's funding threshold for all $t \geq 2$ *among all monotone equilibria,* $\bar{\mu}_t = \underline{\bar{\mu}}_t := \inf_{\sigma \in \Sigma} \bar{\mu}_t(\sigma)$*; and*

(iii) every most-sincere equilibrium is Pareto optimal among all monotone equilibria.

The lemma shows that, by selecting on the most-sincere recipients, we also pin down the donor's behavior. Indeed, equilibria with the most recipient sincerity must also produce the most donor funding. Although most-sincere equilibria may differ off the path of play and in the first period's funding decision, they are otherwise observationally equivalent, generating the same behavior at every belief on the path and thus a unique distribution over outcomes. In this way, the thresholds $\underline{\bar{\mu}}_t$ and $\underline{\mu}_t^*$ fully capture behavior in every most-sincere equilibrium. Moreover, because they bound the corresponding thresholds of every monotone equilibrium from below, the distortions they exhibit understate those of any other monotone equilibrium.

The result also shows that most-sincere equilibria are Pareto optimal among monotone equilibria, i.e., there are no alternative monotone equilibria in which a player can be made better off without making the other worse off. This reinforces our sense in which most-sincere equilibria are the natural class to examine, as efficiency cannot be improved through equilibrium selection.

4 Baseline results

4.1 Comparative statics

We now examine how the equilibrium thresholds change with the model's parameters. The thresholds $\bar{\underline{\mu}}_t$ and $\underline{\mu}_t^*$ depend on time only through the number of remaining periods, and at $\lambda = 0$, as that number grows, the thresholds settle toward limits. At $\lambda = 0$, we denote by $\bar{\underline{\mu}}^{[\infty]} := \lim_{T\to\infty} \bar{\underline{\mu}}_1$ the limiting funding threshold and by $\underline{\mu}^{*[\infty]} := \lim_{T\to\infty} \underline{\mu}_1^*$ the limiting sincerity threshold. The thresholds approach their limits as more periods remain, and thus the limits approximate the thresholds in most periods of a long game.

Proposition 1 (Comparative Statics)**.** *Suppose $\lambda < \lambda^*$. Then,*

1. *in every period t, the thresholds $\bar{\underline{\mu}}_t$ and $\underline{\mu}_t^*$ are weakly increasing in c and weakly decreasing in a, b, and δ; and*

2. *at $\lambda = 0$, the threshold $\bar{\underline{\mu}}^{[\infty]}$ is decreasing in s_0 and s_1, and the threshold $\underline{\mu}^{*[\infty]}$ is decreasing in s_0, and increasing in s_1 if and only if $s_1 > 1/2 + c/(2a)$.*

The payoff parameters move the thresholds in every period, in directions that hold at any sufficiently low level of monitoring capability. When there is no monitoring, $\lambda = 0$, these limits additionally admit closed forms that do not depend on b or δ, which we can use to sign the effects of recipient strength.

Intuitively, increasing the outside payoff c raises what the donor forgoes in each funded period, so the donor is willing to fund only at higher beliefs, and both thresholds rise. In contrast, higher flow payoffs a and b raise the donor's expected flow from funding at every belief, and a higher δ raises the weight on the non-negative continuation surplus that funding preserves, so the thresholds decline in a, b, and δ. The sincerity threshold marks the beliefs from which a defense would drop the donor's posterior below the funding threshold, and thus the two move in tandem.

The effects of recipient strength are less straightforward. Increasing the strength of either type makes funding more attractive at every belief, so the limiting donor's funding threshold falls in both s_0 and s_1. On the other had, the limiting recipient's sincerity threshold responds to strength through two channels. First is the effect on the donor's funding threshold: both a higher s_0 and a higher s_1 lower the funding threshold, which in turn lowers the belief at which an additional defense ends donor support. Second is the effect of a defense on the donor's posterior belief. A higher s_0, holding s_1 fixed, makes a rain relatively less indicative of weakness, reducing the drop in the donor's posterior belief after a defense. Thus, both channels move in the same direction with regard to changes in s_0.

However, this second channel opposes the first with regard to s_1, as an increase to s_1 widens the strength gap between the strong and weak types, making a defense even more indicative of weakness. When s_1 is already large, this second channel dominates. This is because, as the strong type gets even stronger relative to the weak type, a single defense can drop the donor's posterior arbitrarily close to zero. A higher s_1 then raises the limiting sincerity threshold, and the recipient plays sincerely at fewer beliefs. Greater fighting ability can therefore make it *more* difficult for a strong recipient to fight sincerely.

4.2 Welfare under low monitoring

The previous section fixes monitoring capacity and varies the model parameters governing payoffs, and we now examine how changes in monitoring affect recipient welfare holding the payoff parameters fixed. One might expect strong recipients to welcome the donor's scrutiny, as monitoring gives a strong type the chance to prove their strength and separate from the weak. The next result shows it is not this simple.

Proposition 2. *For each $\lambda \in [0, \lambda^*)$, fix any most-sincere equilibrium with funding and sincerity thresholds $\bar{\mu}_t(\lambda)$ and $\mu_t^*(\lambda)$. Denote by $V_{\theta,t}(\mu;\lambda)$ the recipient continuation value* *given λ. Then, for every $t < T$, each θ, and all but finitely many $\mu \in (\bar{\mu}_t(0), \mu_t^*(0))$,*

$$\left.\frac{\partial V_{\theta,t}(\mu;\lambda)}{\partial \lambda}\right|_{\lambda=0} < 0.$$

At the beliefs close to the funding threshold, monitoring cannot deliver the separation a strong type might hope for. Both types attack in the rain at these beliefs, so monitoring reveals only the weather, and both types experience occasional rain. With no monitoring, a recipient who pools on attack is certain to keep funding, as nothing can reveal the rain and the belief stays put until either the end of the game or the sincerity threshold declines below it and sincere play resumes. On the other hand, a little monitoring exposes this recipient to two events. A flow payoff revealing the rain drops the donor's posterior below the funding threshold, and the recipient of either type loses their support. A flow payoff revealing the shine moves the donor's belief upward, which only brings forward the sincere play that was coming anyway. The first event costs the recipient more than the second gains them. Thus, at the margin, both types of recipient would choose no monitoring over a little monitoring, even the strong type.

The result speaks to what donors can and cannot infer from opposition to monitoring when continued support is in doubt. At these beliefs, both types of recipient prefer that the donor gain no additional capacity to observe the war, and would thus oppose the donor investing into even modest oversight, such as an occasional inspection or a noisy intelligence channel. Opposition of donor monitoring is then not informative about the recipient's quality.

Next, we investigate a benchmark case with full monitoring, $\lambda = 1$, for comparison.

4.3 Full monitoring

We close the baseline analysis with a discussion of the full-monitoring benchmark, in which $\lambda = 1$ and hence the donor observes every flow payoff. Because the donor can infer the weather regardless of what the recipient does, the recipient gains nothing by distorting their battlefield tactics. As a result, there exists a monotone equilibrium in which both types of recipient play sincerely at every funded belief. In fact, for any finite horizon, the recipient plays sincerely at every funded belief in essentially every monotone equilibrium (Proposition

7, Appendix D.1).[5]

In the sincere full-monitoring equilibrium, the funding threshold also falls dramatically. When the weather is revealed every period, the donor who funds today learns the recipient's type over time, giving a strong recipient time to prove their strength and stopping only after enough bad news. The chance to learn makes funding worthwhile even at beliefs where its expected flow payoff falls short of c, whereas the donor would never fund such beliefs at $\lambda < \lambda^*$.

Together, the results suggest that the distortion in battlefield tactics is in part due to limited donor monitoring capability, and not only from the conditionality of donor support. Recipients distort at every monitoring capacity below λ^*, while the sincere full-monitoring equilibrium has no distortion at $\lambda = 1$. Monitoring also determines whether funding outcomes differ by type. Under no monitoring, both types are funded for the game's duration, whereas each type loses funding with non-zero probability under full monitoring.

Figure 3 compares the strong type's value function across the cases of no, low, and full monitoring in the infinite-horizon game. With no monitoring, the value is flat across the pooling region, as the strong type always attacks and collects $s_1 a$ every period until the game ends. With low monitoring, the value in this region falls below the solid line, as funding now ends at the first revealed rain, and the same stream of attack payoffs is cut short with positive probability. This reflects the reduction in welfare shown by Proposition 2, which can be seen in the vertical drop from the solid to the dashed line for beliefs just above the no-monitoring funding threshold $\underline{\bar{\mu}}(0)$.

On the other hand, with full monitoring, the donor is willing to fund at lower beliefs and thus the recipient's value function is strictly positive at beliefs where it is zero for the other cases. At beliefs where the recipient would otherwise always attack, sincere play collects the defense payoff that attacking forfeits, while revealed rains put funding at risk. The full-monitoring value thus dips below the no-monitoring value only near the donor's no-monitoring funding threshold, where exit risk is highest. All three cases converge at $\mu = 1$, where sincere play yields $(s_1 a + (1 - s_1)b)/(1 - \delta)$.

[5]Sincere play obtains for all but finitely many values of b, under the convention that beliefs after zero-probability action-payoff observations condition only on the inferred weather.

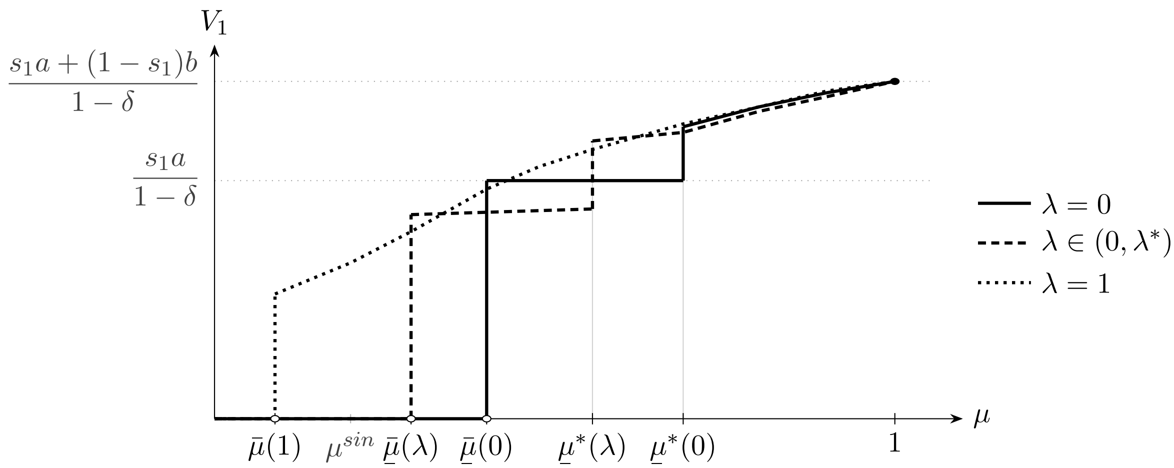

Figure 3: Strong recipient's value function V_1 in the donor's belief over an infinite horizon. *Note:* The solid line reflects no monitoring $\lambda = 0$, the dashed line reflects low-level monitoring $\lambda \in (0, \lambda^*)$, and the dotted line reflects full monitoring with $\lambda = 1$ as constructed in part (i) of Proposition 7. For legibility, levels and spacing are not drawn to scale.

5 Commitment

The baseline results locate the strategic distortion in battlefield tactics in the conditionality of funding and low monitoring capability, as recipients attack in the rain to maintain support they could otherwise lose. A natural potential remedy is for the donor to commit to future funding ahead of time. We now show that limited commitments, even if credible, do not change the recipient's tactics at all.

We adapt the baseline model as follows. In the first period, the donor commits to funding the first k periods regardless of what they subsequently observe, and we refer to these periods as the commitment block. From period $k + 1$ on, the game proceeds as before without commitment for the remaining $T - k$ periods, with play following a most-sincere equilibrium of that continuation game that attains the lowest funding threshold in its first period, which exists by Lemma 1.

Specifically, the timing of the commitment game is as follows, with the new commitment component italicized. Steps 1–4 are identical to those in the baseline model.

1. Nature draws the recipient's strength, $\theta \in \{0, 1\}$, observed by the recipient and for which the donor has a well-specified prior belief, μ_1.

2. The donor stays out or funds, $d_t \in \{0, 1\}$. If they stay out, the game ends.

3. Nature rains or shines $\omega_t \in \{0, 1\}$ with probs. $1 - s_\theta$ and s_θ, respectively, where $1 > s_1 > s_0 > 0$.

4. The recipient observes the weather, then plays defends or attacks, $r_t \in \{0, 1\}$.

5. The donor updates beliefs by Bayes' rule given the recipient's action and flow payoffs if they observe them (which occurs with prob. λ).

   If $t < k$, period $t+1$ proceeds from step 3. If $k \le t < T$, period $t+1$ proceeds from step 2. If $t = T$, total payoffs are realized.

We evaluate the efficacy of commitment in most-sincere equilibria, as a commitment that fails to restore sincerity in the equilibria with the most sincere play fails in every other monotone equilibrium. Moreover, we focus on the case of no monitoring, $\lambda = 0$. This is because when flow payoffs are never observed, commitment is tested against the incentive to distort alone. To see this, note that the donor's beliefs can decline either because of what the recipient does (the recipient takes actions to defend) or despite what it does (the recipient acts to maintain donor beliefs, but the donor observes payoffs through its monitoring capability and infers the weather). The goal of committed support is to overcome the recipient's strategic incentive to distort their tactics, but a commitment unavoidably also insures the recipient against unfavorable monitoring realizations by the donor. We thus isolate the effect of commitment on the recipient's incentive to deviate from sincere play by setting $\lambda = 0$.

Proposition 3. *Let $T < \infty$, $\lambda = 0$, and $\mu_1 \in [\bar{\underline{\mu}}_1(0), \underline{\mu}_1^*(0))$. Then, there exists a $\bar{k} \ge 2$ such that, for every $k < \bar{k}$, the commitment game has a monotone equilibrium and, in every monotone equilibrium, both types of recipient attack at every belief reached in the block.*

We find that giving the donor the ability to commit to funding does not recover sincere recipient fighting. During the block, the recipient cannot lose funding, yet their tactics still move the donor's belief, and the belief at the block's end decides whether funding continues. In every monotone equilibrium of a game with sufficiently limited commitments, both types attack through the entire block.

The incentive to distort battlefield tactics unravels backward through the commitment block. Consider a block of k periods, and suppose the recipient reaches period k. Sincere play would have a defense reveal rain, and the donor enters period $k+1$ believing the recipient too weak to fund, so the recipient attacks in the rain just as they would without a guarantee. Now consider period $k-1$. A defense there reveals the rain all the same, the donor carries that knowledge to the block's end, and funding is nonetheless revoked when discretion returns. The recipient thus attacks in period $k-1$, as well, and the argument continues backward until the game's start. The bound $\bar{k}$ marks the block lengths at which funding after the block is worth more to the recipient than what they expect from sincere play inside it.

Therefore, limited commitments do not help. In fact, we find that as the game lengthens and countries care enough about the future, the unraveling argument covers commitments of *any* fixed duration.

Proposition 4. *Let $\lambda = 0$. Then, for every $K \in \mathbb{N}$, there is a $\bar{\bar{\delta}} < 1$ and $\bar{T} : (0,1) \to \mathbb{N}$ such that, for every $\delta > \bar{\bar{\delta}}$, $T > \bar{T}(\delta)$, $\mu_1 \in [\underline{\bar{\mu}}_1(0), \underline{\bar{\mu}}^{[\infty]}(0)]$, $k \leq K$, and all (a, b, c, s_0, s_1) satisfying Assumptions 1–2, the commitment game has a monotone equilibrium and, in every monotone equilibrium, both types of recipient attack at every belief reached in the commitment block.*

For any commitment block the donor might consider, a long enough war between patient players produces the same distorted battlefield tactics as no commitment at all.[6] Any commitment that does not last to the end of the game leaves a discretionary period for tactics to influence, and if players are sufficiently patient, the incentives unravel backward from the end of the commitment period.

Total commitment and full monitoring accomplish the same thing by different means. The recipient is unable to conceal the weather when the donor observes flow payoffs, and they have no reason to distort their battlefield tactics when the donor commits to fund them for the duration of the war. As a result, sincere fighting is attained in each case. However, as our results show, neither commitment nor monitoring accomplishes this in its partial form: recipients are willing to deviate from sincere play at every sufficiently low monitoring capacity

[6]Note that, because Assumption 2 requires $\delta < 1/(2 - s_0)$, the discount factors above $\bar{\bar{\delta}}$ are admissible only alongside sufficiently high values of s_0, and thus long wars with patient players must come with weak types receiving the shine often enough for the strategic tension to hold.

and sincere play cannot be recovered by a commitment of any fixed length when the war is long enough relative to it and the countries are patient enough.

It is also worth noting that our second commitment result requires the initial belief to be at or below the limiting funding threshold rather than up to the sincerity threshold. This is because, for any of these beliefs, a defense reveals the rain and moves the donor's posterior below the funding threshold of every period at every time horizon, and the unraveling therefore runs through every possible commitment block. Thus, these beliefs are those at which the need for commitment is greatest, as continued support is guaranteed to remain in doubt.

6 Empirical implications

Our results characterize when belligerents intentionally and strategically perform suboptimally on the battlefield, why belligerents are worse off when the donor has the capacity to engage in low-level monitoring, and why funding commitments typically fail to remedy the problem. Each of these claims speaks to an active empirical literature on security assistance.

In this section, we discuss three main empirical implications. First, we illustrate the mechanism of Theorem 1 in the US experience in Afghanistan as well as its implications for what can be learned from the conditionality of external war support. Second, we show how Proposition 2 provides important caveats to the inferences scholars can make from a recipient's opposition to donor monitoring, with Ukraine's opposition to US oversight as the running example. Third, we argue that the incentive structure that sustains Propositions 3 and 4 was present in the case of US support to the Soviet Union during World War II, and offer implications for what this may imply about Soviet conduct at the time.

6.1 Battlefield tactics in Afghanistan (2015-2021)

Theorem 1 identifies two conditions under which the belligerent distorts their battlefield tactics: (i) the donor's monitoring capacity is low and (ii) the donor's belief is high enough to merit support but low enough to leave the belligerent insecure in future funding. At these beliefs, the recipient does not respond to battlefield conditions, sacrificing the occasional payoff of a correct defense to keep the donor from observing negative conditions. Importantly, both the weak and the strong belligerent engage in this strategic distortion, as continued support is in doubt for each at these intermediate beliefs.

Our analysis reveals a new downside of conditionality in external support that has been missed by the empirical literature. Work on security assistance usually focuses on settings where the donor and recipient have divergent interests. This is a more standard principal-agent problem where the donor must extract effort from the recipient, who has an incentive not to comply. In that setting, it is regularly argued that conditionality is the *remedy* to unfavorable distortions by the recipient (Biddle, Macdonald, and Baker, 2018; Berman and Lake, 2019). Instead, we explore a less conventional principal-agent problem in which incentives are perfectly aligned conditional on delegation, i.e., the donor and the recipient want exactly the same war outcome once support has been dealt. In our setting, we find conditionality becomes the *source* of the recipient's distortions.

The observable antecedents of the theorem hold in the case of Afghanistan. The US funded the Afghan National Defense and Security Forces through annual appropriations, and the continuation of that support was never assured, with the option of withdrawal openly contested within the US. After the combat mission ended in 2015, fewer advisers were embedded in the field and those that remained had limited interaction with lower-level units. As a result, much of the new information about battlefield conditions came through reports that the US could not independently verify. For example, the US had limited ability to independently verify what units below the corps level were reporting, and it was found that the tools used for assessment failed to capture intangibles on the ground.[7]

Our theory offers a compelling explanation for why Afghan commanders exhausted their best troops in repeated efforts to reinforce threatened positions and recapture lost ground. These seemingly counterproductive tactics may have been catered towards securing funding despite being suboptimal in pursuing their objectives in war. As US support grew increasingly uncertain in 2021, Afghan commanders committed their most capable force, the Commandos, to recover lost territory. Despite the fact that the Commandos had systems in place for rest and recovery, these systems were regularly ignored.[8] Consistent with our analysis, there was

[7]Refer to SIGAR (2017) and SIGAR (2023).

[8]Refer to United States Department of Defense (2020).

a long-term incentive to engage in this distortionary behavior, as holding the Commandos back would have communicated that Afghan leaders viewed the mission as unattainable.

Additionally, our theory provides an explanation for why the US could not see the collapse of the Afghan forces coming. The conventional wisdom is that the US had a poor assessment of Afghan capability; however, our results suggest that improving the assessment framework alone would not have solved this problem so long as the US maintained low monitoring capability on the ground. Because both the strong and weak recipient distort their observable tactics in the face of insecure funding, the operations the US could have observed would have carried little information about the long-term prospects of the Afghan forces. The US, having no news on which to update, therefore kept its beliefs largely fixed until forces outside the recipient's control revealed the bad news.

6.2 Opposition to monitoring in Ukraine (2022-2024)

Proposition 2 presents the counterintuitive finding that the strong recipient, in addition to the weak one, prefers no monitoring to low levels when their continued support is sufficiently vulnerable. One would expect the strong type to welcome monitoring, as it provides them an opportunity to reveal their strength to the donor. However, when beliefs are intermediate and a negative shock to the donor's belief can cause them to pull their support, even the strong types worry about being unlucky and having the occasional "rainy weather" revealed to the donor. Any marginal gain at low levels of monitoring is guaranteed to be offset by the expected losses incurred by this risk.

This logic has parallels in the case of US support for Ukraine. Although the US supplied crucial intelligence to Ukraine, Ukrainian officials chose to keep most of their operational plans private from the start of the war (Barnes, Schmitt, and Cooper, 2022). US officials in the Biden administration joked that they knew more from spying on the Russians than from their Ukrainian partners (Entous, 2025). The common explanation is that Ukraine did not want to be vetoed by the US, but our theory suggests it may have also been in part due to concerns that their plans would highlight weaknesses and discourage continued support.

In particular, our explanation is consistent with the fact that Ukraine was regularly insecure in their external aid at various times throughout the war. Ukrainian officials were frequently seeking assurances from the US by expressing concerns about their support after the US elections and raising the case of Afghanistan with them directly, fearing they would be abandoned (Entous, 2025). Snyder (1984) proposes this fear of abandonment in the alliance security dilemma, identifying heightened loyalty to the ally and firmness toward the adversary as common reactions. Our theory identifies another reaction in that the recipient strategically keeps information about the war from the donor.

Moreover, Ukraine began sharing their battlefield plans and seeking US and British advice later in the summer of 2022, once it was viewed as being likely that the joint offensive would succeed (Barnes, Schmitt, and Cooper, 2022). Yet, after the failed counteroffensive in 2023 and the stalled continuation of support by the US Congress, Ukraine began keeping operations private again. For example, a covert ground operation into Russia was planned in March 2024 without US knowledge, leading to a confrontation between the CIA station chief in Kyiv and head of Ukrainian military intelligence. Likewise, foreknowledge of the Kursk incursion was kept from the US (Entous, 2025). This fluctuation in Ukrainian transparency with the US is consistent with our model, as recipients prefer less monitoring at intermediate beliefs in which support is precarious. Surprisingly, our model also suggests that these changes in transparency do *not* necessarily reflect changes in the underlying recipient prospects.

6.3 Limited commitments in the Soviet Union (1941-1943)

Next, we provide evidence for our results on commitment. Specifically, Propositions 3 and 4 suggest that limited commitments are unlikely to mitigate the recipient's distortionary behavior. Even while funding commitments are in place, the recipient has the incentive to manage the donor's beliefs by overattacking. When a commitment period does not persist through to the end of the war, there is always a period at which the donor can decide to pull funding, and they may choose to do so if their beliefs have not been managed in the interim. This runs in contrast to typical expectations, where committed support is proposed to free the recipient of their funding concerns and allow them to focus on winning the war.

Our theoretical findings on commitment are illustrated in the case of aid to the Soviet Union during World War II. This included a series of agreements in which US and British support was guaranteed for a specified period. Throughout this arrangement, the US stayed committed to honoring their existing agreements but remained open to whether future support should be continued or not. This approach was a policy of "unconditional aid" (Herring, 1973), but only within the agreement periods. Beyond that, support remained conditioned on evolving US beliefs about Soviet prospects. For example, in a 1941 memorandum to the Secretary of War, President Roosevelt insisted on the importance of providing the USSR with commitments to help with munitions through June 1942; however, considerations on aid beyond that point were left "to be decided at a later date, in the light of the then existing circumstances" (Roosevelt, 1941). Most importantly, Roosevelt felt that aid should be tied to whether the USSR "continues to fight the Axis powers effectively," closely matching the interests and incentives of the donor in our model.

In our theory, when beliefs are intermediate, the donor learns nothing about the recipient's strength during a limited commitment block, as the recipient attacks regardless of the weather to conceal unfavorable conditions. Similarly, during this time, the Soviets were unwilling to share information with their allies, and so US beliefs moved primarily with what could be assessed from observables. For example, roughly halfway through the first agreement, the US embassy in Moscow assessed USSR morale and evaluated their losses and capacity to withstand the coming German offensive (Thurston, 1942). A subsequent agreement was proposed to the USSR prior to the conclusion of the first agreement, and hence its terms reflected beliefs formed before the conclusion of this first commitment block (Roosevelt, 1942). Of course, we cannot be certain that Soviet tactics responded to this incentive, but the details of the case reveal that the underlying incentive structure was in place.

7 Conclusion

This paper studies external support in war when a donor's willingness to continue providing support depends on their beliefs about the recipient's prospects. We develop a dynamic model in which the donor and recipient agree on the optimal battlefield tactics, but the recipient privately observes conditions on the ground and the donor learns about recipient strength from battlefield behavior. We show that, when monitoring is limited, recipients may distort their tactics in order to preserve future support (Theorem 1). When the donor is confident in the recipient, the recipient fights sincerely. At intermediate levels of confidence where funding remains secure today but is vulnerable in the future, both types instead attack regardless of the battlefield conditions in order to regulate the donor's beliefs.

We find that greater fighting ability can make it more difficult for a strong recipient to fight sincerely because a setback becomes more informative of weakness (Proposition 1). We also show that limited monitoring can make both weak and strong recipients worse off (Proposition 2). This is because limited monitoring exposes unfavorable battlefield conditions without mitigating the recipient's incentive to distort their tactics. In contrast, full monitoring reveals battlefield conditions and recovers sincere play (Proposition 7).

Finally, we consider whether commitments to future support can eliminate distortions caused by the conditionality of support. We show that limited commitments do not change recipient behavior because battlefield actions during the commitment period continue to affect the donor's beliefs and whether support continues after the commitment ends (Proposition 3). For any finite commitment length, a sufficiently long war between patient players induces the same behavior as a game with no commitment (Proposition 4). The results suggest that neither low levels of monitoring nor limited commitments resolve the incentives created by the conditionality of external support.

Our theoretical results have empirical implications for how we understand conditional external support in war. Uncertainty over continued US assistance to Afghanistan created incentives for commanders to demonstrate strength through costly offensive operations even when those tactics were counterproductive to their broader war aims. In Ukraine, the results caution against interpreting Ukrainian opposition to US monitoring as evidence of weakness, as even a strong recipient may prefer less monitoring when continued support is insecure. Lastly, the commitment results help explain why US assessments of Soviet prospects during World War II continued to matter for the USSR despite aid agreements already being in place. Together, these cases illustrate how the conditionality of external support can create perverse incentives for belligerents in war.

References

Barnes, Julian, Eric Schmitt, and Helene Cooper. 2022. "The critical moment behind Ukraine's rapid advance." *New York Times* 13.

Beardsley, Kyle. 2012. "UN intervention and the duration of international crises." *Journal of Peace Research* 49(2):335–349.

Benson, Brett V, Adam Meirowitz, and Kristopher W Ramsay. 2014. "Inducing Deterrence Through Moral Hazard in Alliance Contracts." *Journal of Conflict Resolution* 58(2):307–335.

Berman, Eli, and David A. Lake, eds. 2019. *Proxy Wars: Suppressing Violence Through Local Agents*. Ithaca, NY: Cornell University Press.

Besley, Timothy, and Torsten Persson. 2010. "State Capacity, Conflict, and Development." *Econometrica* 78(1):1–34.

Biddle, Stephen, Julia Macdonald, and Ryan Baker. 2018. "Small Footprint, Small Payoff: The Military Effectiveness of Security Force Assistance." *Journal of Strategic Studies* 41(1–2):89–142.

Cunningham, David E. 2010. "Blocking resolution: How external states can prolong civil wars." *Journal of Peace Research* 47(2):115–127.

Dube, Oeindrila, and Suresh Naidu. 2015. "Bases, Bullets, and Ballots: The Effect of US Military Aid on Political Conflict in Colombia." *The Journal of Politics* 77(1):249–267.

Entous, Adam. 2025. "The partnership: The secret history of the war in Ukraine." *New York Times* 29.

Fearon, James D, and David D Laitin. 2003. "Ethnicity, Insurgency, and Civil War." *American Political Science Review* 97(1):75–90.

Gailmard, Sean. 2021. "Imperial Governance and the Growth of Legislative Power in America." *American Journal of Political Science* 65(4):912–925.

Gailmard, Sean, and John W. Patty. 2007. "Slackers and Zealots: Civil Service, Policy Discretion, and Bureaucratic Expertise." *American Journal of Political Science* 51(4):873–889.

Grillo, Edoardo, and Antonio Nicolò. 2025. "Learning the hard way: Conflicts, sanctions and military aid." *Journal of Public Economics* 242:105310.

Herring, George C. 1973. *Aid to Russia, 1941–1946: Strategy, Diplomacy, the Origins of the Cold War*. New York: Columbia University Press.

Kofman, Michael, and Rob Lee. 2023. "Perseverance and Adaptation: Ukraine's Counteroffensive at Three Months." War on the Rocks, September 4.

Krainin, Colin, Caroline Thomas, Thomas Wiseman et al. 2020. "Rational Quagmires: Attrition, Learning, and War." *Quarterly Journal of Political Science* 15(3):369–400.

Krainin, Colin, Kristopher W Ramsay, Bella Wang, and Joseph J Ruggiero. 2022. "Preventive war and sovereign debt." *Conflict Management and Peace Science* 39(5):487–519.

Lipnowski, Elliot, and Joao Ramos. 2020. "Repeated delegation." *Journal of Economic Theory* 188:1–41.

Padró i Miquel, Gerard, and Pierre Yared. 2012. "The Political Economy of Indirect Control." *The Quarterly Journal of Economics* 127(2):947–1015.

Powell, Robert. 2004. "Bargaining and Learning While Fighting." *American Journal of Political Science* 48(2):344–361.

Regan, Patrick M. 2002. "Third-party Interventions and the Duration of Intrastate Conflicts." *Journal of Conflict Resolution* 46(1):55–73.

Roosevelt, Franklin D. 1941. "Memorandum by President Roosevelt to the Secretary of War (Stimson).". August 30, 1941. Foreign Relations of the United States, Diplomatic Papers, 1941, General, The Soviet Union, Volume I, Document 770.

Roosevelt, Franklin D. 1942. "President Roosevelt to the People's Commissar for Foreign Affairs of the Soviet Union (Molotov).". May 29, 1942. Foreign Relations of the United States, Diplomatic Papers, 1942, Europe, Volume III, Document 588.

Sawyer, Katherine, Kathleen Gallagher Cunningham, and William Reed. 2017. "The Role of External Support in Civil War Termination." *Journal of Conflict Resolution* 61(6):1174–1202.

SIGAR. 2017. Reconstructing the Afghan National Defense and Security Forces: Lessons from the U.S. Experience in Afghanistan. Technical Report SIGAR-17-62-LL Special Inspector General for Afghanistan Reconstruction Arlington, VA: .

SIGAR. 2023. Why the Afghan Security Forces Collapsed. Technical Report SIGAR 23-16-IP Special Inspector General for Afghanistan Reconstruction Arlington, VA: .

Snyder, Glenn H. 1984. "The security dilemma in alliance politics." *World Politics* 36(4):461–495.

Spear, Stephen E, and Sanjay Srivastava. 1987. "On Repeated Moral Hazard with Discounting." *The Review of Economic Studies* 54(4):599–617.

Thurston, Walter. 1942. "The Chargé in the Soviet Union (Thurston) to the Secretary of State.". March 20, 1942. Foreign Relations of the United States, Diplomatic Papers, 1942, Europe, Volume III, Document 366.

United States Department of Defense. 2020. Enhancing Security and Stability in Afghanistan. Report to congress United States Department of Defense Washington, DC: .

Online Appendix for "Propping Up Allies With Uncertain Prospects"

Contents

A Preliminaries

Throughout, an equilibrium specifies the donor's belief at every history, with Bayes' rule applied at every belief the donor funds, beliefs concentrated on the support of the current belief, and beliefs otherwise unrestricted at off-path histories the donor reaches only by deviating to fund.

We begin by defining a few key objects pertaining to the donor's beliefs. First, for simplicity, define $s(\mu) := \mu s_1 + (1-\mu)s_0$ the expected probability of shine given donor beliefs μ, noting that it is affine and strictly increasing. This allows us to write the expected flow payoff at belief μ under sincere play by $u^{sin}(\mu) := s(\mu)a + (1-s(\mu))b$. When both types of the recipient play sincerely and the flow payoff is observed, we can also write the posterior beliefs as

$$\mu^+(\mu) = \frac{\mu s_1}{s(\mu)}; \qquad \mu^-(\mu) = \frac{\mu(1-s_1)}{1-s(\mu)}. \tag{A.1}$$

after inferring a shine and rain occurred, respectively.

Let us note two properties of these updates. First, each is invertible, as both are continuous and strictly increasing bijections with strictly positive derivatives. Second, because $\mu^+(\mu) > \mu > \mu^-(\mu)$ for all μ, we have

$$(\mu^-)^{-1}(\mu') > \mu' > (\mu^+)^{-1}(\mu')$$

for every $\mu' \in (0,1)$.

Next, let us define two thresholds in donor beliefs. First, the belief that makes the donor indifferent between the flow payoffs from funding and those from not funding when both types of recipient play sincerely is given by μ^{sin} that solves $u^{sin}(\mu^{sin}) = c$. Second, the belief that makes the donor indifferent about flow payoffs when both types pool on attack regardless of the weather is likewise given by μ^{att} that solves $s(\mu^{att})a = c$. Solving for these indifference beliefs results in

$$\mu^{sin} = \frac{c - s_0 a - (1-s_0)b}{(s_1-s_0)(a-b)} \in (0,1); \qquad \mu^{att} = \frac{c - s_0 a}{(s_1-s_0)a} \in (0,1). \tag{A.2}$$

To see that both beliefs are interior, note first that $c > s_0 a$. Otherwise the denominator of inequality (2), which expands to $(1-s_0)(a-c) - (1-s_1)(c-s_0 a)$, would be positive while its numerator would be at most zero, contradicting $b > 0$. Because $c < s_1 a$, the denominator is positive, and inequality (2) then gives $c - s_0 a > (1-s_0)b$, so both numerators above are positive, and both beliefs are below one because $c < s_1 a$. Note also that $b < (c-s_0 a)/(1-s_0) < a$, because $c < a$, so the defense payoff lies strictly below the attack payoff. Because $c/a > (c-b)/(a-b)$, we can also conclude that $\mu^{att} > \mu^{sin}$. The belief at which the donor is indifferent among flow payoffs when the recipient always pools on attack is strictly larger than that when the recipient plays sincerely.

More generally, with current beliefs μ, we can denote the donor's posterior belief after only action r is observed by $\mu^{r,\emptyset}(\mu)$ and the posterior after both action r and flow payoffs (and hence weather ω can be perfectly inferred) by $\mu^{r,\omega}(\mu)$. By part (i) of Definition 1, the posterior after an attack is weakly higher than that after a defense in any monotone equilibrium, hence for any μ, $\mu^{1,\emptyset}(\mu) \geq \mu^{0,\emptyset}(\mu)$ and $\mu^{1,\omega}(\mu) \geq \mu^{0,\omega}(\mu)$ for every ω. Noting that, at any belief the donor funds, the prior μ is a weighted average of the posteriors after the actions taken with positive probability, part (i) of Definition 1 gives $\mu \in [\mu^{0,\emptyset}(\mu), \mu^{1,\emptyset}(\mu)]$ at every such belief. Likewise, conditional on the weather ω, the updated prior is $\mu^{+}(\mu)$ in the shine and $\mu^{-}(\mu)$ in the rain, and part (i) of Definition 1 likewise gives $\mu^{+}(\mu) \in [\mu^{0,1}(\mu), \mu^{1,1}(\mu)]$ and $\mu^{-}(\mu) \in [\mu^{0,0}(\mu), \mu^{1,0}(\mu)]$ at every such belief.

For a fixed equilibrium, we let $\sigma_{\theta\omega}(\mu) \in [0,1]$ denote the probability that type θ plays the sincere action $r = \omega$ at weather ω when the donor's belief is μ. In this notation, $\sigma_{\theta 1}(\mu)$ is the probability that type θ attacks in the shine, and $\sigma_{\theta 0}(\mu)$ is the probability that type θ defends in the rain.

A.1 Lemma 2

Lemma 2. $\mu^{-}(\mu^{att}) < \mu^{sin}$.

Proof. First, note that

$$\mu^{-}(\mu^{att}) = \frac{\mu^{att}(1-s_1)}{1-s(\mu^{att})} = \frac{(1-s_1)(c-s_0 a)}{(s_1-s_0)(a-c)}$$

where the first equality follows by equation (A.1) and the second equality by (A.2). Further, recall the definition of μ^{sin} in terms of model primitives, given by equation (A.2).

Then, we can see that $\mu^{-}(\mu^{att}) < \mu^{sin}$ if and only if

$$\frac{(1-s_1)(c-s_0 a)}{(s_1-s_0)(a-c)} < \frac{c-s_0 a-(1-s_0)b}{(s_1-s_0)(a-b)}.$$

Solving for b, we recover

$$b < \frac{(s_1 a - c)(c - s_0 a)}{a(1-s_0 s_1) - c(2-s_0-s_1)}$$

which holds by Assumption 1. □

B The donor's problem

Let $\Phi^{[j]}(\mu)$ be the donor's funding surplus at belief μ with j periods remaining in the game when the recipient pools on attack at beliefs below $\mu^{*[j]}$ and plays sincerely at beliefs of at least $\mu^{*[j]}$, with $\Phi^{[0]}(\mu) = 0$, $\Phi^{[j]}(\mu) = \max\{0, \Psi^{[j]}(\mu)\}$, and $\bar{\mu}^{[j]} := \inf\{\mu \in [0,1] : \Psi^{[j]}(\mu) \geq 0\}$ for all $j \geq 1$. The surplus from funding the current period $T-j+1$ and continuing the game optimally is

$$\Psi^{[j]}_{sin}(\mu) = u^{sin}(\mu) - c + \delta\left[s(\mu)\Phi^{[j-1]}(\mu^{+}(\mu)) + (1 - s(\mu))\Phi^{[j-1]}(\mu^{-}(\mu))\right] \tag{A.3}$$

when both types of recipient play sincerely, where action reveals the weather and so beliefs move up and down with probabilities determined by $s(\mu)$, and

$$\Psi^{[j]}_{att}(\mu) = s(\mu)a - c + \delta\left[(1 - \lambda)\Phi^{[j-1]}(\mu) + \lambda s(\mu)\Phi^{[j-1]}(\mu^{+}(\mu))\right] \tag{A.4}$$

when both types pool on attack, where the belief stays where it is with probability $1 - \lambda$ or moves upward with probability $\lambda s(\mu)$. Then, $\Psi^{[j]}(\mu) = \Psi^{[j]}_{sin}(\mu)$ at every $\mu \geq \mu^{*[j]}$ and $\Psi^{[j]}(\mu) = \Psi^{[j]}_{att}(\mu)$ at every $\mu < \mu^{*[j]}$, where $\mu^{*[1]} := 0$ and $\mu^{*[j]} := (\mu^{-})^{-1}(\bar{\mu}^{[j-1]})$ for every $j \geq 2$, as usual applied to the current period. At $j = 1$, $\mu^{*[1]} = 0$ and $\Phi^{[0]} = 0$ give $\Psi^{[1]}(\mu) = u^{sin}(\mu) - c$, so play in the final period is sincere, $\bar{\mu}^{[1]} = \mu^{sin}$. Also note that the surplus is bounded, as $\Phi^{[j]}(\mu) \in [0, (a - c)/(1 - \delta)]$ for all μ.

Because the game ends once the donor pulls funding, $\Phi^{[j]} = \max\{0, \Psi^{[j]}\}$ is the value of the donor's optimal stopping problem against the posited recipient behavior, and the results of this section B are statements about the objects $\Phi^{[j]}$, $\bar{\mu}^{[j]}$, and $\mu^{*[j]}$. For the purposes of this section, recipient play here is an input rather than a best response, and we specify the threshold by $\mu^{*[j]} = (\mu^{-})^{-1}(\bar{\mu}^{[j-1]})$ at each $j \geq 2$. These thresholds are defined by the previous stage's funding threshold, so the recursion is well defined by backward induction and does not yet entail solving for the recipient's best response. We contend with equilibrium in subsequent sections, which cover arbitrary monotone equilibria such that the recipient's threshold need only satisfy $\mu^{*[j]} \geq (\mu^{-})^{-1}(\bar{\mu}^{[j-1]})$ rather than the equality above.

B.1 Proposition 5

With this, we can state the following result.

Proposition 5. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T - j + 1$. There exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, the donor's funding threshold satisfies $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$ for every $j \geq 2$.*

Proof. By definition of $\bar{\mu}^{[j]}$, we have $\Psi^{[j]}(\mu) < 0$ and hence $\Phi^{[j]}(\mu) = 0$ at every $\mu < \bar{\mu}^{[j]}$. Because s is affine with $s(\mu^{att})a = c$, we have $s(\mu)a - c = (s_1 - s_0)a(\mu - \mu^{att})$ at every belief.

First, we show by induction on j that

$$\Psi^{[j]}(\mu) \geq (\delta(1 - \lambda))^{j-1}\left(1 - \frac{c}{a}\right)b - \frac{(s_1 - s_0)a(\mu^{att} - \mu)}{1 - \delta} \tag{A.5}$$

for every λ, $j \geq 1$, and $\mu \leq \mu^{att}$. At $j = 1$, we have

$$\Psi^{[1]}(\mu) = u^{sin}(\mu) - c = s(\mu)(a-b) + b - c$$
$$= \left(1 - \frac{c}{a}\right) b - (s_1 - s_0)(a-b)(\mu^{att} - \mu),$$

where we use the fact that $s(\mu) = c/a - (s_1 - s_0)(\mu^{att} - \mu)$. Then, $\Psi^{[1]}(\mu)$ is at least the right-hand side of inequality (A.5), as $(s_1 - s_0)(a-b) \leq (s_1 - s_0)a \leq (s_1 - s_0)a/(1-\delta)$.

Now fix $j \geq 2$ and suppose inequality (A.5) holds with $j-1$ periods remaining at every $\mu \leq \mu^{att}$. Because $\Phi^{[j-1]} = \max\{0, \Psi^{[j-1]}\}$, the value $\Phi^{[j-1]}(\mu)$ is also at least the right-hand side of (A.5) at $j-1$, whether that quantity is negative or not. If the recipient pools on attack at μ then, because $\Phi^{[j-1]}(\mu^+(\mu)) \geq 0$,

$$\Psi^{[j]}(\mu) \geq (s_1 - s_0)a(\mu - \mu^{att}) + \delta(1-\lambda)\Phi^{[j-1]}(\mu)$$
$$\geq (\delta(1-\lambda))^{j-1}\left(1 - \frac{c}{a}\right) b - (s_1 - s_0)a(\mu^{att} - \mu)\left[1 + \frac{\delta(1-\lambda)}{1-\delta}\right],$$

and the bracketed term equals $(1-\delta\lambda)/(1-\delta) \leq 1/(1-\delta)$, which yields (A.5). If, on the other hand, the recipient plays sincerely at μ, the continuation terms are non-negative and $\Psi^{[j]}(\mu) \geq u^{sin}(\mu) - c \geq (1-c/a)b - (s_1 - s_0)a(\mu^{att} - \mu)$, which is at least the right-hand side of (A.5).

Now, inequality (A.5) is established, and thus $\Psi^{[j]}(\mu) > 0$ at every belief with

$$\mu^{att} - \mu < \frac{(1-\delta)(\delta(1-\lambda))^{j-1}(1-c/a)b}{(s_1 - s_0)a},$$

and every such belief is funded and so lies weakly above $\bar{\mu}^{[j]}$. Because $\delta(1-\lambda) < 1$ and $j \leq T$, we have $(\delta(1-\lambda))^{j-1} \geq (\delta(1-\lambda))^{T-1}$, so for every j,

$$\bar{\mu}^{[j]} \leq \mu^{att} - \frac{(1-\delta)(\delta(1-\lambda))^{T-1}(1-c/a)b}{(s_1 - s_0)a}. \tag{A.6}$$

We know by inequality (2) that $b < (c - s_0 a)/(1-s_0)$, and $c > s_0 a$ implies $(1-c/a)/(1-s_0) \leq 1$. Then, the fraction in (A.6) is less than $(c - s_0 a)/[(s_1 - s_0)a] = \mu^{att}$, and the bound on the funding threshold is strictly above zero.

Now fix $j \geq 2$ and take any $\mu < \bar{\mu}^{[j-1]}$. Then $\Psi^{[j-1]}(\mu) < 0$ and hence $\Phi^{[j-1]}(\mu) = 0$. Moreover, because $\mu^-(\mu') \leq \mu'$ at every belief, $\mu < \bar{\mu}^{[j-1]} \leq (\mu^-)^{-1}(\bar{\mu}^{[j-1]}) = \mu^{*[j]}$, and hence $\Psi^{[j]}(\mu) = \Psi^{[j]}_{att}(\mu)$. With j periods remaining, the surplus from funding at μ is evaluated under pooling on attack. Bounding $s(\mu) \leq 1$ and $\Phi^{[j-1]}(\mu^+(\mu)) \leq (a-c)/(1-\delta)$,

$$\Psi^{[j]}(\mu) = (s_1 - s_0)a(\mu - \mu^{att}) + \delta\lambda s(\mu)\Phi^{[j-1]}(\mu^+(\mu)) \leq -(s_1 - s_0)a(\mu^{att} - \mu) + \frac{\delta\lambda(a-c)}{1-\delta},$$

and because $\mu < \bar{\mu}^{[j-1]}$, inequality (A.6) implies

$$\mu^{att} - \mu > \frac{(1-\delta)(\delta(1-\lambda))^{T-1}(1-c/a)b}{(s_1 - s_0)a},$$

and thus

$$\Psi^{[j]}(\mu) < -(1-\delta)(\delta(1-\lambda))^{T-1}\left(1-\frac{c}{a}\right)b + \frac{\delta\lambda(a-c)}{1-\delta}.$$

Now define

$$\tilde{\lambda} := \min\left\{\frac{1}{2}, \frac{(1-\delta)^2(\delta/2)^{T-1}(1-c/a)b}{\delta(a-c)}\right\} > 0.$$

For every $\lambda < \tilde{\lambda}$, we have $1-\lambda > 1/2$ and hence $(\delta(1-\lambda))^{T-1} > (\delta/2)^{T-1}$, while $\lambda < (1-\delta)^2(\delta/2)^{T-1}(1-c/a)b/[\delta(a-c)]$ implies

$$\frac{\delta\lambda(a-c)}{1-\delta} < (1-\delta)(\delta/2)^{T-1}(1-c/a)b.$$

Putting these together, we have

$$\frac{\delta\lambda(a-c)}{1-\delta} < (1-\delta)\left(\frac{\delta}{2}\right)^{T-1}\left(1-\frac{c}{a}\right)b < (1-\delta)(\delta(1-\lambda))^{T-1}\left(1-\frac{c}{a}\right)b,$$

and therefore $\Psi^{[j]}(\mu) < 0$. Thus no belief below $\bar{\mu}^{[j-1]}$ is funded with j periods remaining, and $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$. □

Note that Proposition 5 pins down where the belief that governs the recipient's sincere threshold relates to μ^{att}. For every λ below its bound, iterating $\bar{\mu}^{[j-1]} \geq \bar{\mu}^{[j-2]} \geq \cdots \geq \bar{\mu}^{[1]} = \mu^{sin}$ and applying Lemma 2 yields

$$\mu^{*[j]} = (\mu^-)^{-1}(\bar{\mu}^{[j-1]}) \geq (\mu^-)^{-1}(\mu^{sin}) > \mu^{att} \tag{A.7}$$

for every $j \geq 2$, so every belief with $\mu \geq \mu^{*[j]}$ has $u^{sin}(\mu) > s(\mu)a > c$.

B.2 Lemma 3

This next result provides the key bound behind the two results that follow: (i) that the donor's surplus is larger under sincere recipient play than when they always attack, and (ii) that it is increasing in the donor's belief. To achieve this, we construct a function $R^{[j]}$ to capture the additional value the donor derives from sincere recipient play.

Lemma 3. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T-j+1$. Define $R^{[j]}(\mu) := \Phi^{[j]}(\mu) - \delta\Phi^{[j-1]}(\mu) - s(\mu)a + c$ for $j \geq 1$. Then, there exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, we have the following:*

(i) If $j \geq 2$ and $\mu \geq \mu^{*[j]}$, then

$$R^{[j]}(\mu) = (1-\delta)(1-s(\mu))b + \delta\left[s(\mu)R^{[j-1]}(\mu^+(\mu)) + (1-s(\mu))R^{[j-1]}(\mu^-(\mu))\right].$$

(ii) $R^{[j]}(\mu) \geq -\delta\lambda(a-c)/(1-\delta)$ for every μ.

Proof. First, take $\mu \geq \mu^{*[j]}$. At $j = 2$, $\mu^{*[1]} = 0$, and at $j \geq 3$, Proposition 5 and the increasing map $(\mu^-)^{-1}$ give $\mu^{*[j]} \geq \mu^{*[j-1]}$, so both types of recipient play sincerely at μ in stages j and $j-1$. By inequality (A.7), $\mu \geq \mu^{*[j]} > \mu^{sin}$, so $\Psi^{[j]}(\mu)$ and $\Psi^{[j-1]}(\mu)$ are each at least $u^{sin}(\mu) - c > 0$, and hence $\Phi^{[j]}(\mu) = \Psi^{[j]}_{sin}(\mu)$ and $\Phi^{[j-1]}(\mu) = \Psi^{[j-1]}_{sin}(\mu)$. Substituting and simplifying into the definition of $R^{[j]}$ recovers part (i).

For $\mu < \mu^{*[j]}$ with $\Psi^{[j]}(\mu) \geq 0$, where both types of recipient pool on attack, the funded surplus is

$$\Phi^{[j]}(\mu) = s(\mu)a - c + \delta\left[(1-\lambda)\Phi^{[j-1]}(\mu) + \lambda s(\mu)\Phi^{[j-1]}(\mu^+(\mu))\right].$$

Subtracting $\delta\Phi^{[j-1]}(\mu) + s(\mu)a - c$ results in

$$R^{[j]}(\mu) = \delta\lambda\left[s(\mu)\Phi^{[j-1]}(\mu^+(\mu)) - \Phi^{[j-1]}(\mu)\right].$$

We now bound $R^{[j]}$ below by induction on j. At $j = 1$ there is no continuation, so $\Phi^{[1]}(\mu) = \max\{0, u^{sin}(\mu) - c\}$ and $\Phi^{[0]}(\mu) = 0$. At every μ with $\Psi^{[1]}(\mu) \geq 0$,

$$R^{[1]}(\mu) = \Phi^{[1]}(\mu) - \delta\Phi^{[0]}(\mu) - (s(\mu)a - c) = (1-s(\mu))b \geq 0.$$

At every μ with $\Psi^{[1]}(\mu) < 0$, we instead have $\Phi^{[1]}(\mu) = 0$ and $R^{[1]}(\mu) = c - s(\mu)a > 0$, because $u^{sin}(\mu) < c$ implies $s(\mu)a < c$. Now suppose $R^{[j-1]}(\mu) \geq -\delta\lambda(a-c)/(1-\delta)$ at every μ. For $\mu < \mu^{*[j]}$ with $\Psi^{[j]}(\mu) \geq 0$,

$$R^{[j]}(\mu) = \delta\lambda\left[s(\mu)\Phi^{[j-1]}(\mu^+(\mu)) - \Phi^{[j-1]}(\mu)\right] \geq \delta\lambda\left(0 - \frac{a-c}{1-\delta}\right) = -\frac{\delta\lambda(a-c)}{1-\delta}.$$

If instead $\Psi^{[j]}(\mu) < 0$, then $\mu < \mu^{*[j]}$ by inequality (A.7), as beliefs of at least $\mu^{*[j]}$ have $\Psi^{[j]}(\mu) \geq u^{sin}(\mu) - c > 0$. Then $\Phi^{[j]}(\mu) = 0$, so $R^{[j]}(\mu) = c - s(\mu)a - \delta\Phi^{[j-1]}(\mu)$, while $\Psi^{[j]}(\mu) = \Psi^{[j]}_{att}(\mu) < 0$ gives

$$c - s(\mu)a > \delta(1-\lambda)\Phi^{[j-1]}(\mu) + \delta\lambda s(\mu)\Phi^{[j-1]}(\mu^+(\mu)).$$

Substituting this into $R^{[j]}(\mu)$ and simplifying,

$$R^{[j]}(\mu) > \delta\lambda\left[s(\mu)\Phi^{[j-1]}(\mu^+(\mu)) - \Phi^{[j-1]}(\mu)\right] \geq -\frac{\delta\lambda(a-c)}{1-\delta}.$$

Moreover, for $\mu \geq \mu^{*[j]}$, part (i) with $R^{[j-1]}(\mu^{\pm}(\mu)) \geq -\delta\lambda(a-c)/(1-\delta)$ implies

$$R^{[j]}(\mu) \geq (1-\delta)(1-s(\mu))b - \frac{\delta^2\lambda(a-c)}{1-\delta} \geq -\frac{\delta^2\lambda(a-c)}{1-\delta} \geq -\frac{\delta\lambda(a-c)}{1-\delta}.$$

These cases exhaust the beliefs, so $R^{[j]}(\mu) \geq -\delta\lambda(a-c)/(1-\delta)$ at every μ and every j, concluding part (ii). Both parts restrict λ only through Proposition 5, so $\tilde{\lambda}$ is the bound of Proposition 5. □

An arbitrary monotone equilibrium may pool on attack at a belief where the posterior after a flow payoff revealing the rain remains funded, adding the continuation term $\delta\lambda(1-s(\mu))\Phi^{[j-1]}(\mu^-(\mu))$ to equation (A.4). The next lemma shows the surplus under sincere play exceeds even this augmented surplus above $\mu^{*[j]}$.

B.3 Lemma 4

Lemma 4. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T - j + 1$. There exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, $j \geq 2$, and $\mu \geq \mu^{*[j]}$,*

$$\Psi^{[j]}_{sin}(\mu) - \Psi^{[j]}_{att}(\mu) > \delta\lambda(1-s(\mu))\Phi^{[j-1]}(\mu^-(\mu)).$$

Proof. Fix $j \geq 2$ and $\mu \geq \mu^{*[j]}$. Then,

$$\begin{aligned}
&\Psi^{[j]}_{sin}(\mu) - \Psi^{[j]}_{att}(\mu) - \delta\lambda(1-s(\mu))\Phi^{[j-1]}(\mu^-(\mu)) \\
&= (1-s(\mu))b + \delta(1-\lambda)\left[s(\mu)\Phi^{[j-1]}(\mu^+(\mu)) + (1-s(\mu))\Phi^{[j-1]}(\mu^-(\mu)) - \Phi^{[j-1]}(\mu)\right] \\
&= (1-s(\mu))b(1-\delta(1-\lambda)) + \delta(1-\lambda)\left[s(\mu)R^{[j-1]}(\mu^+(\mu)) + (1-s(\mu))R^{[j-1]}(\mu^-(\mu))\right],
\end{aligned}$$

where the second equality substitutes the definition of $R^{[j-1]}$ at each belief and the martingale identity $s(\mu)s(\mu^+(\mu)) + (1-s(\mu))s(\mu^-(\mu)) = s(\mu)$, using $\Phi^{[j-1]}(\mu) = \Psi^{[j-1]}(\mu)$, which follows from Proposition 5 and inequality (A.7). At $j = 2$, one period remains, both types play sincerely at every belief, and $\mu^+(\mu) > \mu^{sin}$, so the last-period case in the proof of Lemma 3 gives $R^{[1]}(\mu^+(\mu)) = (1-s(\mu^+(\mu)))b \geq 0$. At $j \geq 3$, $\mu^+(\mu) > \mu \geq \mu^{*[j-1]}$, so part (i) of Lemma 3 applies at $\mu^+(\mu)$, and part (ii) bounds its continuation terms, giving

$$R^{[j-1]}(\mu^+(\mu)) \geq (1-\delta)(1-s_1)b - \frac{\delta^2\lambda(a-c)}{1-\delta},$$

which is non-negative for every $\lambda \leq (1-\delta)^2(1-s_1)b/(\delta^2(a-c))$. At such λ, the $R^{[j-1]}(\mu^+(\mu))$ term is non-negative, part (ii) of Lemma 3 gives $R^{[j-1]}(\mu^-(\mu)) \geq -\delta\lambda(a-c)/(1-\delta)$, and thus

$$\Psi^{[j]}_{sin}(\mu) - \Psi^{[j]}_{att}(\mu) - \delta\lambda(1 - s(\mu))\Phi^{[j-1]}(\mu^-(\mu)) \\ \geq (1 - s(\mu)) \left[b(1 - \delta(1 - \lambda)) - \frac{\delta^2\lambda(1-\lambda)(a-c)}{1-\delta} \right].$$

The bracketed term is then strictly positive at every λ below the bound of the statement. Taking $\tilde{\lambda}$ to be the smaller of the bound of Proposition 5 and $(1-\delta)^2(1-s_1)b/(\delta^2(a-c))$ completes the proof. □

B.4 Lemma 5

Lemma 5. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T - j + 1$. There exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$ and every j, the donor's surplus $\Phi^{[j]}$ is weakly increasing in μ.*

Proof. We argue by induction on j. In the terminal period such that $j = 1$, we have $\Phi^{[1]}(\mu) = \max\{0, u^{sin}(\mu) - c\}$, which is weakly increasing in μ. Now fix $j \geq 2$ and suppose $\Phi^{[j-1]}$ is weakly increasing in μ.

First, take $\mu < \mu^{*[j]}$, where both types of recipient pool on attack. The expected flow $s(\mu)a$ is increasing in μ. Next period, the belief stays at μ with probability $1 - \lambda$, rises to $\mu^+(\mu)$ with probability $\lambda s(\mu)$, and the surplus otherwise ends at zero. A higher μ raises both continuing beliefs and shifts weight from the zero outcome to $\mu^+(\mu)$. Because $\Phi^{[j-1]}$ is weakly increasing and non-negative, both effects raise the continuation term, so equation (A.4) is weakly increasing in μ. For $\mu \geq \mu^{*[j]}$, both types play sincerely. The expected flow $s(\mu)a + (1 - s(\mu))b$ is increasing in μ, and next period's belief is $\mu^+(\mu)$ with probability $s(\mu)$ and $\mu^-(\mu)$ otherwise. A higher μ raises both $\mu^+(\mu)$ and $\mu^-(\mu)$ and shifts weight from $\mu^-(\mu)$ to $\mu^+(\mu)$, so with $\Phi^{[j-1]}$ weakly increasing, equation (A.3) is weakly increasing in μ, as well.

It remains to show $\Psi^{[j]}$ does not fall as μ crosses $\mu^{*[j]}$. At $\mu^{*[j]}$, we have $\Psi^{[j]}(\mu^{*[j]}) = \Psi^{[j]}_{sin}(\mu^{*[j]}) \geq u^{sin}(\mu^{*[j]}) - c > 0$ by inequality (A.7), so $\Phi^{[j]}(\mu^{*[j]}) = \Psi^{[j]}(\mu^{*[j]})$. Because equation (A.4) is weakly increasing in μ, it therefore suffices to show $\Psi^{[j]}_{att}(\mu^{*[j]}) \leq \Psi^{[j]}_{sin}(\mu^{*[j]})$. Lemma 4 and $\Phi^{[j-1]} \geq 0$ give exactly this. Then, $\Psi^{[j]}_{att}(\mu^{*[j]}) < \Psi^{[j]}_{sin}(\mu^{*[j]})$, the surplus at every $\mu < \mu^{*[j]}$ is at most $\Phi^{[j]}(\mu^{*[j]})$, and $\Phi^{[j]}$ is weakly increasing on all of $[0, 1]$.

Taking $\tilde{\lambda}$ to be the bound of Lemma 4, which lies below the bound of Proposition 5, completes the proof. □

B.5 Lemma 6

Lemma 6. *At every $j \geq 1$, $\bar{\mu}^{[j]}(\lambda) \to \bar{\mu}^{[j]}(0)$ and $\mu^{*[j]}(\lambda) \to \mu^{*[j]}(0)$ as $\lambda \to 0$.*

Proof. Throughout the proof, we write $\Psi^{[j]}(\mu; \lambda)$ and $\bar{\mu}^{[j]}(\lambda)$ to make dependence on λ explicit. We show by induction on j that $\bar{\mu}^{[j]}(\lambda) \to \bar{\mu}^{[j]}(0)$ as $\lambda \to 0$ and that there is a finite set of beliefs M_j such that $\Phi^{[j]}(\mu;\lambda) \to \Phi^{[j]}(\mu;0)$ at every $\mu \notin M_j$. At $j=1$, we have $\Psi^{[1]}(\mu;\lambda) = u^{sin}(\mu) - c$ at every λ, and hence $\bar{\mu}^{[1]}(\lambda) = \mu^{sin}$ at every λ and $M_1 = \emptyset$.

Now fix $j \geq 2$ and suppose the claims hold at $j-1$. Define

$$M_j := \{\mu^{*[j]}(0)\} \cup M_{j-1} \cup (\mu^+)^{-1}(M_{j-1}) \cup (\mu^-)^{-1}(M_{j-1}),$$

which is finite because μ^+ and μ^- are strictly increasing. Take any $\mu \notin M_j$. Because $\bar{\mu}^{[j-1]}(\lambda) \to \bar{\mu}^{[j-1]}(0)$ and $(\mu^-)^{-1}$ is continuous, we have $\mu^{*[j]}(\lambda) \to \mu^{*[j]}(0)$, and because $\mu \neq \mu^{*[j]}(0)$, the belief μ lies on the same side of $\mu^{*[j]}(\lambda)$ as of $\mu^{*[j]}(0)$ for every sufficiently small λ. Suppose first $\mu > \mu^{*[j]}(0)$, so that equation (A.3) gives $\Psi^{[j]}(\mu;\lambda)$ at every sufficiently small λ. By the definition of M_j, neither $\mu^+(\mu)$ nor $\mu^-(\mu)$ lies in M_{j-1}, so the inductive hypothesis gives $\Phi^{[j-1]}(\mu^{\pm}(\mu);\lambda) \to \Phi^{[j-1]}(\mu^{\pm}(\mu);0)$, and thus $\Psi^{[j]}(\mu;\lambda) \to \Psi^{[j]}(\mu;0)$. Suppose instead $\mu < \mu^{*[j]}(0)$, so that equation (A.4) gives $\Psi^{[j]}(\mu;\lambda)$ at every sufficiently small λ. Equation (A.4) evaluates $\Phi^{[j-1]}$ at μ with coefficient $\delta(1-\lambda)$ and at $\mu^+(\mu)$ with coefficient $\delta\lambda s(\mu)$. The first term converges to $\delta\Phi^{[j-1]}(\mu;0)$, because $\mu \notin M_{j-1}$, and the second vanishes, because $\Phi^{[j-1]} \leq (a-c)/(1-\delta)$. Then, $\Psi^{[j]}(\mu;\lambda) \to s(\mu)a - c + \delta\Phi^{[j-1]}(\mu;0) = \Psi^{[j]}(\mu;0)$. In either case, $\Phi^{[j]} = \max\{0, \Psi^{[j]}\}$ and $|\max\{0,x\} - \max\{0,y\}| \leq |x-y|$ gives $\Phi^{[j]}(\mu;\lambda) \to \Phi^{[j]}(\mu;0)$.

It remains to show $\bar{\mu}^{[j]}(\lambda) \to \bar{\mu}^{[j]}(0)$. By Proposition 5 and inequality (A.7), we have $\mu^{sin} \leq \bar{\mu}^{[j]}(0) \leq \mu^{att} < \mu^{*[j]}(0)$, and because M_j is finite, we can fix an $\varepsilon > 0$ such that $\bar{\mu}^{[j]}(0) - \varepsilon > 0$ and $\bar{\mu}^{[j]}(0) + \varepsilon < \mu^{*[j]}(0)$, with neither belief in M_j. By definition of $\bar{\mu}^{[j]}(0)$, we have $\Psi^{[j]}(\mu;0) < 0$ at every $\mu < \bar{\mu}^{[j]}(0)$. Now take any $\mu \in (\bar{\mu}^{[j]}(0), \mu^{*[j]}(0))$. Equation (A.4) applies at μ and, at $\lambda = 0$, reads $s(\mu)a - c + \delta\Phi^{[j-1]}(\mu;0)$, which is strictly increasing in μ by Lemma 5. By definition of $\bar{\mu}^{[j]}(0)$, there is a $\mu' \in [\bar{\mu}^{[j]}(0), \mu)$ with $\Psi^{[j]}(\mu';0) \geq 0$, and hence $\Psi^{[j]}(\mu;0) > 0$. Because neither $\bar{\mu}^{[j]}(0) - \varepsilon$ nor $\bar{\mu}^{[j]}(0) + \varepsilon$ lies in M_j, the convergence above gives $\Psi^{[j]}(\bar{\mu}^{[j]}(0) - \varepsilon;\lambda) < 0$ and $\Psi^{[j]}(\bar{\mu}^{[j]}(0) + \varepsilon;\lambda) > 0$ for every sufficiently small λ. The function $\Psi^{[j]}$ is weakly increasing in the belief at every sufficiently small λ, as shown in the proof of Lemma 5, and thus $\bar{\mu}^{[j]}(\lambda)$ lies within ε of $\bar{\mu}^{[j]}(0)$ for every sufficiently small λ. Because ε can be taken arbitrarily small, $\bar{\mu}^{[j]}(\lambda) \to \bar{\mu}^{[j]}(0)$, and $\mu^{*[j]}(\lambda) \to \mu^{*[j]}(0)$ follows as above, which completes the induction. $\square$

C Properties of monotone equilibria

In this subsection, we fix an arbitrary monotone equilibrium, with a funding threshold $\bar{\mu}^{[j]}$ and a belief at which the recipient goes from pooling on attack to defending in the rain that satisfies $\mu^{*[j]} \geq (\mu^-)^{-1}(\bar{\mu}^{[j-1]})$, where $j = T - t + 1$ denotes the number of periods remaining.

Two value formulas recur throughout. For each type θ and horizon j, define

$$\bar{v}_\theta^{[j]} := (s_\theta a + (1-s_\theta)b)\frac{1-\delta^j}{1-\delta}; \qquad \underline{v}_\theta^{[j]} := s_\theta a \frac{1-\delta^j(1-\lambda(1-s_\theta))^j}{1-\delta(1-\lambda(1-s_\theta))}.$$

The first is the value of play that is always funded and sincere, and the second is the value of attacking in every period when funding is pulled if and only if a flow payoff revealing the rain is observed. Because any strategy receives a flow payoff of at most a in the shine, at most b in the rain, and zero once funding ends, the recipient's value at any belief satisfies $V_\theta^{[j]}(\mu) \leq \bar{v}_\theta^{[j]}$.

C.1 Lemma 7

Lemma 7. *There exists a $\tilde{\lambda} > 0$ such that, for both types θ, every $j \geq 1$, and every $\lambda < \tilde{\lambda}$, $\delta(\bar{v}_\theta^{[j]} - \underline{v}_\theta^{[j]}) < b$.*

Proof. The difference $\bar{v}_\theta^{[j]} - \underline{v}_\theta^{[j]}$ is at most its limit as j grows,

$$\frac{s_\theta a + (1-s_\theta)b}{1-\delta} - \frac{s_\theta a}{1-\delta(1-\lambda(1-s_\theta))} = \frac{(1-s_\theta)b}{1-\delta} + s_\theta a \frac{\delta\lambda(1-s_\theta)}{(1-\delta)\left(1-\delta(1-\lambda(1-s_\theta))\right)},$$

because the first formula falls short of its limit by more than the second. Bounding the final denominator below by $1-\delta$, the limit is at most

$$\frac{(1-s_\theta)b}{1-\delta} + \frac{\delta\lambda s_\theta a(1-s_\theta)}{(1-\delta)^2},$$

and thus

$$b - \delta(\bar{v}_\theta^{[j]} - \underline{v}_\theta^{[j]}) \geq b\frac{1-\delta(2-s_\theta)}{1-\delta} - \frac{\delta^2\lambda s_\theta a(1-s_\theta)}{(1-\delta)^2},$$

which is strictly positive at every $\lambda < (1-\delta)b\left(1-\delta(2-s_\theta)\right)/(\delta^2 s_\theta a(1-s_\theta))$, a bound that is strictly positive for both types by Assumption 2. Taking $\tilde{\lambda}$ to be the smaller of the two bounds completes the proof. □

C.2 Lemma 8

The following result establishes that a funded recipient of type θ can guarantee at least the value $\underline{v}_\theta^{[j]}$.

Lemma 8. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T-j+1$. Fix a monotone equilibrium and suppose that, at every j' with $2 \leq j' \leq j$ and every belief $\mu \geq \bar{\mu}^{[j']}$, the donor's posterior after an attack is at least $\bar{\mu}^{[j'-1]}$ both when only the action is observed, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j'-1]}$, and when the flow payoff reveals a shine, $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j'-1]}$. Then, $V_\theta^{[j]}(\mu) \geq \underline{v}_\theta^{[j]}$ for each type θ and every $\mu \geq \bar{\mu}^{[j]}$.*

Proof. Fix a type θ and consider the deviation of attacking in every remaining period regardless of the weather, which collects an expected flow payoff of $s_\theta a$ each period. With probability λs_θ the flow payoff reveals a shine and the posterior is $\mu^{1,1}(\mu)$, with probability $1-\lambda$ only the action is observed and the posterior is $\mu^{1,\emptyset}(\mu)$, and with probability $\lambda(1-s_\theta)$ the flow payoff reveals the rain. By supposition, the first two posteriors are at least $\bar{\mu}^{[j-1]}$ and funding continues, so funding survives each period with probability at least $1-\lambda(1-s_\theta)$. We argue by induction on j that the deviation is worth at least $\underline{v}_\theta^{[j]}$ at every $\mu \geq \bar{\mu}^{[j]}$. At $j=1$, the deviation yields $s_\theta a = \underline{v}_\theta^{[1]}$. At $j \geq 2$, the surviving posteriors are at least $\bar{\mu}^{[j-1]}$, where the deviation is worth at least $\underline{v}_\theta^{[j-1]}$ by the inductive hypothesis, and the continuation after a revealed rain is worth at least zero, so the deviation is worth at least $s_\theta a + \delta(1-\lambda(1-s_\theta))\underline{v}_\theta^{[j-1]} = \underline{v}_\theta^{[j]}$. The equilibrium value is at least the value of this deviation. $\square$

C.3 Lemma 9

Next, we establish that both types attack in the shine at every funded belief in any monotone equilibrium, so that a defense can only occur in the rain.

Lemma 9. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T-j+1$. Suppose $j \geq 2$, fix a monotone equilibrium, and suppose that, at every j' with $2 \leq j' \leq j-1$ and every belief $\mu \geq \bar{\mu}^{[j']}$, the donor's posterior after an attack is at least $\bar{\mu}^{[j'-1]}$ both when only the action is observed, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j'-1]}$, and when the flow payoff reveals a shine, $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j'-1]}$. Then, there exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, both types of recipient attack in the shine at every belief at which $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$ and $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j-1]}$.*

Proof. Take $\tilde{\lambda}$ to be the bound of Lemma 7, and fix a type θ and a belief μ at which $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$ and $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j-1]}$. Attacking in the shine yields a flow payoff of a and, whether or not the flow payoff is observed, a funded posterior at which Lemma 8 applies at $j-1$, so attacking is worth at least $a + \delta \underline{v}_\theta^{[j-1]}$. Defending yields a flow payoff of zero and a continuation of at most $\bar{v}_\theta^{[j-1]}$. By Lemma 7, the payoff from attacking exceeds the payoff from defending by at least $a - \delta(\bar{v}_\theta^{[j-1]} - \underline{v}_\theta^{[j-1]}) > a - b > 0$. $\square$

C.4 Lemma 10

Now we show that, at low levels of monitoring, both types must play the same pure action, attack or defend, in the rain at every funded belief. Specifically, we show that (i) the strong type will not mix, (ii) the weak type will not be the only one to defend, (iii) the strong type will not be the only one to defend, whether the weak attacks purely or mixes.

Lemma 10. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T-j+1$. Suppose $j \geq 2$, fix a monotone equilibrium with $\bar{\mu}^{[j-1]} > 0$, and suppose that, at every j' with $2 \leq j' \leq j-1$ and every belief $\mu \geq \bar{\mu}^{[j']}$, the donor's posterior after an attack is at least $\bar{\mu}^{[j'-1]}$ both when only the action is observed, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j'-1]}$, and when the flow payoff reveals a shine, $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j'-1]}$. Then, there exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, both types of recipient play the same action in the rain with probability one, $\sigma_{00}(\mu) = \sigma_{10}(\mu) \in \{0,1\}$, at every funded belief μ at which $\mu \geq \bar{\mu}^{[j-1]}$, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$, and $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j-1]}$.*

Proof. At $\mu = 1$ the donor's belief is degenerate, the posterior is 1 after every observation, and the belief 1 is always funded. Defending in the rain yields a flow payoff of b and the continuation value at the belief 1, whereas attacking yields a flow payoff of zero and the same continuation value, and thus both types defend in the rain and we have $\sigma_{00}(1) = \sigma_{10}(1) = 1$.

Now fix a funded belief $\mu < 1$ at which $\mu \geq \bar{\mu}^{[j-1]}$, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$, and $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j-1]}$. The conditions of Lemma 9 thus hold, so both types attack in the shine at μ. Then, a defense occurs only in the rain, and the posterior after a defense is the same regardless of whether the flow payoff is observed, $\mu^{0,\emptyset}(\mu) = \mu^{0,0}(\mu)$. Because the recipient can play sincerely in the current period regardless of how play continues, the value at any funded belief is at least $s_\theta a + (1 - s_\theta)b$. The conditions of Lemma 8 also hold at $j - 1$, so $V_\theta^{[j-1]} \geq \underline{v}_\theta^{[j-1]}$ at every belief weakly above $\bar{\mu}^{[j-1]}$.

First, we show the strong type does not mix in the rain. Suppose $\sigma_{10}(\mu) \in (0, 1)$, and then the strong type is indifferent between defending and attacking in the rain,

$$b + \delta V_1^{[j-1]}(\mu^{0,0}(\mu)) = \delta \left[\lambda V_1^{[j-1]}(\mu^{1,0}(\mu)) + (1 - \lambda)V_1^{[j-1]}(\mu^{1,\emptyset}(\mu))\right]. \tag{A.8}$$

Take first $\mu^{0,0}(\mu) \geq \bar{\mu}^{[j-1]}$. Each continuation value on the right-hand side of equality (A.8) is at most $\bar{v}_1^{[j-1]}$, whereas $V_1^{[j-1]}(\mu^{0,0}(\mu)) \geq \underline{v}_1^{[j-1]}$, and thus the equality requires $b \leq \delta(\bar{v}_1^{[j-1]} - \underline{v}_1^{[j-1]})$, which fails at every λ below the bound of Lemma 7, and equality (A.8) cannot hold. Take instead $\mu^{0,0}(\mu) < \bar{\mu}^{[j-1]}$. Because a defense occurs only in the rain, we have $\mu^{0,\emptyset}(\mu) = \mu^{0,0}(\mu) < \bar{\mu}^{[j-1]}$, and thus a defense ends funding regardless of whether the flow payoff is observed, and the left-hand side of equality (A.8) equals b. Because $V_1^{[j-1]}(\mu^{1,0}(\mu)) \geq 0$ and $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$ is funded, where playing sincerely for a single period yields an expected flow payoff of at least $s_1 a + (1 - s_1)b$, the right-hand side of equality (A.8) is at least $\delta(1 - \lambda)(s_1 a + (1 - s_1)b)$. Then, the equality requires $b \geq \delta(1 - \lambda)(s_1 a + (1 - s_1)b)$, which fails for every $\lambda < 1 - b/[\delta(s_1 a + (1 - s_1)b)]$. In either case the strong type cannot be indifferent, and thus $\sigma_{10}(\mu) \in \{0, 1\}$.

Next, suppose $\sigma_{10}(\mu) = 0$ and, for contradiction, $\sigma_{00}(\mu) > 0$. A defense in the rain then occurs only for the weak type, and thus $\mu^{0,0}(\mu) = 0 < \bar{\mu}^{[j-1]}$, funding ends, and defending yields a payoff of b. Attacking yields at least $\delta(1 - \lambda)V_0^{[j-1]}(\mu^{1,\emptyset}(\mu))$, and $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$ is funded, where playing sincerely for a single period yields an expected flow payoff of at least $s_0 a + (1 - s_0)b$. Then, attacking is worth at least $\delta(1 - \lambda)(s_0 a + (1 - s_0)b)$, which exceeds b for every $\lambda < 1 - b/[\delta(s_0 a + (1 - s_0)b)]$, and the right-hand side of this bound is strictly positive by inequality (3) of Assumption 1. The weak type then strictly prefers attacking, a contradiction, and thus $\sigma_{00}(\mu) = 0$.

Finally, suppose $\sigma_{10}(\mu) = 1$ and, for contradiction, $\sigma_{00}(\mu) < 1$. If $\sigma_{00}(\mu) = 0$, a defense occurs only for the strong type, and thus $\mu^{0,0}(\mu) = 1$, and the weak type's preference for attacking in the rain requires

$$\delta\left[\lambda V_0^{[j-1]}(\mu^{1,0}(\mu)) + (1-\lambda)V_0^{[j-1]}(\mu^{1,\emptyset}(\mu))\right] \geq b + \delta V_0^{[j-1]}(1).$$

The left-hand side is at most $\delta\bar{v}_0^{[j-1]}$, whereas the belief 1 is always funded and Lemma 8 gives $V_0^{[j-1]}(1) \geq \underline{v}_0^{[j-1]}$. The required preference thus needs $\delta(\bar{v}_0^{[j-1]} - \underline{v}_0^{[j-1]}) \geq b$, which fails at every λ below the bound of Lemma 7, a contradiction. With $\sigma_{00}(\mu) < 1$ maintained and $\sigma_{00}(\mu) = 0$ ruled out, we have $\sigma_{00}(\mu) \in (0, 1)$, and the weak type is indifferent between attacking and defending in the rain. Because $\sigma_{10}(\mu) = 1$, an attack whose flow payoff reveals the rain identifies the weak type, and thus $\mu^{1,0}(\mu) = 0 < \bar{\mu}^{[j-1]}$ and the continuation there is zero. The indifference then reads

$$(1-\lambda)V_0^{[j-1]}(\mu^{1,\emptyset}(\mu)) - V_0^{[j-1]}(\mu^{0,0}(\mu)) = \frac{b}{\delta}. \tag{A.9}$$

Take first $\mu^{0,0}(\mu) < \bar{\mu}^{[j-1]}$. Then, we have $V_0^{[j-1]}(\mu^{0,0}(\mu)) = 0$, and equality (A.9) gives $V_0^{[j-1]}(\mu^{1,\emptyset}(\mu)) = b/[\delta(1-\lambda)] > 0$, and thus $\mu^{1,\emptyset}(\mu)$ is funded. Because the strong type defends in the rain, attacking must yield a weakly worse payoff for the strong type,

$$b = b + \delta V_1^{[j-1]}(\mu^{0,0}(\mu)) \geq \delta(1-\lambda)V_1^{[j-1]}(\mu^{1,\emptyset}(\mu)) \geq \delta(1-\lambda)(s_1 a + (1-s_1)b),$$

which fails for every $\lambda < 1 - b/[\delta(s_1 a + (1-s_1)b)]$ as above, a contradiction. Take instead $\mu^{0,0}(\mu) \geq \bar{\mu}^{[j-1]}$. Then, Lemma 8 gives $V_0^{[j-1]}(\mu^{0,0}(\mu)) \geq \underline{v}_0^{[j-1]}$ while $V_0^{[j-1]}(\mu^{1,\emptyset}(\mu)) \leq \bar{v}_0^{[j-1]}$, and thus equality (A.9) requires $b/\delta \leq \bar{v}_0^{[j-1]} - \underline{v}_0^{[j-1]}$, which fails at every λ below the bound of Lemma 7, a contradiction, and thus $\sigma_{00}(\mu) = 1$.

Every case above ends in a contradiction when λ lies below a strictly positive bound, and Lemma 9 holds below its own strictly positive bound. Taking $\tilde{\lambda}$ to be the smallest of all of these bounds, we have $\sigma_{00}(\mu) = \sigma_{10}(\mu) \in \{0, 1\}$ for every $\lambda < \tilde{\lambda}$ at every funded belief μ at which $\mu \geq \bar{\mu}^{[j-1]}$, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$, and $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j-1]}$. □

C.5 Lemma 11

This next result shows every monotone equilibrium's funding threshold is interior and at most μ^{att}, the a priori bound for our threshold comparisons and limit arguments.

Lemma 11. *Let $T < \infty$ and j denote the number of periods remaining so that the current period is $T - j + 1$. Fix a monotone equilibrium and suppose that, at every j' with $2 \leq j' \leq j - 1$ and every belief $\mu \geq \bar{\mu}^{[j']}$, the donor's posterior after an attack is at least $\bar{\mu}^{[j'-1]}$ both when only the action is observed, $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j'-1]}$, and when the flow payoff reveals a shine, $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j'-1]}$. Then, there exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, the funding threshold at j lies in $(0, \mu^{att}]$.*

Proof. First, we show the belief 0 is never funded for any j. At $\mu = 0$ the donor's belief is degenerate and the posterior is 0 after every observation. Because attacking in the shine yields a and defending in the rain yields b, the donor's expected flow payoff is at most $s_0a+(1-s_0)b$ under any recipient play, and Assumption 1 implies $s_0a+(1-s_0)b<c$.

We now argue by induction on j that the donor strictly prefers not to fund at $\mu=0$. At $j=1$, funding yields at most $s_0a+(1-s_0)b<c$. Now fix $j\geq 2$ and suppose not funding is strictly preferred at $\mu=0$ with $j-1$ periods remaining. Funding at $\mu=0$ yields a flow of at most $s_0a+(1-s_0)b$ and returns the donor to the belief 0 with $j-1$ periods remaining, where not funding is strictly preferred, so

$$s_0a+(1-s_0)b+\delta c\frac{1-\delta^{j-1}}{1-\delta}<c+\delta c\frac{1-\delta^{j-1}}{1-\delta}=c\frac{1-\delta^{j}}{1-\delta},$$

completing the induction. The belief 0 is therefore never funded for any j, and $\bar{\mu}^{[j]}>0$.

Next, we show the donor funds at every belief $\mu>\mu^{att}$, and thus $\bar{\mu}^{[j]}\leq\mu^{att}$. Fix $\mu>\mu^{att}$ and consider the recipient's play in the shine at μ, whether μ is funded or the donor reaches it by deviating to fund there. At $j=1$ there is no continuation, and thus attacking yields a flow payoff of $a>0$ whereas defending yields zero, and both types attack in the shine. Now take $j\geq 2$ and consider three cases for the posteriors after an attack at μ.

First, suppose $\mu^{1,\emptyset}(\mu)\geq\bar{\mu}^{[j-1]}$. The supposition here is the supposition of Lemma 8 at $j-1$, and thus $V_\theta^{[j-1]}\geq\underline{v}_\theta^{[j-1]}$ at every belief weakly above $\bar{\mu}^{[j-1]}$. Attacking in the shine yields a flow payoff of a and, when only the action is observed, which occurs with probability $1-\lambda$, the funded posterior $\mu^{1,\emptyset}(\mu)$, whereas defending yields a flow payoff of zero and a continuation of at most $\bar{v}_\theta^{[j-1]}$. Then, the payoff from attacking exceeds the payoff from defending by at least

$$a-\delta\left(\bar{v}_\theta^{[j-1]}-\underline{v}_\theta^{[j-1]}\right)-\delta\lambda\underline{v}_\theta^{[j-1]}.$$

By Lemma 7, the middle term is below b at every λ below its bound, and $\underline{v}_\theta^{[j-1]}\leq s_\theta a/(1-\delta)\leq(s_1a+(1-s_1)b)/(1-\delta)$, so the payoff from attacking exceeds the payoff from defending by more than $a-b-\delta\lambda(s_1a+(1-s_1)b)/(1-\delta)$, which is strictly positive for every $\lambda<(a-b)(1-\delta)/[\delta(s_1a+(1-s_1)b)]$, itself strictly positive because $a>b$. Both types attack in the shine.

Second, suppose $\mu^{1,\emptyset}(\mu)<\bar{\mu}^{[j-1]}\leq\mu^{1,1}(\mu)$. By part (i) of Definition 1, we have $\mu^{0,\emptyset}(\mu)\leq\mu^{1,\emptyset}(\mu)<\bar{\mu}^{[j-1]}$, and thus funding ends after either action when only the action is observed. Attacking yields a flow payoff of a and, when the flow payoff reveals the shine, the funded posterior $\mu^{1,1}(\mu)$, where Lemma 8 bounds the continuation value from below, whereas defending yields a flow payoff of zero and, when the flow payoff reveals the shine, a continuation value of at most $\bar{v}_\theta^{[j-1]}$. Then, the payoff from attacking exceeds the payoff from defending by at least $a-\delta\lambda(\bar{v}_\theta^{[j-1]}-\underline{v}_\theta^{[j-1]})$, which exceeds $a-\lambda b$ at every λ below the bound of Lemma 7, strictly positive because $a>b$. Both types attack in the shine.

Third, suppose $\mu^{1,\emptyset}(\mu) < \bar{\mu}^{[j-1]}$ and $\mu^{1,1}(\mu) < \bar{\mu}^{[j-1]}$. By part (i) of Definition 1, every posterior after either action at the shine lies below $\bar{\mu}^{[j-1]}$, and thus funding ends after every realization and the payoff comparison involves flow payoffs only. Attacking yields a whereas defending yields zero, and both types attack in the shine.

In every case, both types attack in the shine at μ. Then, funding for one period yields the donor an expected flow payoff of at least $s(\mu)a$, and because s is increasing and $s(\mu^{att})a = c$, we have $s(\mu)a > c$. Because the donor can stop funding in every subsequent period, the value of funding at μ is at least

$$s(\mu)a + \delta c \frac{1-\delta^{j-1}}{1-\delta} > c + \delta c \frac{1-\delta^{j-1}}{1-\delta} = c\frac{1-\delta^{j}}{1-\delta},$$

the value of not funding. The donor therefore funds at every $\mu > \mu^{att}$, and $\bar{\mu}^{[j]} \leq \mu^{att}$. Taking $\tilde{\lambda}$ to be the smaller of the bound of Lemma 7 and $(a-b)(1-\delta)/[\delta(s_1 a + (1-s_1)b)]$ completes the proof. □

C.6 Proposition 6

Proposition 6. *Let $T < \infty$. There exists a $\tilde{\lambda} > 0$ such that, for every $\lambda < \tilde{\lambda}$, the donor's funding threshold in any monotone equilibrium is weakly decreasing over time, with $\bar{\mu}_t \geq \bar{\mu}_{t+1}$ for all $t < T$.*

Proof. Let $j = T - t + 1$ denote the number of periods remaining, and write the thresholds as $\bar{\mu}^{[j]}$, so that the claim is $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$ for every $j \geq 2$. Let $\tilde{\lambda}$ be the smallest of the bounds of Lemmas 9, 10, and 11 and the two bounds

$$1 - \frac{b}{\delta(s_0 a + (1-s_0)b)} \qquad \text{and} \qquad \frac{a(1-\delta)}{\delta(s_1 a + (1-s_1)b)}.$$

Each is strictly positive, and thus $\tilde{\lambda} > 0$. Fix $\lambda < \tilde{\lambda}$.

At $j = 1$ there is no continuation, and thus attacking in the shine yields a flow payoff of a whereas defending yields zero, and defending in the rain yields a flow payoff of b whereas attacking yields zero. Because $a > 0$ and $b > 0$, both types play sincerely at every funded belief. The donor's flow payoff from funding is then $u^{sin}(\mu)$ against c, and because the funded set is an interval $[\bar{\mu}^{[1]}, 1]$ by part (iii) of Definition 1, optimality gives $\bar{\mu}^{[1]} = \mu^{sin}$.

We argue by induction on $j \geq 2$ that $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$, assuming the ordering holds below j. Under this hypothesis, we have $\bar{\mu}^{[j-1]} \geq \bar{\mu}^{[1]} = \mu^{sin} > 0$. Moreover, at every $k < j$ and every funded belief $\mu \geq \bar{\mu}^{[k]} \geq \bar{\mu}^{[k-1]}$, the prior lies between the two posteriors after each action, and part (i) of Definition 1 implies the posterior after an attack is weakly larger than the posterior after a defense. Then, we have $\mu^{1,\emptyset}(\mu) \geq \mu \geq \bar{\mu}^{[k-1]}$. Likewise, conditional on the flow payoff revealing a shine, the prior updates to $\mu^{+}(\mu)$, which lies between the two posteriors after each action, and part (i) of Definition 1 then implies $\mu^{1,1}(\mu) \geq \mu^{+}(\mu) > \mu \geq \bar{\mu}^{[k-1]}$. This verifies the suppositions of Lemmas 8, 9, 10, and 11 at every $k \leq j-1$, and Lemma 11 gives $\bar{\mu}^{[k]} \in (0, \mu^{att}]$ at every $k \leq j-1$.

The supposition of Lemma 11 at j is the supposition of Lemma 8 at $j-1$, which is verified above, and thus Lemma 11 applies at j as well. Then, we have $\bar{\mu}^{[j]} \leq \mu^{att}$. By $\bar{\mu}^{[j-1]} \geq \mu^{sin}$ and Lemma 2, we have $(\mu^-)^{-1}(\bar{\mu}^{[j-1]}) > \mu^{att}$, and thus every belief in $(\mu^{att}, (\mu^-)^{-1}(\bar{\mu}^{[j-1]}))$ is funded.

First, we show that both types attack in the rain at every belief in this interval. Fix $\mu_1 \in (\mu^{att}, (\mu^-)^{-1}(\bar{\mu}^{[j-1]}))$. The belief μ_1 is funded, and the argument above gives $\mu^{1,\emptyset}(\mu_1) \geq \mu_1 > \mu^{att} \geq \bar{\mu}^{[j-1]}$ and $\mu^{1,1}(\mu_1) \geq \mu^+(\mu_1) > \mu_1$. Then, Lemmas 9 and 10 apply at μ_1, and thus both types attack in the shine, and in the rain both types play the same action with probability one, $\sigma_{00}(\mu_1) = \sigma_{10}(\mu_1) \in \{0, 1\}$.

Suppose both types defend in the rain at μ_1. A defense then occurs only in the rain and at the same rate for both types, and thus its posterior is $\mu^-(\mu_1)$ regardless of whether the flow payoff is observed. Because $\mu_1 < (\mu^-)^{-1}(\bar{\mu}^{[j-1]})$ and μ^- is increasing, we have $\mu^-(\mu_1) < \bar{\mu}^{[j-1]}$, and thus funding ends and defending yields b. Now consider a deviation to attacking in the rain. Under the supposed play an attack occurs only in the shine, and thus when the donor observes only the action, Bayes' rule gives $\mu^{1,\emptyset}(\mu_1) = \mu^+(\mu_1) > \mu_1 > \bar{\mu}^{[j-1]}$, which is funded, where one sincere period guarantees an expected flow payoff of at least $s_0 a + (1-s_0)b$. Then, the deviation is worth at least $\delta(1-\lambda)(s_0 a + (1-s_0)b)$, which exceeds b because $\lambda < 1 - b/[\delta(s_0 a + (1-s_0)b)]$, a contradiction. Then, both types attack in the rain at μ_1.

Next suppose, for contradiction, that some belief $\mu_0 < \bar{\mu}^{[j-1]}$ is funded. By part (ii) of Definition 1, the probability that each type defends in the rain is weakly increasing in the donor's belief, and it is zero at $\mu_1 > \mu_0$. Then, both types attack in the rain at μ_0. In the shine at μ_0, suppose some type defends with positive probability. Both actions then occur with positive probability, and thus the prior is a weighted average of the two posteriors after each action, and part (i) of Definition 1 implies $\mu^{0,\emptyset}(\mu_0) \leq \mu_0 < \bar{\mu}^{[j-1]}$. Then, the continuation after a defense is zero when only the action is observed. Because any strategy receives a flow payoff of at most a in the shine and at most b in the rain, and because $a > b$ implies $s_\theta a + (1-s_\theta)b \leq s_1 a + (1-s_1)b$, the continuation after a defense whose flow payoff is observed is at most $(s_1 a + (1-s_1)b)/(1-\delta)$. Then, defending in the shine is worth at most $\delta\lambda(s_1 a + (1-s_1)b)/(1-\delta)$, which is less than a because $\lambda < a(1-\delta)/[\delta(s_1 a + (1-s_1)b)]$, whereas attacking yields at least a, a contradiction. Then, both types attack in the shine at μ_0, and thus both types attack regardless of the weather.

An attack is then uninformative, with $\mu^{1,\emptyset}(\mu_0) = \mu_0 < \bar{\mu}^{[j-1]}$, and after an attack whose flow payoff reveals the rain, the posterior is $\mu^{1,0}(\mu_0) = \mu^-(\mu_0) < \bar{\mu}^{[j-1]}$. An attack in the rain at μ_0 therefore yields a flow payoff of zero and ends funding at both posteriors, and thus its payoff is zero, whereas defending yields at least $b > 0$. Both types then strictly prefer defending in the rain at μ_0, whereas we showed above that both types attack in the rain at

μ_0, a contradiction. Then, no belief below $\bar{\mu}^{[j-1]}$ is funded, which gives $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$ and completes the induction. □

D Full monitoring benchmark

When $\lambda = 1$, the donor observes the recipient's flow payoff, so the action-payoff pair reveals the weather. Bayes' rule then determines the posterior whenever the observed pair occurs with positive probability. If a pair is assigned probability zero by both types, we treat the commonly unused action as uninformative about type conditional on the revealed weather. That is, the donor assigns the posterior that would follow from observing the weather alone. We maintain this convention throughout this subsection.

Assumption 3 (Off-path beliefs)**.** *The donor's beliefs follow Bayes' rule given the recipient's strategies at every history, including histories the donor reaches only by deviating to fund. After an observation of an action-payoff pair assigned probability zero by both types of recipient, the donor's posterior is the posterior induced by the weather revealed by that pair.*

D.1 Proposition 7

Proposition 7 (Full Monitoring)**.** *Let $\lambda = 1$ and suppose Assumption 3 holds. Then, we have the following:*

(i) For every $T \leq \infty$, there is a monotone equilibrium in which both types of recipient play sincerely at every belief.

(ii) Fix $T < \infty$. For all but finitely many values of b, in every monotone equilibrium both types of recipient play sincerely at every funded belief, all monotone equilibria share the same funding thresholds, and every monotone equilibrium is most-sincere.

Proof. At $\lambda = 1$, the donor observes the flow payoff in every period. An attack yields a flow payoff of a in the shine and zero in the rain, and a defense yields b in the rain and zero in the shine, and thus every pair of action and flow payoff identifies the weather. Consider the strategy profile in which both types of recipient play sincerely at every belief, the donor's posterior is $\mu^+(\mu)$ after every observation identifying a shine and $\mu^-(\mu)$ after every observation identifying the rain, and the donor funds according to a threshold constructed below. Both types play the same action in each realization of the weather, and thus the action carries no information beyond the weather, the posteriors agree with Bayes' rule on the path of play, and the posterior after an attack equals the posterior after a defense at every belief.

First, sincere play is optimal for both types. In the rain, either action moves the belief to $\mu^-(\mu)$, so both actions yield the same continuation value and defending gains the flow payoff $b > 0$, and in the shine, either action moves the belief to $\mu^+(\mu)$ and attacking gains the flow payoff $a > 0$. Neither comparison involves the donor's funding, so no one-shot deviation from sincere play is profitable, and because flow payoffs are bounded and discounted, sincere play is optimal for both types at every belief and every $T \leq \infty$.

For $T < \infty$, consider the donor's funding problem against sincere play. Funding at belief μ yields the expected flow payoff $u^{sin}(\mu)$ and moves the belief to $\mu^+(\mu)$ with probability $s(\mu)$ and to $\mu^-(\mu)$ otherwise, and thus the surplus from funding and continuing optimally is equation (A.3) applied at every belief, with j periods remaining. We show by induction on j that $\Psi^{[j]}$ is continuous and strictly increasing in μ. At $j = 1$, we have $\Psi^{[1]} = u^{sin} - c$, which is affine and strictly increasing. Now fix $j \geq 2$ and suppose the claim holds at $j - 1$. Then, $\Phi^{[j-1]} = \max\{0, \Psi^{[j-1]}\}$ is continuous and weakly increasing, and μ^+ and μ^- are continuous and strictly increasing, and thus equation (A.3) is continuous in μ. A higher μ strictly raises $u^{sin}(\mu)$, raises both posteriors $\mu^+(\mu)$ and $\mu^-(\mu)$, and shifts weight from $\mu^-(\mu)$ to $\mu^+(\mu)$, where $\Phi^{[j-1]}(\mu^+(\mu)) \geq \Phi^{[j-1]}(\mu^-(\mu))$ by the inductive hypothesis, and thus $\Psi^{[j]}$ is strictly increasing, which completes the induction. By Assumption 1, we have $\Psi^{[j]}(1) \geq u^{sin}(1) - c > 0$, and the set of beliefs with $\Psi^{[j]}(\mu) \geq 0$ is a closed interval $[\bar{\mu}_t, 1]$ with $\bar{\mu}_t < 1$. The continuation $\Phi^{[j]} = \max\{0, \Psi^{[j]}\}$ is the value of continuing optimally, and thus funding on this interval is optimal, and a deviation to funding at an unfunded belief yields $\Psi^{[j]}(\mu) < 0$, because both types play sincerely regardless of funding.

For $T = \infty$, the donor's problem is stationary. Against sincere play, the surplus from funding and continuing optimally is stationary and solves

$$\Psi(\mu) = u^{sin}(\mu) - c + \delta \left[s(\mu)\Phi(\mu^+(\mu)) + (1 - s(\mu))\Phi(\mu^-(\mu))\right] \tag{A.10}$$

with $\Phi = \max\{0, \Psi\}$. The map sending a bounded function Φ to $\max\{0, \Psi\}$ with Ψ given by equation (A.10) is a contraction with modulus δ under the supremum norm, because the continuation weights sum to δ and $|\max\{0, x\} - \max\{0, y\}| \leq |x - y|$, so it has a unique bounded fixed point Φ. The image of the zero function is $\Phi^{[1]}$, and the image of $\Phi^{[j-1]}$ is $\Phi^{[j]}$, so the iterates from the zero function are the $\Phi^{[j]}$ and they converge uniformly to Φ. Each iterate is continuous and weakly increasing in μ, by the induction above, and thus so are the uniform limit Φ and, by equation (A.10), the fixed surplus Ψ. Because $\Phi \geq 0$, equation (A.10) gives $\Psi(1) \geq u^{sin}(1) - c > 0$, and the set of beliefs with $\Psi(\mu) \geq 0$ is a closed interval $[\bar{\mu}, 1]$ with $\bar{\mu} < 1$. Flow payoffs are bounded and discounted, and thus a funding strategy admitting no profitable one-shot deviation is optimal. At a funded belief, the value of funding and continuing optimally exceeds the value of stopping by $\Psi(\mu) \geq 0$, and at an unfunded belief, funding once and continuing optimally falls short of stopping by $-\Psi(\mu) > 0$, because both types play sincerely regardless of funding. Then, the threshold rule is optimal.

It remains to verify that the strategy profile is monotone. The posterior after an attack equals the posterior after a defense at every belief. The probability that each type defends in the rain is one at every belief, and thus weakly increasing in the donor's belief. The funded set is an interval $[\bar{\mu}_t, 1]$ in every period. Then, the strategy profile satisfies each part of Definition 1 and is a monotone equilibrium at every $T \leq \infty$, which is part (i).

Now fix $T < \infty$ and any monotone equilibrium. Call a funded belief $\mu \in (0,1)$ informative in period t if the action distributions of the two types at μ differ in some realization of the weather. In the terminal period, neither action carries a continuation value, defending in the rain gains the flow payoff b, attacking in the shine gains the flow payoff a, and thus both types play sincerely and no terminal belief is informative. The belief 0 is never funded, because the belief never leaves 0 and the donor's expected flow payoff there lies below c under every play, by Assumption 1. At the belief 1, every posterior equals 1, and thus both types play sincerely, as in the terminal period.

Suppose some belief is informative, and fix an informative belief μ_ν in the latest period t with one. At every funded interior belief in every period after t, both types attack with a common probability in the shine and defend with a common probability in the rain, the common probabilities cancel in Bayes' rule when the action and flow payoff pair has positive probability, the restriction above applies when it does not, and every observation moves the belief by the weather alone, to $\mu^+(\mu)$ after a shine and to $\mu^-(\mu)$ after the rain. Then, at every such belief, either action yields the same posterior in each realization of the weather, and both types play sincerely, as in the terminal period. The recipient therefore plays sincerely at every funded belief after period t.

Fix a type θ and a belief $\tilde{\mu}$ in period $t+1$. Under sincere play from $\tilde{\mu}$, funding continues until the belief first falls below the threshold of its period, the event that funding survives k further periods depends only on the first k realizations of the weather, and the flow payoff of a funded period depends only on the weather of that period, which is independent of the earlier realizations conditional on the type. Then,

$$V_\theta(\tilde{\mu}) = (s_\theta a + (1 - s_\theta)b)\, D_\theta(\tilde{\mu}), \tag{A.11}$$

where $D_\theta(\tilde{\mu})$ denotes the expected discounted number of funded periods from $\tilde{\mu}$ under the weather distribution of type θ. Each finite weather sequence moves the belief through compositions of μ^+ and μ^- applied to $\tilde{\mu}$, and funding survives the sequence if and only if each composition value is at least the threshold of its period. Each composition is strictly increasing, and thus D_θ is weakly increasing. At every funded belief, D_θ is at least one, and at every unfunded belief, D_θ is zero. The value $D_\theta(\tilde{\mu})$ is a sum with one term for each surviving weather sequence, and the term for a sequence of length k is the probability that type θ generates the sequence, discounted by δ^k. Then, every value of D_θ, and every difference of two values, lies in a finite set determined by δ, s_θ, and $T - t$ alone.

Suppose the action distributions at μ_ν differ in the shine. If some type mixes in the shine, that type is indifferent, and by equation (A.11),

$$a = \delta\,(s_\theta a + (1 - s_\theta)b)\left[D_\theta(\mu^{0,1}(\mu_\nu)) - D_\theta(\mu^{1,1}(\mu_\nu))\right],$$

which is at most zero because $\mu^{1,1}(\mu_\nu) \geq \mu^{0,1}(\mu_\nu)$ by part (i) of Definition 1 and D_θ is weakly increasing. If instead one type attacks and the other defends, part (i) of Definition 1 requires the attacker to be the strong type, the defense reveals the weak type, and the weak type gains $a + \delta V_0(\mu^{1,1}(\mu_\nu)) > 0$ over $\delta V_0(0) = 0$ by deviating to an attack. Then, the action distributions at μ_ν differ in the rain, and, by part (i) of Definition 1 and Bayes' rule, the strong type attacks in the rain with a strictly higher probability than the weak type.

Suppose the strong type attacks in the rain with probability one. The defense then reveals the weak type, and $\mu^{0,0}(\mu_\nu) = 0$. If the weak type mixes, indifference and equation (A.11) give $b = \delta\,(s_0 a + (1 - s_0)b)\, D_0(\mu^{1,0}(\mu_\nu))$, and the right-hand side is zero if $\mu^{1,0}(\mu_\nu)$ is unfunded and at least $\delta\,(s_0 a + (1 - s_0)b)$ if funded, and each is impossible by Assumption 1. If instead the weak type defends with probability one, the attack reveals the strong type, and the weak type gains $\delta V_0(1) \geq \delta\,(s_0 a + (1 - s_0)b) > b$ over $b + \delta V_0(0) = b$ by deviating to an attack, by Assumption 1.

Then, the strong type mixes in the rain, and indifference and equation (A.11) give

$$b = \delta\,(s_1 a + (1 - s_1)b)\left[D_1(\mu^{1,0}(\mu_\nu)) - D_1(\mu^{0,0}(\mu_\nu))\right]. \tag{A.12}$$

If $\mu^{1,0}(\mu_\nu)$ is unfunded, the right-hand side is at most zero. If $\mu^{0,0}(\mu_\nu)$ is unfunded and $\mu^{1,0}(\mu_\nu)$ is funded, the right-hand side is at least $\delta\,(s_1 a + (1 - s_1)b) > \delta\,(s_0 a + (1 - s_0)b) > b$ by Assumption 1. Then, both posteriors are funded. The bracketed term of equation (A.12) lies in a finite set that does not depend on b, and $b/\left[\delta\,(s_1 a + (1 - s_1)b)\right]$ is strictly increasing in b, and thus each member of the finite set satisfies equation (A.12) at most at one value of b. Then, at all but finitely many values of b, no informative belief exists.

At such b, both types play sincerely at every funded belief in every period. Now fix a belief the donor's strategy leaves unfunded in some period, and consider the history at which the donor deviates to fund it. By Assumption 3, the donor's posteriors there follow Bayes' rule given the recipient's play, with commonly unused action-payoff pairs updated by the weather alone. If the two types induce the same action distribution in every weather realization, both actions induce the same weather posterior in each realization, continuation play depends only on the belief, and the argument of part (i) makes sincere play strictly optimal for both types. If instead the distributions differ in some realization, the arguments above apply unchanged at the deviation history, because the continuation values satisfy equation (A.11) at every belief and the bracketed term of equation (A.12) draws from the same finite set, and thus, at all but finitely many b, no such belief exists. Then, at all but finitely many b, both types play sincerely at every belief at which the donor funds, on or off the path of play, the donor's surplus from funding is equation (A.3) at every belief in every period, and their funding threshold in every period is $\bar{\mu}_t$.

Let μ_t^* denote the infimum of the set of funded beliefs at which both types defend in the rain. In every monotone equilibrium, both types defend in the rain at every funded belief, and thus $\mu_t^* = \bar{\mu}_t$. Then, the infimum of Definition 2 is attained at every monotone equilibrium, and every monotone equilibrium is most-sincere, which concludes part (ii). $\square$

E Proofs

This section provides proofs for all statements presented in the main text. We restate the results for convenience.

E.1 Theorem 1

Theorem 1. *Fix $T \leq \infty$. There exists a $\lambda^* \in (0,1)$ such that, for all $\lambda < \lambda^*$,*

(i) a monotone equilibrium exists;

(ii) in every monotone equilibrium, there exists a threshold $\bar{\mu}_t \in (0,1)$ such that the donor funds in period t if and only if $\mu \geq \bar{\mu}_t$;

(iii) in every monotone equilibrium, there exists a threshold $\mu_t^* \in [(\mu^-)^{-1}(\bar{\mu}_{t+1}), 1]$ for all $t < T$, with $(\mu^-)^{-1}(\bar{\mu}_{t+1}) > \bar{\mu}_t$, such that both types of recipient play sincerely for any $\mu > \mu_t^*$ and pool on attack for any $\mu \in [\bar{\mu}_t, \mu_t^*)$ in each period t.

Proof. We prove the theorem first for the infinite horizon and then for all finite horizon. In each case, we first characterize play in any monotone equilibrium, which yields parts (ii) and (iii), and then we construct the equilibrium for part (i).

$T = \infty$ characterization. Suppose first that $T = \infty$. We first derive bounds on the recipient's continuation value in any stationary monotone equilibrium. Fix a funded belief μ. When the flow payoff is not observed, the current belief is a weighted average of the posteriors after an attack and a defense. Definition 1 part (i) orders these posteriors, and thus $\mu^{1,\emptyset}(\mu) \geq \mu$. This also covers an attack assigned probability zero, because Bayes' rule then gives $\mu^{0,\emptyset}(\mu) = \mu$. Recall from equation (A.1) that $\mu^+(\mu)$ is the donor's posterior after learning that the weather is the shine, before using the recipient's action to update further. If an attack occurs with positive probability conditional on the shine, $\mu^+(\mu)$ is a weighted average of $\mu^{1,1}(\mu)$ and $\mu^{0,1}(\mu)$. Definition 1 part (i) then gives $\mu^{1,1}(\mu) \geq \mu^+(\mu)$. If an attack occurs with probability zero conditional on the shine, Bayes' rule instead gives $\mu^{0,1}(\mu) = \mu^+(\mu)$, and Definition 1 part (i) again gives $\mu^{1,1}(\mu) \geq \mu^+(\mu)$. Because $s_1 > s_0$, equation (A.1) also gives $\mu^+(\mu) \geq \mu$. Thus, when the current belief is funded, the posterior after an attack is funded both when the flow payoff is not observed and when it reveals the shine.

Let $V_\theta(\mu)$ denote the stationary continuation value of type θ at a funded belief. The recipient receives at most $s_\theta a + (1 - s_\theta)b$ in each period, and thus

$$\frac{s_\theta a}{1 - \delta(1 - \lambda(1 - s_\theta))} \leq V_\theta(\mu) \leq \frac{s_\theta a + (1 - s_\theta)b}{1 - \delta}. \tag{A.13}$$

To obtain the lower bound, suppose the recipient always attacks. In each period, the recipient then receives the expected flow payoff $s_\theta a$. Funding continues after an unobserved flow payoff and after a flow payoff revealing the shine, which occurs with probability $1-\lambda(1-s_\theta)$. Summing these discounted payoffs gives the lower bound.

At $\lambda=0$, Assumption 2 gives $\delta(1-s_\theta)b/(1-\delta)<b$ for both types. Inequality (3) gives $\delta s_0 a>[1-\delta+\delta s_0]b>(1-\delta)b$. Because $s_1>s_0$, we also have $\delta s_1 a>(1-\delta)b$. At $\lambda=0$, the left-hand side of inequality (A.16) equals $a-\delta(1-s_\theta)b/(1-\delta)$, which is greater than $a-b>0$ by Assumption 2. Each inequality below is therefore strict at $\lambda=0$ and continuous in λ. There is a strictly positive bound on λ below which, for both types,

$$\delta\left[\frac{s_\theta a+(1-s_\theta)b}{1-\delta}-\frac{s_\theta a}{1-\delta(1-\lambda(1-s_\theta))}\right]<b, \tag{A.14}$$

$$\delta(1-\lambda)\frac{s_\theta a}{1-\delta(1-\lambda(1-s_\theta))}>b, \tag{A.15}$$

$$a+\delta(1-\lambda)\frac{s_\theta a}{1-\delta(1-\lambda(1-s_\theta))}-\delta\frac{s_\theta a+(1-s_\theta)b}{1-\delta}>0. \tag{A.16}$$

For the remainder of the infinite-horizon argument, restrict λ below this bound.

Consider any stationary monotone equilibrium. At $\mu=0$, beliefs are fixed at zero forever. The recipient therefore plays sincerely whenever the donor funds. Equation (A.2) gives $s_0 a+(1-s_0)b<c$, and thus funding is unprofitable at $\mu=0$. The stationary funding threshold is strictly positive.

Next, take any belief $\mu>\mu^{att}$ and suppose the weather is the shine. If the posterior after an attack when the flow payoff is not observed is funded, the lower and upper bounds in equation (A.13) imply that an attack exceeds a defense by at least inequality (A.16). If the posterior after an attack when the flow payoff is not observed is unfunded but the posterior after an attack whose flow payoff reveals the shine is funded, Definition 1 part (i) implies that the posterior after a defense when the flow payoff is not observed is also unfunded. The payoff from an attack then exceeds the payoff from a defense by more than $a-\lambda b>0$, by inequality (A.14). If neither posterior after an attack is funded, Definition 1 part (i) implies that neither posterior after a defense is funded, and an attack exceeds a defense by the flow payoff a. Then, both types attack in the shine at every belief above μ^{att}. Funding at such a belief gives the donor an expected flow payoff of at least $s(\mu)a>c$. The donor can end support in the next period, and thus funding is strictly profitable. It follows that the stationary funding threshold satisfies $\bar{\mu}\in(0,\mu^{att}]$.

Now consider recipient play at a funded belief. Fix a funded belief and suppose the weather is the shine. Both posteriors after an attack are funded, as shown above. The bounds in equation (A.13) and inequality (A.14) imply that an attack exceeds a defense by more than $a-b>0$. Then, both types attack in the shine.

A defense therefore reveals the rain whenever it occurs with positive probability. Once the defense reveals the rain, observing the flow payoff adds no information, and the posterior after the defense is the same whether or not the flow payoff is observed. If this posterior is funded, the bounds in equation (A.13) and inequality (A.14) make the defense strictly optimal for both types. If it is unfunded, an attack leads to a funded posterior when the flow payoff is not observed and inequality (A.15) makes the attack strictly optimal for both types. Then, whenever a defense occurs with positive probability, both types defend with probability one. Otherwise, both types attack with probability one. Thus, both types choose the same pure action in the rain at every funded belief.

At $\mu = 1$, the posterior remains at one forever on. Because $\bar{\mu} \leq \mu^{att} < 1$, the donor funds at this belief. The bounds in equation (A.13) and inequality (A.14) then make a defense strictly optimal in the rain for both types. Definition 1 part (ii) therefore implies that the funded beliefs at which both types defend in the rain form an upper interval. Let μ^* be the lower endpoint of this interval.

Now fix a funded belief $\mu < (\mu^-)^{-1}(\bar{\mu})$. Suppose both types defend in the rain. A defense then reveals the rain, and Bayes' rule gives the posterior $\mu^-(\mu) < \bar{\mu}$. Thus, funding ends and the defense gives b. Because both types attack in the shine, an attack observed without the flow payoff can only have followed the shine. Bayes' rule then gives the posterior $\mu^+(\mu)$. We have $\mu^+(\mu) > \mu \geq \bar{\mu}$, and thus this posterior is funded. Inequality (A.15) therefore makes the attack strictly optimal. Then, both types attack in the rain at every such belief, and $\mu^* \geq (\mu^-)^{-1}(\bar{\mu}) > \bar{\mu}$. This proves parts (ii) and (iii) for every stationary monotone equilibrium.

$\boldsymbol{T = \infty}$ **construction.** We now construct a stationary monotone equilibrium. Fix a candidate funding threshold $\mu' \in (0, \mu^{att}]$. Suppose both types attack in the shine at every belief and defend in the rain if and only if $\mu \geq (\mu^-)^{-1}(\mu')$. Let $V_D^{\mu',\lambda}(\mu)$ denote the donor's best-response continuation value against this recipient strategy. The donor can end support and receive $c/(1-\delta)$. At a belief where the recipient plays sincerely, the donor's value satisfies

$$V_D^{\mu',\lambda}(\mu) = \max\left\{\frac{c}{1-\delta}, u^{sin}(\mu) + \delta\left[s(\mu)V_D^{\mu',\lambda}(\mu^+(\mu)) + (1-s(\mu))V_D^{\mu',\lambda}(\mu^-(\mu))\right]\right\}. \quad \text{(A.17)}$$

At a belief where both types pool on attack, an unobserved flow payoff leaves the belief unchanged. If the donor funds at the belief, substituting the continuation value at the unchanged belief and rearranging gives the second expression below. If that expression is below $c/(1-\delta)$, substituting $c/(1-\delta)$ after the unobserved flow payoff makes funding strictly worse than ending support. Then, in either case, the donor's value satisfies

$$V_D^{\mu',\lambda}(\mu) = \max\left\{\frac{c}{1-\delta}, \frac{s(\mu)a + \delta\lambda\left[s(\mu)V_D^{\mu',\lambda}(\mu^+(\mu)) + (1-s(\mu))V_D^{\mu',\lambda}(\mu^-(\mu))\right]}{1-\delta(1-\lambda)}\right\}. \quad \text{(A.18)}$$

Because the coefficient in equation (A.18) is $\delta\lambda/[1-\delta(1-\lambda)] \leq \delta$, repeated substitution from any bounded function therefore converges uniformly to a unique bounded solution of equations (A.17) and (A.18).

We next show that this solution is weakly increasing in the belief. Begin repeated substitution from the constant function $a/(1-\delta)$. Suppose the function at some step is W, and that W is weakly increasing and weakly above $s(\mu)a/(1-\delta)$ at every belief. The expected value of s at the two Bayesian posteriors equals its current value, and thus the posterior average of the function is weakly above $s(\mu)a/(1-\delta)$. Under sincere play, the funding expression is therefore weakly above $u^{sin}(\mu)+\delta s(\mu)a/(1-\delta) \geq s(\mu)a/(1-\delta)$. When pooling on attack, substituting the same lower bound into equation (A.18) gives $s(\mu)a/(1-\delta)$. Then, substitution into either equation preserves the lower bound.

Because μ^+, μ^-, and s are increasing, the distribution of the posterior rises in first-order stochastic dominance as μ rises, and thus the posterior average of any weakly increasing function is weakly increasing in the current belief. Then, the funding expression in each of equations (A.17) and (A.18) is weakly increasing on the region where it applies.

At the recipient's threshold $\mu = (\mu^-)^{-1}(\mu')$, the difference between the sincere and the pooling funding expressions equals

$$(1-s(\mu))b \tag{A.19}$$

$$+\frac{\delta(1-\lambda)}{1-\delta(1-\lambda)}\left[(1-\delta)\left[s(\mu)W(\mu^+(\mu))+(1-s(\mu))W(\mu^-(\mu))\right]-s(\mu)a\right], \tag{A.20}$$

where W is the continuation-value function inserted at that step of repeated substitution. The expression in square brackets is non-negative by the lower bound just established. The funding expression therefore does not fall when the recipient changes from pooling on attack to sincere play. Every function produced by repeated substitution is weakly increasing, and so is the uniform limit $V_D^{\mu',\lambda}$.

We next establish continuity in μ' and λ. Because the recipient's threshold moves with μ', to hold this threshold fixed while comparing candidate thresholds, map each original belief μ to the belief $\tilde{\mu}$ satisfying

$$\frac{\tilde{\mu}}{1-\tilde{\mu}} = \frac{\mu^{att}}{1-\mu^{att}}\frac{\mu}{1-\mu}\frac{1-\mu'}{\mu'}. \tag{A.21}$$

Equation (A.21) maps the candidate μ' to μ^{att} and the recipient's threshold $(\mu^-)^{-1}(\mu')$ to $(\mu^-)^{-1}(\mu^{att})$. Bayes' rule multiplies the odds of a belief by s_1/s_0 after the shine and by $(1-s_1)/(1-s_0)$ after the rain. Then, the change of beliefs preserves the posterior maps. Solving equation (A.21) for the original belief gives

$$\mu = \frac{\mu'(1-\mu^{att})\tilde{\mu}}{\mu'(1-\mu^{att})\tilde{\mu}+(1-\mu')\mu^{att}(1-\tilde{\mu})}.$$

On every closed interval of candidate thresholds contained in $(0,1)$, the right-hand side is uniformly continuous in μ' and $\tilde{\mu}$. After beliefs are indexed by $\tilde{\mu}$, the recipient's threshold and the posterior arguments in equations (A.17) and (A.18) no longer depend on μ'. The flow probability $s(\mu)$ changes uniformly with μ'.

Take two pairs of candidate thresholds and monitoring levels in this closed interval. Inequality (1) gives $c < a$, and every flow payoff from funding is at most a. Then, the fixed continuation values are bounded above by $a/(1-\delta)$ at both pairs. Let ε be the maximum difference between the values produced by the two donor equations when the same continuation function bounded by $a/(1-\delta)$ is inserted at both pairs, after beliefs are indexed by $\tilde{\mu}$. As the two pairs converge, $\varepsilon \to 0$. If d denotes the maximum difference between their fixed continuation values, then the two donor equations imply $d \leq \delta d + \varepsilon$, and hence $d \leq \varepsilon/(1-\delta)$. Then, the fixed continuation values converge uniformly, and the funding expression in equation (A.18) evaluated at the candidate threshold is continuous in μ' and λ.

At $\lambda = 0$, this funding expression equals $s(\mu')a/(1-\delta)$. It equals $c/(1-\delta)$ at $\mu' = \mu^{att}$ and is strictly below $c/(1-\delta)$ at every $\mu' < \mu^{att}$. Now take $\lambda > 0$ and the candidate $\mu' = \mu^{att}$. At the favorable posterior $\mu^+(\mu^{att})$, funding for one period and ending support afterward is strictly profitable because either pooling or sincere play gives an expected flow payoff strictly above c. Then, $V_D^{\mu^{att},\lambda}(\mu^+(\mu^{att})) > c/(1-\delta)$, while the value at every other belief is weakly above $c/(1-\delta)$. Substitution into equation (A.18) shows that funding is strictly profitable at the candidate μ^{att} whenever $\lambda > 0$.

Fix any $\mu'' \in (0, \mu^{att})$ and any $\lambda < (1-\delta)\,(c - s(\mu'')a)\,/\,(\delta s(\mu'')a)$, a strictly positive bound folded into λ^*. At any candidate $\mu' \leq \mu''$, if the supremum of $V_D^{\mu',\lambda}$ over beliefs weakly below μ' exceeded $c/(1-\delta)$, the funding branch of equation (A.18) would bound it by $s(\mu'')a(1-\delta+\delta\lambda)/(1-\delta)^2 < c/(1-\delta)$, so $V_D^{\mu',\lambda} = c/(1-\delta)$ at every belief weakly below μ', and the funding expression at μ' then lies strictly below $c/(1-\delta)$.

Below the recipient's threshold, the funding expression in equation (A.18) is strictly increasing in the belief. The current flow payoff is strictly increasing, and the posterior average of the weakly increasing value $V_D^{\bar{\mu},\lambda}$ is weakly increasing. Because the donor is indifferent at $\bar{\mu}$, funding is strictly unprofitable below $\bar{\mu}$ and strictly profitable between $\bar{\mu}$ and $(\mu^-)^{-1}(\bar{\mu})$. At the recipient's threshold, equation (A.19) shows that the sincere funding expression is weakly above the pooling funding expression. The sincere expression is weakly increasing above the threshold. Then, funding if and only if $\mu \geq \bar{\mu}$ is a best response for the donor.

Reduce the positive bound on λ, if necessary, so that this donor threshold exists and inequalities (A.14)–(A.16) hold. For the infinite-horizon case, let λ^* denote this bound and fix $\lambda < \lambda^*$. Fix the donor threshold $\bar{\mu}$ obtained above and let $\mu^* = (\mu^-)^{-1}(\bar{\mu})$. The donor funds if and only if $\mu \geq \bar{\mu}$. At funded beliefs below μ^*, both types attack in both realizations of the weather. At funded beliefs weakly above μ^*, both types play sincerely. At $\mu = 0$, both types play sincerely and the posterior remains zero after every observation.

At funded beliefs, beliefs follow Bayes' rule after every positive-probability observation. Below μ^*, the donor assigns posterior zero after an off-path defense. Weakly above μ^*, both types play sincerely. Thus, when the flow payoff is not observed, an attack reveals the shine and gives posterior $\mu^+(\mu)$, while a defense reveals the rain and gives posterior $\mu^-(\mu)$. When the flow payoff is observed after an off-path action at funded beliefs weakly above μ^*, the donor assigns the posterior induced by the weather revealed by the action and flow payoff pair. These beliefs satisfy Definition 1 part (i).

Fix a funded belief below μ^*. In the rain, a defense gives b and leads to posterior zero. An attack preserves funding whenever the flow payoff is not observed and gives more than b by inequality (A.15). In the shine, an attack gives a, while a defense gives zero and leads to posterior zero. Then, both types attack in both realizations of the weather.

Now fix a funded belief weakly above μ^*. The definition of μ^* gives $\mu^-(\mu) \geq \bar{\mu}$, and therefore both weather posteriors are funded. In the rain, the bounds in equation (A.13) and inequality (A.14) make the defense strictly optimal. In the shine, equation (A.13), inequality (A.14), and $a > b$ make the attack strictly optimal. Then, both types play sincerely.

It remains to specify play after the donor deviates to fund at an interior belief $\mu < \bar{\mu}$. Both types attack in both realizations of the weather. After an attack whose flow payoff is not observed or reveals the shine, the donor assigns posterior $\bar{\mu}$. After an attack whose flow payoff reveals the rain, or after any defense, the donor assigns posterior zero. These beliefs satisfy Definition 1 part (i). In the rain, an attack reaches the funded belief $\bar{\mu}$ whenever the flow payoff is not observed and gives more than b by inequality (A.15), while a defense gives b and ends funding. In the shine, an attack gives a, while a defense gives zero and ends funding. Then, both types attack.

At funded beliefs, the donor's strategy is optimal by the best-response argument above. The donor is indifferent at $\bar{\mu}$, so the continuation value there is $c/(1-\delta)$. The continuation value at posterior zero is also $c/(1-\delta)$ because the donor ends support. If the donor deviates to fund at an interior belief $\mu < \bar{\mu}$, both types attack and every continuation posterior gives this same value. The difference between funding and ending support is therefore $s(\mu)a - c < 0$. At $\mu = 0$, sincere play gives the donor $s_0 a + (1 - s_0)b < c$ and leaves the belief unchanged, and thus funding is also unprofitable.

The donor funds on $[\bar{\mu}, 1]$. Across funded beliefs, the probability that either type defends in the rain is zero below μ^* and one weakly above μ^*. The profile satisfies Definition 1 parts (i)–(iii). Flow payoffs are bounded and discounted, and the one-shot deviation checks above therefore establish optimality for the recipient. Then, the profile forms a stationary equilibrium, which concludes the proof of part (i) when $T = \infty$.

***$T < \infty$* characterization.** Now let $T < \infty$ and $j = T - t + 1$ denote the number of periods remaining. Recall that the recipient's value at any belief satisfies $V_\theta^{[j]}(\mu) \leq \bar{v}_\theta^{[j]}$ and, under the conditions of Lemma 8, $V_\theta^{[j]}(\mu) \geq \underline{v}_\theta^{[j]}$, while Lemma 7 gives $\delta(\bar{v}_\theta^{[j]} - \underline{v}_\theta^{[j]}) < b$ at any j and every λ below its bound. Now let λ^* be the smallest of the bound of Lemma 7, the bound derived at the end of the proof of part (i), the bounds of Proposition 5 and Lemmas

9, 5, 10, and 11, and the two bounds

$$1 - \frac{b}{\delta(s_0 a + (1-s_0)b)} \qquad \text{and} \qquad \frac{a(1-\delta)}{\delta(s_1 a + (1-s_1)b)}.$$

Each is strictly positive, and thus $\lambda^* \in (0,1)$. Note that λ^* lies weakly below every bound invoked in the proof of Proposition 6.

Fix $\lambda < \lambda^*$. We first establish parts (ii) and (iii) for any monotone equilibrium. Fix any monotone equilibrium. By Proposition 6, its funding thresholds are weakly decreasing over time, $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$ for every $j \geq 2$. As in the proof of Proposition 6, at every $j \geq 2$ and every funded belief $\mu \geq \bar{\mu}^{[j]}$, we have $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$ and $\mu^{1,1}(\mu) \geq \bar{\mu}^{[j-1]}$. These inequalities verify the posterior conditions of Lemmas 8, 9, 10, and 11. Lemma 11 then gives $\bar{\mu}^{[j]} \in (0, \mu^{att}]$ for every j, so the remaining condition $\bar{\mu}^{[j-1]} > 0$ in Lemma 10 also holds.

Part (ii) now follows. The donor funds in period t if and only if $\mu \geq \bar{\mu}_t$, by Definition 1, and by Lemma 11 and $\mu^{att} < 1$, we have $\bar{\mu}_t \in (0,1)$ in every period.

For part (iii), fix $t < T$, so that $j \geq 2$. Both types play the same pure action in the rain at every funded belief by Lemma 10, and by part (ii) of Definition 1, if both types defend in the rain at a funded belief, they defend at every larger belief. At $\mu = 1$ every posterior is 1 and the belief is always funded, so defending in the rain yields b and the same continuation as attacking, and both types defend there. The funded beliefs at which both types defend in the rain therefore form an interval containing 1, with lower endpoint $\mu^{*[j]}$. At a funded belief $\mu < (\mu^-)^{-1}(\bar{\mu}^{[j-1]})$ at which both types defend in the rain, a defense reveals the rain by Lemma 9, moves the belief to $\mu^-(\mu) < \bar{\mu}^{[j-1]}$, ends funding, and yields b, whereas attacking preserves the funded posterior $\mu^{1,\emptyset}(\mu) \geq \bar{\mu}^{[j-1]}$ when only the action is observed, where one sincere period yields at least $s_0 a + (1-s_0)b$, and thus attacking is worth at least $\delta(1-\lambda)(s_0 a + (1-s_0)b) > b$, a contradiction. Then, $\mu^{*[j]} \geq (\mu^-)^{-1}(\bar{\mu}^{[j-1]})$, both types pool on attack at every $\mu \in [\bar{\mu}^{[j]}, \mu^{*[j]})$ by Lemma 9, and both types play sincerely at every $\mu > \mu^{*[j]}$. It remains to show $(\mu^-)^{-1}(\bar{\mu}^{[j-1]}) > \bar{\mu}^{[j]}$. At $j = 1$ there is no continuation, both types play sincerely at every belief, the donor's flow payoff from funding is $u^{sin}(\mu) - c$, and thus $\bar{\mu}^{[1]} = \mu^{sin}$. By Proposition 6, we then have $\bar{\mu}^{[j-1]} \geq \mu^{sin}$, so Lemmas 11 and 2 give $\mu^-(\bar{\mu}^{[j]}) \leq \mu^-(\mu^{att}) < \mu^{sin} \leq \bar{\mu}^{[j-1]}$, and applying the increasing function $(\mu^-)^{-1}$ yields $\bar{\mu}^{[j]} < (\mu^-)^{-1}(\bar{\mu}^{[j-1]})$. Then, the attack region $[\bar{\mu}^{[j]}, \mu^{*[j]})$ is nonempty. In the terminal period itself, both types play sincerely at every funded belief, as shown above. Then, the attack region is empty and there is no threshold μ^*_T, which is why part (iii) restricts to $t < T$. This establishes parts (ii) and (iii).

***$T < \infty$* construction.** In the remainder of the proof, the symbols $\Psi^{[j]}$, $\Phi^{[j]}$, $\bar{\mu}^{[j]}$, and $\mu^{*[j]}$ denote the objects of equations (A.3) and (A.4), with $\mu^{*[1]} = 0$ and $\mu^{*[j]} = (\mu^-)^{-1}(\bar{\mu}^{[j-1]})$ for $j \geq 2$, and no longer the thresholds of the monotone equilibrium fixed above. Consider the following strategy profile. The donor funds at belief μ with j periods remaining if and only if $\Psi^{[j]}(\mu) \geq 0$. At $j = 1$, both types of recipient play sincerely at every belief. At $j \geq 2$, both types attack in the shine at every belief, and in the rain, both types attack at beliefs $\mu < \mu^{*[j]}$ and defend at beliefs $\mu \geq \mu^{*[j]}$. Beliefs follow Bayes' rule on the path of play. A defense leads to the posterior $\mu^-(\mu)$ in both realizations of the flow payoff, as does an attack whose flow payoff reveals the rain. At a belief the donor funds only by deviating, the posterior after an attack is as follows. For $\mu \in [\bar{\mu}^{[j-1]}, \bar{\mu}^{[j]})$, it is given by Bayes' rule when both types attack regardless of the weather, and thus the posterior is μ when only the action is observed and $\mu^+(\mu)$ when the flow payoff reveals a shine. For $\mu < \bar{\mu}^{[j-1]}$, the posterior is $\bar{\mu}^{[j-1]}$ both when only the action is observed and when the flow payoff reveals a shine, and it is $\mu^-(\mu)$ when the flow payoff reveals the rain. These specifications agree with Bayes' rule wherever the observation is on the path of play, and in every case the posterior after an attack weakly exceeds the posterior after a defense.

We first show the donor's funded set is the interval $[\bar{\mu}^{[j]}, 1]$. Against the strategy profile's recipient play, the donor's surplus from funding and continuing optimally is $\Psi^{[j]}$. To see this, note that equation (A.3) gives the flow payoff and continuation values when both types play sincerely, equation (A.4) gives them when both types pool on attack, and play at $j = 1$ is sincere by construction. The proof of Lemma 5 shows equation (A.4) is weakly increasing below $\mu^{*[j]}$, and by inequality (A.7) equation (A.3) is strictly positive at and above $\mu^{*[j]}$. Then, the set of beliefs at which $\Psi^{[j]}(\mu) \geq 0$ is an interval with infimum $\bar{\mu}^{[j]}$ and right endpoint 1. We show by induction that $\Psi^{[j]}$ is right-continuous in μ. The terminal case is continuous, right-continuity passes through the maximum with zero and through composition with the continuous increasing maps $\mu^{\pm}$, and every belief has a right-neighborhood on which $\Psi^{[j]}$ coincides with one of equations (A.3) and (A.4), which completes the induction. Then, $\Psi^{[j]}(\bar{\mu}^{[j]})$ equals the limit of $\Psi^{[j]}(\mu)$ as μ decreases to $\bar{\mu}^{[j]}$, which is at least zero, and the funded set is $[\bar{\mu}^{[j]}, 1]$.

The thresholds are ordered by Proposition 5, $\bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$ at every $j \geq 2$, with $\bar{\mu}^{[1]} = \mu^{sin} > 0$. By inequality (A.6),

$$\bar{\mu}^{[j]} \leq \mu^{att} - \frac{(1-\delta)(\delta(1-\lambda))^{T-1}(1-c/a)b}{(s_1 - s_0)a} < \mu^{att},$$

and thus every threshold is interior. Moreover, inequality (A.7) gives $\mu^{*[j]} > \mu^{att} > \bar{\mu}^{[j]}$, and thus the attack region $[\bar{\mu}^{[j]}, \mu^{*[j]})$ is nonempty at every $j \geq 2$.

Next, we establish that the bound of Lemma 8 applies. The donor funds on the interval $[\bar{\mu}^{[j]}, 1]$, and after an attack at a funded belief, the posterior is μ itself when only the action is observed and both types pool on attack, is $\mu^+(\mu)$ when only the action is observed and both types play sincerely, and is $\mu^+(\mu)$ when the flow payoff reveals a shine. Each of these is at least $\bar{\mu}^{[j-1]}$. The first is because $\mu \geq \bar{\mu}^{[j]} \geq \bar{\mu}^{[j-1]}$ by Proposition 5, and the others are because $\mu^+(\mu) > \mu$. The conditions of Lemma 8 therefore hold, and the recipient's value at every funded belief is at least the bound of Lemma 8.

We can now verify the recipient's play at funded beliefs, beginning with the shine. Fix a funded belief μ and $j \geq 2$. Attacking yields a flow payoff of a, and the posterior remains funded regardless of whether the flow payoff is observed, whereas defending yields a flow payoff of zero. Then, by Lemmas 8 and 7, the payoff from attacking exceeds the payoff from defending by at least $a - \delta(\bar{v}_\theta^{[j-1]} - \underline{v}_\theta^{[j-1]}) \geq a - b > 0$. Both types attack in the shine.

In the rain, take first $\mu \in [\bar{\mu}^{[j]}, \mu^{*[j]})$. Defending lowers the belief to $\mu^-(\mu) < \mu^-(\mu^{*[j]}) = \bar{\mu}^{[j-1]}$, funding ends, and the payoff is b. Attacking yields a flow payoff of zero, and with probability $1 - \lambda$ the donor observes only the action, the posterior stays at $\mu \geq \bar{\mu}^{[j-1]}$, and funding continues. Playing sincerely next period yields an expected flow payoff of at least $s_\theta a + (1 - s_\theta)b \geq s_0 a + (1 - s_0)b$. Then, attacking is worth at least $\delta(1 - \lambda)(s_0 a + (1 - s_0)b)$, which exceeds b because $\lambda < 1 - b/[\delta(s_0 a + (1 - s_0)b)]$. Both types attack. Now take $\mu \geq \mu^{*[j]}$. Defending moves both posteriors to $\mu^-(\mu) \geq \bar{\mu}^{[j-1]}$, funding continues, and defending is worth the flow payoff b together with the discounted continuation, whereas attacking yields a flow payoff of zero and a continuation of at most $\bar{v}_\theta^{[j-1]}$. Then, by Lemmas 8 and 7, the payoff from defending exceeds the payoff from attacking by at least $b - \delta(\bar{v}_\theta^{[j-1]} - \underline{v}_\theta^{[j-1]}) > 0$. Both types defend.

The recipient moves at an unfunded belief only if the donor deviates to fund there. We now verify that the strategy profile's play at such beliefs is a best response to the specified posteriors, and that the donor's one-shot deviation is unprofitable against it. Fix $j \geq 2$, and take first $\mu \in [\bar{\mu}^{[j-1]}, \bar{\mu}^{[j]})$, if this interval is nonempty. In the rain, defending lowers the belief to $\mu^-(\mu) < \mu^-(\mu^{*[j]}) = \bar{\mu}^{[j-1]}$, because $\mu < \bar{\mu}^{[j]} \leq \mu^{att} < \mu^{*[j]}$. Then, funding ends and defending yields b, whereas after an attack the belief stays at $\mu \geq \bar{\mu}^{[j-1]}$ when only the action is observed, funding continues, and a single period of sincere play yields at least $s_0 a + (1 - s_0)b$. Then, attacking is worth at least $\delta(1 - \lambda)(s_0 a + (1 - s_0)b) > b$. In the shine, attacking yields at least a whereas defending yields zero. Both types therefore attack, and the donor's deviation yields the flow payoff and continuation values in equation (A.4). Its value is $\Psi^{[j]}(\mu)$, which is strictly negative below $\bar{\mu}^{[j]}$ because the set of beliefs at which $\Psi^{[j]}(\mu) \geq 0$ is the interval $[\bar{\mu}^{[j]}, 1]$, as established above.

Now take $\mu < \bar{\mu}^{[j-1]}$. Against the specified posteriors, after an attack in the rain the belief is $\bar{\mu}^{[j-1]}$ when only the action is observed and funding continues. The recipient can obtain at least $s_0 a + (1 - s_0)b$ at the funded belief $\bar{\mu}^{[j-1]}$ by playing sincerely once, and thus attacking is worth at least $\delta(1 - \lambda)(s_0 a + (1 - s_0)b) > b$, whereas defending yields b and ends funding. In the shine, attacking yields at least a whereas defending yields zero. Both types therefore attack. The donor's deviation then yields the flow payoff $s(\mu)a - c$ and, at their specified posteriors, the continuation value $\Phi^{[j-1]}(\bar{\mu}^{[j-1]}) = \Psi^{[j-1]}(\bar{\mu}^{[j-1]})$ after an attack that is unobserved or reveals a shine, and zero otherwise. Its value is

$$s(\mu)a - c + \delta(1 - \lambda + \lambda s(\mu))\Psi^{[j-1]}(\bar{\mu}^{[j-1]}).$$

By inequality (A.6), we have

$$c - s(\mu)a > c - s(\bar{\mu}^{[j-1]})a = (s_1 - s_0)a(\mu^{att} - \bar{\mu}^{[j-1]}) \geq (1-\delta)(\delta(1-\lambda))^{T-1}(1-c/a)b,$$

and thus the deviation is strictly unprofitable whenever $\delta\Psi^{[j-1]}(\bar{\mu}^{[j-1]})$ is at most $(1-\delta)(\delta(1-\lambda))^{T-1}(1-c/a)b$. To bound $\Psi^{[j-1]}(\bar{\mu}^{[j-1]})$, we show by induction on j that, for every pair of beliefs $\mu' < \mu \leq \mu^{att}$,

$$\Psi^{[j]}(\mu) - \Psi^{[j]}(\mu') \leq \left[s(\mu) - s(\mu')\right] \frac{a}{(1-\delta)^2} + \frac{\delta\lambda(a-c)}{(1-\delta)^2}.$$

At $j = 1$, we have $\Psi^{[1]}(\mu) - \Psi^{[1]}(\mu') = [s(\mu) - s(\mu')](a-b)$, and $a - b < a/(1-\delta)^2$ gives the claim. Now fix $j \geq 2$ and suppose the claim holds at $j-1$. Both beliefs lie below $\mu^{att} < \mu^{*[j]}$, and thus equation (A.4) applies to each, and differencing,

$$\Psi^{[j]}(\mu) - \Psi^{[j]}(\mu') = [s(\mu) - s(\mu')]a + \delta(1-\lambda)\left[\Phi^{[j-1]}(\mu) - \Phi^{[j-1]}(\mu')\right] + \delta\lambda\left[s(\mu)\Phi^{[j-1]}(\mu^+(\mu)) - s(\mu')\Phi^{[j-1]}(\mu^+(\mu'))\right].$$

Consider the first bracketed term. If $\Psi^{[j-1]}(\mu') \geq 0$, then $\Psi^{[j-1]}(\mu) \geq 0$ as well, because $\Psi^{[j-1]}$ is weakly increasing by the proof of Lemma 5, so $\Phi^{[j-1]} = \max\{0, \Psi^{[j-1]}\}$ agrees with $\Psi^{[j-1]}$ at both beliefs, and the difference equals $\Psi^{[j-1]}(\mu) - \Psi^{[j-1]}(\mu')$. If $\Psi^{[j-1]}(\mu) \leq 0$, the difference is zero. Otherwise, the difference is $\Psi^{[j-1]}(\mu)$, which is less than $\Psi^{[j-1]}(\mu) - \Psi^{[j-1]}(\mu')$. In every case, the first bracketed term is at most $\Psi^{[j-1]}(\mu) - \Psi^{[j-1]}(\mu')$, and the inductive hypothesis then bounds it by $[s(\mu) - s(\mu')]a/(1-\delta)^2 + \delta\lambda(a-c)/(1-\delta)^2$. For the second bracketed term, adding and subtracting $s(\mu')\Phi^{[j-1]}(\mu^+(\mu))$ and using $\Phi^{[j-1]} \in [0, (a-c)/(1-\delta)]$ with $s(\mu') \leq 1$ gives an upper bound of $[s(\mu) - s(\mu')](a-c)/(1-\delta) + (a-c)/(1-\delta)$. Combining the three bounds gives a coefficient on $s(\mu) - s(\mu')$ of at most $a/(1-\delta)^2$ and the terms free of $s(\mu) - s(\mu')$ sum to at most $\delta\lambda(a-c)/(1-\delta)^2$, which completes the induction. Now take any $\mu' < \bar{\mu}^{[j-1]}$, which exists because $\bar{\mu}^{[j-1]} > 0$, and note $\Psi^{[j-1]}(\mu') < 0$. The claim at $j-1$ with $\mu = \bar{\mu}^{[j-1]}$ then gives

$$\Psi^{[j-1]}(\bar{\mu}^{[j-1]}) < \left[s(\bar{\mu}^{[j-1]}) - s(\mu')\right] \frac{a}{(1-\delta)^2} + \frac{\delta\lambda(a-c)}{(1-\delta)^2}.$$

Because this holds for every $\mu' < \bar{\mu}^{[j-1]}$ and s is continuous, we have $\Psi^{[j-1]}(\bar{\mu}^{[j-1]}) \leq \delta\lambda(a-c)/(1-\delta)^2$. The deviation is therefore strictly unprofitable whenever

$$\frac{\delta^2\lambda(a-c)}{(1-\delta)^2} \leq (1-\delta)(\delta(1-\lambda))^{T-1}(1-c/a)b,$$

which rearranges to $\lambda/(1-\lambda)^{T-1} \leq (1-\delta)^3\delta^{T-3}b/a$. If $\lambda \leq 1/2$, then $(1-\lambda)^{T-1} \geq 2^{-(T-1)}$, and then $\lambda/(1-\lambda)^{T-1} \leq 2^{T-1}\lambda$. The required inequality therefore holds for every λ below the smaller of $1/2$ and $(1-\delta)^3\delta^{T-3}b/(2^{T-1}a)$, which is the strictly positive bound of part (i) included in the definition of λ^*. Then, $\lambda < \lambda^*$ implies the deviation is strictly unprofitable. At $j=1$, the strategy profile plays sincerely at every belief, and funding at an unfunded $\mu < \mu^{sin}$ yields $u^{sin}(\mu) - c < 0$.

In the terminal period, there is no continuation, and thus attacking in the shine yields a flow payoff of $a > 0$ whereas defending yields zero, and defending in the rain yields a flow payoff of $b > 0$ whereas attacking yields zero, so the prescribed terminal play is optimal.

It remains to check that the strategy profile is monotone. The posterior after an attack weakly exceeds the posterior after a defense at every belief, by the specification of the posteriors above. The probability of a defense in the rain is zero at funded beliefs below $\mu^{*[j]}$ and one at beliefs weakly above it, weakly increasing in μ, as verified above. The funded set is the interval $[\bar{\mu}^{[j]}, 1]$, as established above. Then, the strategy profile satisfies each part of Definition 1 and is a monotone equilibrium, and part (i) follows. □

E.2 Lemma 1

Lemma 1. *Suppose $\lambda < \lambda^*$. Then,*

(i) most-sincere equilibria exist;

(ii) every most-sincere equilibrium has the lowest donor's funding threshold for all $t \geq 2$ among all monotone equilibria, $\bar{\mu}_t = \underline{\bar{\mu}}_t := \inf_{\sigma\in\Sigma} \bar{\mu}_t(\sigma)$; and

(iii) every most-sincere equilibrium is Pareto optimal among all monotone equilibria.

Proof. We proceed in two parts, first with any finite horizon and then with the infinite horizon.

$\boldsymbol{T < \infty}$**.** Suppose first that $T < \infty$. Fix any monotone equilibrium. Objects carrying the superscript $[j]$ are those of equations (A.3) and (A.4), and objects carrying the subscript t are those of the fixed equilibrium, with $j = T - t + 1$ the number of periods remaining. We show by induction on j that the donor's surplus from funding is at most $\Psi^{[j]}(\mu)$ at every funded belief μ, and their continuation surplus is at most $\Phi^{[j]}(\mu)$ at every belief, and thus that $\bar{\mu}_t \geq \bar{\mu}^{[j]}$ and $\mu_t^* \geq \mu^{*[j]}$. At $j=1$, both types play sincerely, the donor's surplus from funding is $u^{sin}(\mu) - c = \Psi^{[1]}(\mu)$, and thus $\bar{\mu}_T \geq \mu^{sin} = \bar{\mu}^{[1]}$. Now fix $j \geq 2$ and suppose the claims hold below j. By the inductive hypothesis, the donor's continuation value with $j-1$ periods remaining is at most $\Phi^{[j-1]}$ at every belief, and, by part (iii) of Theorem 1, the recipient pools on attack below μ_t^* and plays sincerely above, with $\mu_t^* \geq (\mu^-)^{-1}(\bar{\mu}_{t+1}) \geq (\mu^-)^{-1}(\bar{\mu}^{[j-1]}) = \mu^{*[j]}$. At every $\mu < \mu^{*[j]}$, we have $\mu^-(\mu) < \bar{\mu}^{[j-1]} \leq \bar{\mu}_{t+1}$ by the inductive hypothesis, and thus a flow payoff revealing the rain ends funding and the donor's surplus is at most equation (A.4). At every $\mu \in [\mu^{*[j]}, \mu_t^*)$, the donor's surplus is at most $\Psi_{att}^{[j]}(\mu) + \delta\lambda(1-s(\mu))\Phi^{[j-1]}(\mu^-(\mu)) < \Psi_{sin}^{[j]}(\mu)$ by Lemma 4. At every $\mu \geq \mu_t^*$, the donor's surplus is at most $\Psi_{sin}^{[j]}(\mu)$. Then, every funded belief has $\Psi^{[j]}(\mu) \geq 0$, and thus

$\bar{\mu}_t \geq \bar{\mu}^{[j]}$. The donor's continuation surplus is zero at unfunded beliefs and equals the funding surplus at funded beliefs, in both cases at most $\Phi^{[j]}(\mu)$, which completes the induction.

The equilibrium constructed in the proof of Theorem 1 has funding thresholds $\bar{\mu}^{[j]}$ and thresholds $(\mu^-)^{-1}(\bar{\mu}^{[j-1]}) = \mu^{*[j]}$. Then, $\underline{\mu}_t^* = \mu^{*[j]}$, the infimum is attained, and most-sincere equilibria exist, which is part (i). Moreover, fix any most-sincere equilibrium and any $t \geq 2$, with $j = T - t + 1$. Part (iii) of Theorem 1 at period $t-1$ gives $\mu_{t-1}^* \geq (\mu^-)^{-1}(\bar{\mu}_t)$, most-sincerity gives $\mu_{t-1}^* = (\mu^-)^{-1}(\bar{\mu}^{[j]})$, and applying the increasing map μ^- yields $\bar{\mu}_t \leq \bar{\mu}^{[j]}$. Together with $\bar{\mu}_t \geq \bar{\mu}^{[j]}$ in every monotone equilibrium, this yields $\underline{\bar{\mu}}_t = \bar{\mu}^{[j]}$ at every $t \geq 2$, and hence part (ii).

Finally, fix a most-sincere equilibrium and any monotone equilibrium that yields a weakly higher payoff at the initial prior μ_1 for the donor and for the recipient. Denote by $V'_{D,t'}$ and $V'_{\theta,t'}$ the period-t' value functions of the donor and of the recipient of type θ in the latter equilibrium, and let t' denote the first period in which the two equilibria induce different play at the belief realized in that period, with $t' = T+1$ if no such period arrives. Before period t', the two equilibria induce the same actions at the same beliefs, and thus the same flow payoffs and the same distribution over realized beliefs. Then, the donor's values at the initial prior differ by the expected discounted difference at period t',

$$\Phi^{[T]}(\mu_1) - V'_{D,1}(\mu_1) = \mathrm{E}\left[\delta^{t'-1}\left(\Phi^{[T-t'+1]}(x) - V'_{D,t'}(x)\right)\right], \tag{A.22}$$

where x denotes the belief realized in period t' and the term at $t' = T+1$ equals zero. The induction above gives $\Phi^{[T-t'+1]}(x) \geq V'_{D,t'}(x)$ at every belief, the hypothesis gives $V'_{D,1}(\mu_1) \geq \Phi^{[T]}(\mu_1)$, and thus every realization of equation (A.22) satisfies $\Phi^{[T-t'+1]}(x) = V'_{D,t'}(x)$.

Now fix any realization with $t' \leq T$, so that play differs at x in period t'. The induction above gives $\bar{\mu}_{t'} \geq \bar{\mu}^{[j]}$ and $\mu_{t'}^* \geq \mu^{*[j]}$ at $j = T - t' + 1$, and thus play at x can differ in only two ways, either the most-sincere equilibrium funds at x while the other does not, or both fund and the most-sincere equilibrium plays sincerely at x while the other pools on attack. Suppose the donor funds at x in both equilibria and the recipient pools on attack at x in the other equilibrium while playing sincerely in the most-sincere equilibrium. The belief x is funded in both equilibria, and thus $x \geq \bar{\mu}^{[j]}$ and $\Phi^{[j]}(x) = \Psi^{[j]}(x)$, which equals $\Psi_{sin}^{[j]}(x)$ because $x \geq \mu^{*[j]}$. The donor's surplus at x in the other equilibrium is at most $\Psi_{att}^{[j]}(x) + \delta\lambda(1 - s(x))\Phi^{[j-1]}(\mu^-(x))$ by the induction above, which lies strictly below $\Psi_{sin}^{[j]}(x)$ by Lemma 4, and thus $V'_{D,t'}(x) < \Phi^{[j]}(x)$, which contradicts the equality above. Then, at every realization with $t' \leq T$, the donor funds at x in the most-sincere equilibrium and not in the other, and $V'_{D,t'}(x) = 0 = \Phi^{[j]}(x)$ gives $\Psi^{[j]}(x) = 0$.

The recipients' values at the initial prior likewise differ by the expected discounted difference at period t',

$$V'_{\theta,1}(\mu_1) - V_{\theta,1}(\mu_1) = \mathrm{E}_\theta\left[\delta^{t'-1}\left(V'_{\theta,t'}(x) - V_{\theta,t'}(x)\right)\right], \tag{A.23}$$

where $V_{\theta,t'}$ denotes the value function of type θ in the most-sincere equilibrium, the expectation conditions on type θ, and the term at $t' = T + 1$ equals zero. At every realization with $t' \leq T$, funding ends at x in the other equilibrium and $V'_{\theta,t'}(x) = 0$, whereas the belief x is funded in the most-sincere equilibrium and $V_{\theta,t'}(x)$ is at least the bound of Lemma 8, which is strictly positive. Every realized belief is interior, and thus each type realizes every path with positive probability, and the right-hand side of equation (A.23) is strictly negative for each type whenever some realization has $t' \leq T$. Averaging equation (A.23) over the types with weights μ_1 and $1 - \mu_1$ gives the recipient's payoff difference at the initial prior, which is non-negative by hypothesis, and thus every realization has $t' = T + 1$, the two equilibria induce the same play at every realized belief, and the payoffs of every player coincide at the initial prior. Then, no monotone equilibrium yields any player a strictly higher payoff than a most-sincere equilibrium, concluding part (iii).

$\boldsymbol{T = \infty}$**.** Now suppose $T = \infty$, where every monotone equilibrium is stationary, and write $V_D^{\mu',\lambda}$ for the donor's best-response value of the proof of Theorem 1 against the recipient that attacks in the shine at every belief and defends in the rain if and only if $\mu \geq (\mu^-)^{-1}(\mu')$.

First, the funding threshold $\bar{\mu}$ of every monotone equilibrium satisfies

$$\bar{\mu} \geq \mu^{att} - \frac{\delta\lambda(a-c)}{(1-\delta)(s_1 - s_0)a}. \tag{A.24}$$

By the proof of Theorem 1, both types pool on attack at $\bar{\mu}$, a flow payoff revealing the rain ends funding, and otherwise the belief stays at $\bar{\mu}$ or moves to $\mu^+(\bar{\mu})$. The donor's value V_D in this equilibrium is at least $c/(1-\delta)$ at $\bar{\mu}$ and satisfies

$$V_D(\bar{\mu}) = s(\bar{\mu})a + \delta\left[(1-\lambda)V_D(\bar{\mu}) + \lambda s(\bar{\mu})V_D(\mu^+(\bar{\mu})) + \lambda(1-s(\bar{\mu}))\frac{c}{1-\delta}\right],$$

and substituting $V_D(\bar{\mu}) \geq c/(1-\delta)$ and $V_D(\mu^+(\bar{\mu})) \leq a/(1-\delta)$ and rearranging gives inequality (A.24). Because $(\mu^-)^{-1}(\mu^{att}) > \mu^{att}$, a strictly positive bound on λ, folded into λ^*, then makes $(\mu^-)^{-1}(\bar{\mu}) > \mu^{att}$ in every monotone equilibrium.

Second, $V_D \leq V_D^{\bar{\mu},\lambda}$ at every funded belief. Substitute $V_D^{\bar{\mu},\lambda}$ as the continuation value into the donor's one-step value under the equilibrium's recipient play, with funding on $[\bar{\mu}, 1]$, and, at any belief, call the one-step value when both types always attack the one-step attack value and the one-step value when both types play sincerely the one-step sincere value. At unfunded beliefs the substitution returns $c/(1-\delta) \leq V_D^{\bar{\mu},\lambda}$; where both recipients always attack, it differs from the one-step attack value of $V_D^{\bar{\mu},\lambda}$ only in the rain-revealed continuation $c/(1-\delta) \leq V_D^{\bar{\mu},\lambda}$; where both are sincere, the one-step values coincide. Where the equilibrium's recipient always attacks and the recipient defining $V_D^{\bar{\mu},\lambda}$ is sincere, we have $\mu \geq (\mu^-)^{-1}(\bar{\mu}) > \mu^{att}$, and it suffices that

$$V_D^{\underline{\bar{\mu}},\lambda}(\mu) - \Big[s(\mu)V_D^{\underline{\bar{\mu}},\lambda}(\mu^+(\mu)) + (1-s(\mu))V_D^{\underline{\bar{\mu}},\lambda}(\mu^-(\mu))\Big] \le (1-s(\mu))b, \tag{A.25}$$

because then the one-step attack value falls short of the one-step sincere value, which is at most $V_D^{\underline{\bar{\mu}},\lambda}(\mu)$. The bracketed average is at least $s(\mu)a/(1-\delta)$, because $V_D^{\underline{\bar{\mu}},\lambda} \ge s(\cdot)a/(1-\delta)$ at every belief, as established in the proof of Theorem 1. If $V_D^{\underline{\bar{\mu}},\lambda}(\mu) = c/(1-\delta)$, the left-hand side of inequality (A.25) is at most zero because $s(\mu)a > c$; otherwise $V_D^{\underline{\bar{\mu}},\lambda}(\mu)$ is the one-step sincere value and the left-hand side is $u^{sin}(\mu)$ less $1-\delta$ times the average, at most $(1-s(\mu))b$. Repeated substitution, which is monotone and converges to V_D, gives $V_D \le V_D^{\underline{\bar{\mu}},\lambda}$.

Third, let $\underline{\bar{\mu}}$ be the smallest candidate $\mu' \in (0, \mu^{att}]$ at which the funding expression of equation (A.18) under $V_D^{\mu',\lambda}$, evaluated at μ', is at least $c/(1-\delta)$. This expression is continuous in μ' and weakly profitable at the threshold constructed in the proof of Theorem 1.

Fix any $\mu'' \in (0, \mu^{att})$ and any $\lambda < (1-\delta)\,(c - s(\mu'')a)\,/\,(\delta s(\mu'')a)$, a strictly positive bound folded into λ^*. At any candidate $\mu' \le \mu''$, if the supremum of $V_D^{\mu',\lambda}$ over beliefs weakly below μ' exceeded $c/(1-\delta)$, the funding branch of equation (A.18) would bound it by $s(\mu'')a(1-\delta+\delta\lambda)/(1-\delta)^2 < c/(1-\delta)$. Then, $V_D^{\mu',\lambda} = c/(1-\delta)$ at every belief weakly below μ', and the funding expression at μ' lies strictly below $c/(1-\delta)$.

Then, $\underline{\bar{\mu}} \ge \mu''$ exists, every smaller candidate is strictly unprofitable, and continuity gives the donor indifference of the construction at $\underline{\bar{\mu}}$. The verification in the proof of Theorem 1 then applies and delivers a monotone equilibrium with funding threshold $\underline{\bar{\mu}}$ and sincerity threshold $\underline{\mu}^* := (\mu^-)^{-1}(\underline{\bar{\mu}})$.

In any monotone equilibrium, funding at $\bar{\mu}$ is optimal, and its value there is the same funding expression evaluated at $V_D \le V_D^{\underline{\bar{\mu}},\lambda}$. Then, $\bar{\mu} \ge \underline{\bar{\mu}}$ and $\mu^* \ge (\mu^-)^{-1}(\bar{\mu}) \ge \underline{\mu}^*$, the infimum of the sincerity thresholds is attained, and most-sincere equilibria exist, which completes part (i).

Moreover, every most-sincere equilibrium has $(\mu^-)^{-1}(\bar{\mu}) \le \mu^* = \underline{\mu}^*$, and thus $\bar{\mu} = \underline{\bar{\mu}}$ and it funds on $[\underline{\bar{\mu}}, 1]$. A defense at $\underline{\mu}^*$ lands on the funded belief $\underline{\bar{\mu}}$ and is strictly optimal by equation (A.13) and inequality (A.14), so the equilibrium defends in the rain exactly at beliefs weakly above $\underline{\mu}^*$, and the donor's value is $V_D^{\underline{\bar{\mu}},\lambda}$ at every funded belief.

For part (iii), fix an initial prior μ_1, a most-sincere equilibrium, with value functions V_D^{sin} and V_θ^{sin}, and any monotone equilibrium, with value functions V_D' and V_θ', that yields the donor and the recipient a weakly higher payoff at μ_1. Substituting $V_D^{\underline{\bar{\mu}},\lambda}$ into the donor's one-step value under the latter equilibrium's recipient play and funding rule returns at most $V_D^{\underline{\bar{\mu}},\lambda}$ at every belief, by the cases of the second step, using $\mu^* \ge \underline{\mu}^*$ from the third step in the sincere case; in the only new case, always attacking at $\mu \ge \underline{\mu}^*$, inequality (A.25) at $\underline{\bar{\mu}}$ places the one-step attack value below the one-step sincere value $V_D^{\underline{\bar{\mu}},\lambda}(\mu)$ by at least $(1-s(\mu))b(1-\delta(1-\lambda)) > 0$. Then, $V_D' \le V_D^{\underline{\bar{\mu}},\lambda}$ at every belief. Let t' be the first period in which the two equilibria induce different play at the realized belief, with $t' = \infty$ if none arrives, and x the belief realized in period t'. Values are bounded and discounted, so

$$V_D^{sin}(\mu_1) - V_D'(\mu_1) = \mathrm{E}[\delta^{t'-1}(V_D^{sin}(x) - V_D'(x))],$$

with the event $t' = \infty$ contributing zero, and likewise for the recipient of each type θ under E_θ. Play at x can differ in only two ways: only the most-sincere equilibrium funds at x, or both fund and the latter equilibrium pools on attack at $x \geq \underline{\mu}^*$ while the most-sincere equilibrium plays sincerely. Either way x is funded in the most-sincere equilibrium, where $V_D^{sin} = V_D^{\bar{\mu},\lambda} \geq V_D'$, so every realization term is non-negative, the hypothesis makes the expectation at most zero, and every realization with $t' < \infty$ has $V_D^{sin}(x) = V_D'(x)$. The second kind is impossible: $V_D'(x)$ is its one-step attack value at V_D', at most that at $V_D^{\bar{\mu},\lambda}$, which falls short of $V_D^{sin}(x)$ as above. Then, at every realization with $t' < \infty$, funding ends at x in the latter equilibrium, $V_\theta'(x) = 0$, and $V_\theta^{sin}(x)$ is at least the strictly positive lower bound of equation (A.13). Every belief on the common path is interior, each type realizes every path with positive probability, and each type's decomposition is strictly negative whenever $t' < \infty$ with positive probability. Averaging over the types with weights μ_1 and $1 - \mu_1$ gives the recipient's payoff difference at μ_1, which is strictly negative as well, contradicting the hypothesis. Then, $t' = \infty$ almost surely, payoffs coincide at μ_1, and no monotone equilibrium yields any player a strictly higher payoff, concluding part (iii) at $T = \infty$. □

E.3 Proposition 1

Proposition 1 (Comparative Statics). *Suppose $\lambda < \lambda^*$. Then,*

1. *in every period t, the thresholds $\underline{\bar{\mu}}_t$ and $\underline{\mu}_t^*$ are weakly increasing in c and weakly decreasing in a, b, and δ; and*
2. *at $\lambda = 0$, the threshold $\underline{\bar{\mu}}^{[\infty]}$ is decreasing in s_0 and s_1, and the threshold $\underline{\mu}^{*[\infty]}$ is decreasing in s_0, and increasing in s_1 if and only if $s_1 > 1/2 + c/(2a)$.*

Proof. The proof of Lemma 1 gives $\underline{\bar{\mu}}_t = \bar{\mu}^{[j]}$ and $\underline{\mu}_t^* = \mu^{*[j]}$ at every t, with $j = T - t + 1$ periods remaining. Equations (A.3) and (A.4) do not depend on T, and thus define $\bar{\mu}^{[j]}$ at every $j \geq 1$. The bound of Proposition 5 is positive at every T, and $\lambda = 0$ lies below each. Then, the sequence $\bar{\mu}^{[j]}$ is nondecreasing by Proposition 5, bounded above by μ^{att} by inequality (A.6), and convergent. By continuity of $(\mu^-)^{-1}$, $\mu^{*[j]}$ converges to $\mu^{*[\infty]}$. We also have $\bar{\mu}^{[\infty]} \geq \bar{\mu}^{[1]} = \mu^{sin}$, and $\mu^{sin} > 0$ by equation (A.2).

Next, we prove the characterization of $\bar{\mu}^{[\infty]}$. Take first $\lambda = 0$ and fix any $\mu \in (\mu^{sin}, \mu^{att})$. By inequality (A.7), we have $\mu^{*[j]} > \mu^{att} > \mu$ at every $j \geq 2$, and thus equation (A.4) applies at μ and, at $\lambda = 0$,

$$\Psi^{[j]}(\mu) = s(\mu)a - c + \delta \max\left\{0, \Psi^{[j-1]}(\mu)\right\}.$$

If $\Psi^{[j-1]}(\mu) \leq 0$ at some j, then $\Psi^{[j']}(\mu) = s(\mu)a - c < 0$ at every $j' > j$. If instead $\Psi^{[j']}(\mu) > 0$ at every $j' < j$, then iterating,

$$\Psi^{[j]}(\mu) = (s(\mu)a - c)\frac{1-\delta^{j-1}}{1-\delta} + \delta^{j-1}\left(u^{sin}(\mu) - c\right),$$

which is negative at every sufficiently large j because $s(\mu)a < c$. In either case, $\Psi^{[j]}(\mu) < 0$ at every sufficiently large j, and thus $\bar{\mu}^{[j]} \geq \mu$ eventually and $\bar{\mu}^{[\infty]} \geq \mu$. Then, $\bar{\mu}^{[\infty]} = \mu^{att}$.

Now take $\lambda \in (0, \lambda^*)$ and fix any $\mu \leq \mu^{att}$ and any $j \geq 2$. By inequality (A.7), we have $\mu < \mu^{*[j]}$, and thus equation (A.4) applies at μ. Whichever of equations (A.3) and (A.4) applies at $\mu^+(\mu)$, its flow payoff is at least $s(\mu^+(\mu))a - c$ and its continuation values are non-negative, and thus $\Phi^{[j-1]}(\mu^+(\mu)) \geq s(\mu^+(\mu))a - c$. Then, dropping the non-negative term $\delta(1-\lambda)\Phi^{[j-1]}(\mu)$ from equation (A.4),

$$\Psi^{[j]}(\mu) \geq s(\mu)a - c + \delta\lambda s(\mu)\left[s(\mu^+(\mu))a - c\right].$$

The right-hand side is continuous in μ, and at $\mu = \mu^{att}$ it equals $\delta\lambda s(\mu^{att})\left[s(\mu^+(\mu^{att}))a - c\right]$, which is strictly positive because $\mu^+(\mu^{att}) > \mu^{att}$. Then, the right-hand side is strictly positive on an interval of beliefs with upper end μ^{att}, every belief in the interval is funded at every $j \geq 2$, and $\bar{\mu}^{[1]} = \mu^{sin} < \mu^{att}$. The thresholds $\bar{\mu}^{[j]}$ therefore remain bounded away from μ^{att} at every j.

For part (i), fix two values $c' > c$ of the donor's flow cost, holding the remaining primitives fixed, and suppose $\lambda < \lambda^*$ holds at both. We show by induction on j that $\Psi_{c'}^{[j]}(\mu) \leq \Psi_c^{[j]}(\mu)$ at every μ, that $\bar{\mu}_{c'}^{[j]} \geq \bar{\mu}_c^{[j]}$, and thus that $\mu_{c'}^{*[j+1]} \geq \mu_c^{*[j+1]}$, where subscripts denote the value of c. At $j = 1$, we have $\Psi^{[1]}(\mu) = u^{sin}(\mu) - c$, which is decreasing in c, and $\bar{\mu}^{[1]} = \mu^{sin}$, which is increasing in c. Now fix $j \geq 2$ and suppose all three claims hold below j. At every $\mu < \mu_c^{*[j]}$, both versions take equation (A.4), and $\Psi_{c'}^{[j]}(\mu) \leq \Psi_c^{[j]}(\mu)$ by the inductive hypothesis. At every $\mu \geq \mu_{c'}^{*[j]}$, both versions take equation (A.3), and the comparison is identical. At every $\mu \in [\mu_c^{*[j]}, \mu_{c'}^{*[j]})$, the c' version takes equation (A.4) and the c version equation (A.3), and thus $\Psi_{c'}^{[j]}(\mu)$ is at most equation (A.4) evaluated at c as above, which is at most $\Psi_c^{[j]}(\mu)$ by Lemma 4 and $\Phi^{[j-1]} \geq 0$. Then, every belief with $\Psi_{c'}^{[j]}(\mu) \geq 0$ has $\Psi_c^{[j]}(\mu) \geq 0$, and thus $\bar{\mu}_{c'}^{[j]} \geq \bar{\mu}_c^{[j]}$, and $\mu_{c'}^{*[j+1]} \geq \mu_c^{*[j+1]}$ because $(\mu^-)^{-1}$ is increasing. The arguments for a, for b, and for δ are symmetric with the inequalities reversed. A larger a raises the flow payoff of equations (A.3) and (A.4) and lowers μ^{sin}, a larger b raises the flow payoff of equation (A.3), leaves equation (A.4) unchanged, and lowers μ^{sin}, and a larger δ raises the non-negative continuation term of each and leaves μ^{sin} unchanged. In each argument, only the version with the larger parameter value plays sincerely on the interval between the two values of $\mu^{*[j]}$, and there the surplus at the smaller value is at most equation (A.4) at the larger value, which is at most equation (A.3) at the larger value by Lemma 4 and $\Phi^{[j-1]} \geq 0$.

Now let $T = \infty$, and suppose $\lambda < \lambda^*$ at both parameter values. By Lemma 1, the thresholds are $\underline{\mu}$, the smallest candidate $\mu' \in (0, \mu^{att}]$ at which funding is weakly profitable in the sense of its proof, and $\underline{\mu}^* = (\mu^-)^{-1}(\underline{\bar{\mu}})$, where $(\mu^-)^{-1}$ is increasing and free of a, b, c, and δ. For a fixed candidate, the recipient strategy defining $V_D^{\mu',\lambda}$ is the same at both parameter values, and $V_D^{\mu',\lambda}$ net of $c/(1-\delta)$ is the unique bounded fixed point of a monotone map whose flow terms $s(\mu)a - c$ and $u^{sin}(\mu) - c$ fall in c and rise in a and, for the sincere term, in b, and whose non-negative continuation terms carry weight δ, which a larger δ raises. Repeated substitution therefore orders the fixed points across parameter values, and with them the funding expression at every candidate. Take $c' > c$: every candidate weakly profitable at c' either exceeds $\mu_c^{att} \geq \underline{\bar{\mu}}_c$ or is weakly profitable at c, so $\underline{\bar{\mu}}_{c'} \geq \underline{\bar{\mu}}_c$. For a larger value of a, b, or δ, the smaller value's threshold $\underline{\bar{\mu}}$ is either weakly profitable at the larger value or exceeds μ^{att} at the larger value, and either way the larger value's threshold is weakly smaller. In every case, $\underline{\mu}^*$ moves with $\underline{\bar{\mu}}$.

For part (ii), the characterization above gives $\bar{\mu}^{[\infty]} = \mu^{att}$ at $\lambda = 0$. The derivative of μ^{att} in s_0 has the sign of $c/a - s_1$ and in s_1 the sign of $s_0 a - c$, both negative by Assumption 1. Solving $\mu^-(x) = \mu^{att}$ for x gives

$$\mu^{*[\infty]} = \frac{(1-s_0)(c-s_0 a)}{(s_1 - s_0)\,[a(1-s_1) + c - s_0 a]},$$

whose derivative in s_0 has the sign of $[a(2s_0 - 1) - c]\,(1-s_1)(s_1 a - c) < 0$, because $a(2s_0 - 1) \leq s_0 a < c$, and whose derivative in s_1 has the sign of $(1-s_0)(c - s_0 a)(2s_1 a - a - c)$, which is positive if and only if $2s_1 > 1 + c/a$. ☐

E.4 Proposition 2

Proposition 2. *For each $\lambda \in [0, \lambda^*)$, fix any most-sincere equilibrium with funding and sincerity thresholds $\bar{\mu}_t(\lambda)$ and $\mu_t^*(\lambda)$. Denote by $V_{\theta,t}(\mu;\lambda)$ the recipient continuation value given λ. Then, for every $t < T$, each θ, and all but finitely many $\mu \in (\bar{\mu}_t(0), \mu_t^*(0))$,*

$$\left.\frac{\partial V_{\theta,t}(\mu;\lambda)}{\partial \lambda}\right|_{\lambda=0} < 0.$$

Proof. Suppose first that $T < \infty$ and fix $t < T$. By the proof of Lemma 1, the thresholds satisfy $\bar{\mu}_{t'}(\lambda) = \bar{\mu}^{[T-t'+1]}(\lambda)$ at every $t' \geq 2$ and $\mu_{t'}^*(\lambda) = (\mu^-)^{-1}(\bar{\mu}_{t'+1}(\lambda))$ at every $t' < T$, and thus these thresholds converge to their values at $\lambda = 0$ by Lemma 6. In the first period, the belief is the initial prior, which does not immediately end the game by assumption. To see that the first identity extends to $t' = 1$, note that the donor's first-period funding surplus is $\Psi^{[T]}(\mu;\lambda)$ at every belief that no composition of μ^+ and μ^- carries to a threshold, which is strictly positive above $\bar{\mu}^{[T]}(\lambda)$, and thus part (iii) of Definition 1 gives $\bar{\mu}_1(\lambda) \leq \bar{\mu}^{[T]}(\lambda)$, while the reverse inequality holds in every monotone equilibrium.

By part (iii) of Theorem 1, both types play the same action at every funded belief in each realization of the weather, and thus every observation moves the donor's belief from its current value $\tilde{\mu}$ to $\tilde{\mu}$, $\mu^+(\tilde{\mu})$, or $\mu^-(\tilde{\mu})$. Then, at any λ, every belief realized in period t' arises from applying the maps μ^+ and μ^- to the period-t belief at most $t'-t$ times.

These compositions are finitely many and strictly increasing, and thus only finitely many beliefs in $(\bar{\mu}_t(0), \mu_t^*(0))$ are carried by some composition to some threshold $\bar{\mu}_{t'}(0)$ or $\mu_{t'}^*(0)$. Fix any μ in this interval that is not. Then, each belief realized in a period t' lies strictly above or strictly below $\bar{\mu}_{t'}(0)$, and likewise for $\mu_{t'}^*(0)$, and by the convergence above there is a $\bar{\lambda} > 0$ such that each realized belief lies on the same side of $\bar{\mu}_{t'}(\lambda)$ and of $\mu_{t'}^*(\lambda)$ at every $\lambda \in [0, \bar{\lambda})$.

Fix $\lambda \in [0, \bar{\lambda})$ and any belief $\tilde{\mu}$ realized in a period t' with $\tilde{\mu} > \bar{\mu}_{t'}(\lambda)$, which the donor funds. By Proposition 6, we have $\tilde{\mu} > \bar{\mu}_{t'}(\lambda) \geq \bar{\mu}_{t'+1}(\lambda)$, and $\mu^+(\tilde{\mu}) > \tilde{\mu}$, and thus the donor funds at $\tilde{\mu}$ and at $\mu^+(\tilde{\mu})$ in period $t'+1$. If $\tilde{\mu} < \mu_{t'}^*(\lambda)$, both types pool on attack by part (iii) of Theorem 1, the belief stays at $\tilde{\mu}$ when no flow payoff is observed and moves to $\mu^+(\tilde{\mu})$ when the flow payoff reveals a shine, and $\tilde{\mu} < \mu_{t'}^*(\lambda) = (\mu^-)^{-1}(\bar{\mu}_{t'+1}(\lambda))$ gives $\mu^-(\tilde{\mu}) < \bar{\mu}_{t'+1}(\lambda)$, and thus a flow payoff revealing the rain ends funding. Then,

$$V_{\theta,t'}(\tilde{\mu};\lambda) = s_\theta a + \delta\left[(1-\lambda)V_{\theta,t'+1}(\tilde{\mu};\lambda) + \lambda s_\theta V_{\theta,t'+1}(\mu^+(\tilde{\mu});\lambda)\right]. \tag{A.26}$$

If instead $\tilde{\mu} > \mu_{t'}^*(\lambda)$, both types play sincerely, the posterior is $\mu^+(\tilde{\mu})$ after an attack and $\mu^-(\tilde{\mu})$ after a defense regardless of whether the flow payoff is observed, and $\tilde{\mu} > (\mu^-)^{-1}(\bar{\mu}_{t'+1}(\lambda))$ gives $\mu^-(\tilde{\mu}) > \bar{\mu}_{t'+1}(\lambda)$, and thus both posteriors remain funded. Then,

$$V_{\theta,t'}(\tilde{\mu};\lambda) = s_\theta\left[a + \delta V_{\theta,t'+1}(\mu^+(\tilde{\mu});\lambda)\right] + (1-s_\theta)\left[b + \delta V_{\theta,t'+1}(\mu^-(\tilde{\mu});\lambda)\right]. \tag{A.27}$$

In the terminal period, both types play sincerely, there is no continuation, and $V_{\theta,T}(\tilde{\mu};\lambda) = s_\theta a + (1-s_\theta)b$ at every funded $\tilde{\mu}$. The terminal value is constant in λ, and equations (A.26) and (A.27) build each value from later values by sums and products of polynomials in λ, and thus, by backward induction, $V_{\theta,t'}(\tilde{\mu};\lambda)$ is a polynomial in λ on $[0,\bar{\lambda})$ at every realized belief $\tilde{\mu} > \bar{\mu}_{t'}(\lambda)$. Denote by $D_{\theta,t'}(\tilde{\mu})$ its derivative at $\lambda = 0$.

We show $D_{\theta,t'}(\tilde{\mu}) \leq 0$ at every such belief by backward induction, where $D_{\theta,T}(\tilde{\mu}) = 0$ because the terminal value is constant in λ. At a belief with $\tilde{\mu} > \mu_{t'}^*(0)$, differentiating equation (A.27) at $\lambda = 0$ gives

$$D_{\theta,t'}(\tilde{\mu}) = \delta s_\theta D_{\theta,t'+1}(\mu^+(\tilde{\mu})) + \delta(1-s_\theta)D_{\theta,t'+1}(\mu^-(\tilde{\mu})) \leq 0$$

by the inductive hypothesis. At a belief with $\tilde{\mu} < \mu_{t'}^*(0)$, differentiating equation (A.26) at $\lambda = 0$ gives

$$D_{\theta,t'}(\tilde{\mu}) = \delta D_{\theta,t'+1}(\tilde{\mu}) + \delta\left[s_\theta V_{\theta,t'+1}(\mu^+(\tilde{\mu});0) - V_{\theta,t'+1}(\tilde{\mu};0)\right]. \tag{A.28}$$

An attack yields the recipient a flow payoff of a in the shine and zero in the rain, a defense yields b in the rain and zero in the shine, and thus the expected flow payoff of the recipient in any period is at most $s_\theta a + (1 - s_\theta)b$ and

$$V_{\theta,t'+1}(\mu^+(\tilde{\mu}); 0) \leq (s_\theta a + (1 - s_\theta)b)\frac{1 - \delta^{T-t'}}{1 - \delta}.$$

At $\lambda = 0$, the belief stays put in every period pooling on attack, and in every sincere period the realized belief lies above the sincere threshold of its period, so its downward update lands weakly above the next period's funding threshold. Together with $\tilde{\mu} > \bar{\mu}_{t'}(\lambda) \geq \bar{\mu}_{t''}(\lambda)$ at every $t'' > t'$ by Proposition 6, every belief reached from $\tilde{\mu}$ remains funded and the expected flow payoff is at least $s_\theta a$ in every period, whether play pools or is sincere. Then,

$$V_{\theta,t'+1}(\tilde{\mu}; 0) \geq s_\theta a \frac{1 - \delta^{T-t'}}{1 - \delta}.$$

Then, the bracketed term of equation (A.28) is at most

$$s_\theta(s_\theta a + (1 - s_\theta)b)\frac{1 - \delta^{T-t'}}{1 - \delta} - s_\theta a \frac{1 - \delta^{T-t'}}{1 - \delta} = -s_\theta(1 - s_\theta)(a - b)\frac{1 - \delta^{T-t'}}{1 - \delta} < 0,$$

and thus $D_{\theta,t'}(\tilde{\mu}) < 0$, which completes the induction. The belief μ is realized in period t with $\mu > \bar{\mu}_t(\lambda)$ and $\mu < \mu_t^*(0)$, and thus $D_{\theta,t}(\mu) < 0$.

Now let $T = \infty$, and write $\bar{\mu}(\lambda)$, $\mu^*(\lambda) = (\mu^-)^{-1}(\bar{\mu}(\lambda))$, and $V_\theta(\mu; \lambda)$ for the stationary thresholds and values of Lemma 1, so that it suffices to prove the derivative inequality at every $\mu \in (\bar{\mu}(0), \mu^*(0))$. By the proof of Theorem 1, we have $\bar{\mu}(\lambda) < \mu^{att}$ at every $\lambda > 0$ and $\bar{\mu}(0) = \mu^{att}$, so $\mu^*(\lambda) < \mu^*(0)$, and inequality (A.24) gives $\bar{\mu}(\lambda) \to \bar{\mu}(0)$ and $\mu^*(\lambda) \to \mu^*(0)$ as $\lambda \to 0$. Fix a type θ, a belief $\mu \in (\bar{\mu}(0), \mu^*(0))$, and $\bar{\lambda} \in (0, \lambda^*)$ with $\mu < \mu^*(\lambda)$ at every $\lambda < \bar{\lambda}$. At every $\lambda \in [0, \bar{\lambda})$, both types pool on attack at the funded belief μ, the belief stays at μ when no flow payoff is observed, moves to the funded belief $\mu^+(\mu)$ when a flow payoff reveals a shine, and a flow payoff revealing the rain ends funding, so

$$V_\theta(\mu; \lambda) = \frac{s_\theta a + \delta\lambda s_\theta V_\theta(\mu^+(\mu); \lambda)}{1 - \delta(1 - \lambda)},$$

and, with $V_\theta(\mu; 0) = s_\theta a/(1 - \delta)$, direct computation gives

$$\frac{V_\theta(\mu; \lambda) - V_\theta(\mu; 0)}{\lambda} = \frac{\delta s_\theta \left[(1 - \delta)V_\theta(\mu^+(\mu); \lambda) - a\right]}{(1 - \delta(1 - \lambda))\,(1 - \delta)}. \tag{A.29}$$

Next, $V_\theta(\mu^+(\mu); \lambda) \to V_\theta(\mu^+(\mu); 0)$ as $\lambda \to 0$. Couple the equilibria across λ on a common weather sequence and common independent uniform draws, with the flow payoff observed when the draw is at most λ. Both types play the same action at every funded belief, and thus every belief realized from $\mu^+(\mu)$ within N periods lies among the at most N-fold compositions of μ^+ and μ^- applied to $\mu^+(\mu)$, a finite set. Each such belief retains its funding and play at every sufficiently small λ: beliefs weakly above μ^{att} are funded at every λ and beliefs below are eventually unfunded, beliefs above $\mu^*(0)$ are sincere at every λ and, at beliefs weakly below that are funded, both types eventually attack regardless of the weather, and at $\mu^*(0)$ itself a defense lands on the funded belief $\mu^-(\mu^*(0)) = \bar{\mu}(0) \geq \bar{\mu}(\lambda)$, so both types are sincere there at every λ by equation (A.13) and inequality (A.14). Then, for every N there is a $\lambda_N > 0$ such that, at every $\lambda < \lambda_N$, the coupled paths from $\mu^+(\mu)$ coincide through period N unless some of the first N draws is at most λ, an event of probability at most $N\lambda$, and values are bounded by $a/(1-\delta)$, so $|V_\theta(\mu^+(\mu); \lambda) - V_\theta(\mu^+(\mu); 0)| \leq (N\lambda + \delta^N)2a/(1-\delta)$, and taking $\lambda \to 0$ and then $N \to \infty$ gives the convergence. Then, taking $\lambda \to 0$ in equation (A.29), the derivative exists and equals

$$\left.\frac{\partial V_\theta(\mu; \lambda)}{\partial \lambda}\right|_{\lambda=0} = \frac{\delta s_\theta \left[(1-\delta)V_\theta(\mu^+(\mu); 0) - a\right]}{(1-\delta)^2} \leq -\frac{\delta s_\theta(1-s_\theta)(a-b)}{(1-\delta)^2} < 0,$$

by equation (A.13) and $a > b$. □

E.5 Proposition 3

Now, fix a game of length T and an initial belief μ_1. For each type θ, define

$$F_\theta(j) := s_\theta a \frac{1-\delta^{j-1}}{1-\delta} + \delta^{j-1}\left(s_\theta a + (1-s_\theta)b\right) \tag{A.30}$$

at every $j \geq 1$. For a candidate block length $k \leq T-1$, consider the two conditions

$$\mu^-(\mu_1) < \bar{\mu}^{[T-k]}, \tag{A.31}$$

$$\delta^k F_0(T-k) > b + \delta(1-s_0)b\frac{1-\delta^{k-1}}{1-\delta}, \tag{A.32}$$

and let $\bar{k} := m+1$, where m is the largest integer in $\{1, \ldots, T-1\}$ such that conditions (A.31) and (A.32) hold after replacing k by every $k' \leq m$.

Proposition 3. *Let $T < \infty$, $\lambda = 0$, and $\mu_1 \in [\underline{\mu}_1(0), \underline{\mu}_1^*(0))$. Then, there exists a $\bar{k} \geq 2$ such that, for every $k < \bar{k}$, the commitment game has a monotone equilibrium and, in every monotone equilibrium, both types of recipient attack at every belief reached in the block.*

Proof. We show the proposition holds with the block bound $\bar{k}$ defined at the start of this subsection.

At $k = 1$, condition (A.31) is $\mu^-(\mu_1) < \bar{\mu}^{[T-1]}$, which is $\mu_1 < \mu_1^*(0)$, and condition (A.32) is $\delta F_0(T-1) > b$. Rearranging inequality (3) gives $s_0 a/(1-\delta) \geq s_0 a + (1-s_0)b$, and thus $F_0(j) \geq s_0 a + (1-s_0)b$ at every j. Then, $\delta F_0(T-1) \geq \delta\left(s_0 a + (1-s_0)b\right) > b$, and thus $\bar{k} \geq 2$. Fix $k < \bar{k}$. From period $k+1$, play follows the most-sincere equilibrium selected in the definition of the commitment game, its funding threshold with $T-k$ periods remaining is $\bar{\mu}^{[T-k]}$ by Lemma 1, and thus a recipient is funded after the block if and only if the block-end belief is at least $\bar{\mu}^{[T-k]}$.

We record two facts about every monotone equilibrium of the commitment game. First, at every belief $\tilde{\mu}$ and every period, the posterior after an attack is at least $\tilde{\mu}$. In either probability case, part (i) of Definition 1 gives the claim. Second, at every belief at least $\bar{\mu}^{[T-k]}$ with j periods remaining in the game, the value of a recipient of type θ satisfies

$$F_\theta(j) \le V_\theta \le (s_\theta a + (1-s_\theta)b)\,\frac{1-\delta^j}{1-\delta}. \tag{A.33}$$

The upper bound holds at every belief, because no play yields an expected flow payoff above $s_\theta a + (1-s_\theta)b$ in any period. For the lower bound, consider the deviation that attacks whenever funded in every remaining nonterminal period, including the discretionary periods after the block, and plays sincerely at the terminal period. An attack weakly raises the donor's belief by the first fact, the funding thresholds weakly decline over time, and thus the deviation preserves funding through the terminal period once the block-end belief is at least $\bar{\mu}^{[T-k]}$. It yields the expected flow payoff $s_\theta a$ in each of the first $j-1$ remaining periods and $s_\theta a + (1-s_\theta)b$ at the terminal period, and collecting the flows gives $F_\theta(j)$.

Next, we construct a monotone equilibrium. At every block-period belief $\mu \in (0,1)$, both types attack in both realizations of the weather, the donor's posterior after an attack is her prior, and her posterior after the off-path defense is zero. At each belief $\mu \in \{0,1\}$, both types play sincerely and the posterior remains the prior after every observation. From period $k+1$, play follows the most-sincere equilibrium selected in the definition of the commitment game. Part (i) of Definition 1 holds at every belief. Across the funded beliefs in $(0,1)$ within the block, the probability of defending in the rain is zero, which satisfies part (ii), part (ii) holds across such beliefs in the discretionary periods because the continuation is a monotone equilibrium, and no belief in $(0,1)$ arises off the path within the block. Funding is exogenous within the block and follows the most-sincere threshold after it, which satisfies part (iii) as adapted to the commitment game. Then, the profile is monotone. On the path of play, the belief remains μ_1, which is at least $\bar{\mu}^{[T-k]}$ because $\mu_1 \ge \underline{\bar{\mu}}_1(0) = \bar{\mu}^{[T]} \ge \bar{\mu}^{[T-k]}$ by the proof of Lemma 1 and Proposition 5. Consider a deviation by a recipient of type θ to a defense in a block period with $n-1$ committed periods remaining afterward. The deviation yields at most the flow payoff b now, at most the expected flow payoff $s_\theta a + (1-s_\theta)b$ in each remaining committed period, and a continuation of zero after the block, because the posterior is zero, every subsequent posterior is zero, and the belief zero lies below $\bar{\mu}^{[T-k]}$. Equilibrium play yields the expected flow payoff $s_\theta a$ in each subsequent committed period and the continuation $V_\theta(\mu_1)$ after the block, which is at least $F_\theta(T-k)$ by inequality (A.33). The deviation is therefore unprofitable if

$$\delta^n F_\theta(T-k) \ge b + \delta(1-s_\theta)b\frac{1-\delta^{n-1}}{1-\delta},$$

whose left-hand side falls and right-hand side rises as n grows, and thus inequality (A.32) at $n=k$ covers every deviation by the weak type, and the strong type's inequality follows from the weak type's because $F_1 \geq F_0$ and $(1-s_1)b \leq (1-s_0)b$. Then, the profile is a monotone equilibrium.

The donor initiates the commitment in period 1. By inequality (A.31) and Proposition 5, in each of the first k periods we have

$$\mu^-(\mu_1) < \bar{\mu}^{[T-k]} \leq \bar{\mu}^{[j-1]},$$

where j is the number of periods remaining. Thus, $\mu_1 < (\mu^-)^{-1}(\bar{\mu}^{[j-1]}) = \mu^{*[j]}$, so the uncommitted most-sincere equilibrium also pools on attack throughout the first k periods. The commitment and no-commitment games therefore induce the same play through the commitment block and, by construction, the same continuation afterward. Their funding surpluses consequently coincide at $\Psi^{[T]}(\mu_1)$, which is non-negative because $\mu_1 \geq \underline{\bar{\mu}}_1(0) = \bar{\mu}^{[T]}$.

Next, we show that, in every monotone equilibrium, at every block-period belief $\tilde{\mu} < \bar{\mu}^{[T-k]}$ with $n \geq 1$ committed periods remaining, every posterior realized at $\tilde{\mu}$ lies below $\bar{\mu}^{[T-k]}$, the block-end belief lies below $\bar{\mu}^{[T-k]}$ with probability one under the weak type's play, and the weak type's continuation value is at most $(s_0 a + (1-s_0)b)\frac{1-\delta^n}{1-\delta}$. We argue by induction on n, with the convention that at $n=0$ the block has ended, the value above is zero, and the block-end belief is $\tilde{\mu}$ itself. Now fix $n \geq 1$ and suppose the claims hold at $n-1$. Suppose first that the attack has positive probability at $\tilde{\mu}$ and the posterior after an attack is at least $\bar{\mu}^{[T-k]}$. A weak recipient who attacks now and in every remaining period reaches the block's end with a belief at least $\bar{\mu}^{[T-k]}$, by the first fact, and thus collects at least

$$\delta s_0 a \frac{1-\delta^{n-1}}{1-\delta} + \delta^n F_0(T-k),$$

in addition to a flow payoff of a when the current weather shines. A weak recipient who defends collects at most b now, at most the expected flow payoff $s_0 a + (1-s_0)b$ in each remaining committed period, and nothing after the block, because the posterior after a defense lies weakly below $\tilde{\mu}$ by part (i) of Definition 1, so the inductive hypothesis applies. The defense is therefore strictly unprofitable for the weak type in both realizations of the weather if

$$\delta^n F_0(T-k) > b + \delta(1-s_0)b\frac{1-\delta^{n-1}}{1-\delta},$$

whose left-hand side falls and right-hand side rises as n grows, and thus condition (A.32) at $n=k$ covers every $n \leq k$. Then, the weak type attacks with probability one at $\tilde{\mu}$, Bayes' rule places the odds of the posterior after an attack at most at the odds of $\tilde{\mu}$, and the posterior lies below $\bar{\mu}^{[T-k]}$, a contradiction. Thus every posterior realized at $\tilde{\mu}$ lies below $\bar{\mu}^{[T-k]}$, because the posterior after an attack does, and the posterior after a defense either lies weakly below the belief $\tilde{\mu}$, when the attack has positive probability, or equals $\tilde{\mu}$, when it does not. Then, the inductive hypothesis applies at every belief realized in the following period, the block-end belief lies below $\bar{\mu}^{[T-k]}$ with probability one, and the weak type's value is at most $(s_0a+(1-s_0)b)\frac{1-\delta^n}{1-\delta}$, because no play yields a higher expected flow payoff in any period and the value after the block is zero. This completes the induction.

Finally, suppose some monotone equilibrium departs from pooled attack at some belief reached in the block, and consider the first block period with a departure, with $n \geq 1$ committed periods remaining. Every earlier block period pools on attack, and thus the belief entering this period is μ_1, which is at least $\bar{\mu}^{[T-k]}$ because $\mu_1 \geq \underline{\bar{\mu}}_1(0) = \bar{\mu}^{[T]} \geq \bar{\mu}^{[T-k]}$ by the proof of Lemma 1 and Proposition 5. Suppose some type θ defends with positive probability in some realization of the weather, and suppose first that the posterior after a defense is at least $\bar{\mu}^{[T-k]}$. Both posteriors then lie at least at $\bar{\mu}^{[T-k]}$, inequality (A.33) bounds the continuation values from both, and the two bounds differ by at most $(1-s_\theta)b\frac{1-\delta^{T-k+n-2}}{1-\delta}$. By Assumption 2, δ times this difference is less than b, and thus each type strictly prefers the defense in the rain and the attack in the shine, play is sincere, and the posterior after a defense is $\mu^-(\mu_1)$, which lies below $\bar{\mu}^{[T-k]}$ by condition (A.31), a contradiction. Suppose instead that the posterior after a defense lies below $\bar{\mu}^{[T-k]}$. The induction above bounds the weak type's continuation after a defense, the posterior after an attack is at least $\mu_1 \geq \bar{\mu}^{[T-k]}$, and inequality (A.33) bounds the weak type's continuation after an attack from below, so the comparison of the previous paragraph applies unchanged, condition (A.32) makes the defense strictly unprofitable for the weak type in both realizations of the weather, and the weak type attacks with probability one. The defense then comes only from the strong type, the posterior after a defense is one, and the posterior after an attack is interior, which violates part (i) of Definition 1. Then, no type defends at μ_1, which contradicts the departure. Thus, in every monotone equilibrium, both types of recipient attack at every belief reached in the commitment block. $\square$

E.6 Proposition 4

Proposition 4. *Let $\lambda = 0$. Then, for every $K \in \mathbb{N}$, there is a $\bar{\delta} < 1$ and $\bar{T} : (0,1) \to \mathbb{N}$ such that, for every $\delta > \bar{\delta}$, $T > \bar{T}(\delta)$, $\mu_1 \in [\underline{\bar{\mu}}_1(0), \underline{\bar{\mu}}^{[\infty]}(0)]$, $k \leq K$, and all (a, b, c, s_0, s_1) satisfying Assumptions 1–2, the commitment game has a monotone equilibrium and, in every monotone equilibrium, both types of recipient attack at every belief reached in the commitment block.*

Proof. Fix $K \in \mathbb{N}$, and let $\bar{\delta} \in (2/3, 1)$ be such that $\delta^K(3 - 2/\delta) > 4(1-\delta)$ at every $\delta > \bar{\delta}$, which exists because the left-hand side rises to one and the right-hand side falls to zero as $\delta \to 1$.

Now fix any primitives satisfying Assumptions 1 and 2 with $\delta > \bar{\delta}$.

By Lemma 1, we have $\underline{\bar{\mu}}_1(0) = \bar{\mu}^{[T]}$ at every T, and thus $\underline{\bar{\mu}}^{[\infty]}(0) = \lim_{j\to\infty} \bar{\mu}^{[j]}$ at $\lambda = 0$, and the proof of Proposition 1 shows this limit equals μ^{att}. The value $\lambda = 0$ lies below the bound of Proposition 5, so inequality (A.7) gives $\underline{\mu}_1^*(0) = \mu^{*[T]} \geq (\mu^-)^{-1}(\mu^{sin}) > \mu^{att}$ at every $T \geq 2$, and thus every $\mu_1 \in [\underline{\bar{\mu}}_1(0), \underline{\bar{\mu}}^{[\infty]}(0)] = [\underline{\bar{\mu}}_1(0), \mu^{att}]$ satisfies the hypothesis of Proposition 3 at every $T \geq 2$.

Condition (A.31) holds at every T and block length, because $\mu^-(\mu_1) \leq \mu^-(\mu^{att}) < \mu^{sin} \leq \bar{\mu}^{[T-k']}$ by Lemma 2 and Proposition 5. By Assumption 2, we have $\delta(1-s_0) < 1-\delta$, and thus the right-hand side of condition (A.32) is less than $2b$ at every block length.

Inequalities (3) and Assumption 2 give $b < s_0 a$. Fix a block length $k' \leq K$ and let $\bar{T}(\delta) := K + 1 + \lceil \ln 2 / \ln(1/\delta) \rceil$, which depends only on δ and K. At every $T > \bar{T}(\delta)$, equation (A.30) gives

$$\delta^{k'} F_0(T-k') \geq \delta^K s_0 a \frac{1-\delta^{T-k'-1}}{1-\delta} > 4 s_0 a \left(1 - \delta^{T-K-1}\right) \geq 2 s_0 a > 2b,$$

because $\delta^K > 4(1-\delta)$, $\delta^{T-K-1} \leq 1/2$ at every $T > \bar{T}(\delta)$, and $b < s_0 a$. Then, condition (A.32) holds at every $k' \leq K$, both conditions hold at every $k' \leq K$, and thus $\bar{k} > K$. Proposition 3, which holds with this $\bar{k}$, gives both claims at every block length $k \leq K$. □