WORKING PAPERS

N° 1712

February 2026

“Peace Talk and Conflict Traps”

Andrei Gyarmathy and Georgy Lukyanov

Toulouse School of Economics

Peace Talk and Conflict Traps

Andrei Gyarmathy* Georgy Lukyanov†

Abstract

Costly pre-play messages can deter unnecessary wars—but the same messages can also entrench stalemates once violence begins. We develop an overlapping-generations model of a security dilemma with persistent group types (normal vs. bad), one-sided private signaling by the current old to the current young, and noisy private memory of the last encounter. We characterize a stationary equilibrium in which, for an intermediate band of signal costs, normal old agents *mix* on sending a costly reassurance only after an alarming private history; the signal is kept marginally persuasive by endogenous receiver cutoffs and strategic mimicking by bad types. Signaling strictly reduces the hazard of conflict *onset*; conditional on onset, duration is unchanged in the private model but *increases* once a small probability of publicity (leaks) creates a public record of failed reconciliation. With publicity, play generically absorbs in a *peace trap* or a *conflict trap*. We discuss welfare and policy: when to prefer back-channels versus public pledges.

Keywords: signaling; cheap talk; security dilemma; overlapping generations; publicity; audience costs; reputation; conflict traps.
JEL Codes: D74; D83; C73; C72; D82.

1 Introduction

Costly pre-play messages can deter unnecessary wars, yet the very same messages can entrench stalemates once violence begins. This paper shows how both forces arise from one set of micro-foundations when communication interacts with persistent heterogeneity and thin memory.

We develop an overlapping-generations model of conflict between two groups. In each period an old member of one group meets a young member of the other and they play a symmetric security-dilemma (stag–hunt) stage game with actions "no conflict" and "conflict." Each group's type (normal or bad) is persistent; bad types always choose conflict, while normal types prefer peace if sufficiently reassured. Before actions, only the current old can

*CREST, École Polytechnique. Email: andrei.gyarmathy@polytechnique.edu.
†Toulouse School of Economics. Email: georgy.lukyanov@tse-fr.eu. We thank Geoffrey Barrows, Margarita Kirneva, Yukio Koriyama, David Kraus, Darius Lazea, Alessandro Riboni, George-Ioan Stoica, Hannah Tollié and Anna Vlasova for useful comments an ddiscussions. All remaining errors are our own.

send a costly private signal to the current young. Agents carry a one-bit, noisy private memory of what they believe they saw when young. We then allow a small probability that a costly signal leaks publicly, creating a simple reputation state that future agents observe.

First, we characterize a mixed-signaling equilibrium in the private model. For a non-degenerate interval of signal costs, a normal old mixes on whether to send the costly reassurance only after an alarming private history, while never signaling after a benign history. The receiver's posterior following a signal is kept at a cutoff by strategic mimicking from bad types, and the sender's mixing probability is pinned by an indifference condition that equates the expected gain from persuading a normal receiver to the cost of signaling. The result delivers explicit receiver response probabilities on either side of a policy threshold and yields a sharp comparative-statics map in the cost of signaling (Theorem 1 and Corollary 1).

Our second contribution concerns onset, duration, and welfare in the private model. Allowing costly signals strictly lowers the hazard of conflict onset among normal–normal meetings because some alarming histories are defused by reassurance that is marginally persuasive. Conditional on onset, however, the expected duration of a conflict spell among normals is unchanged: with purely private messages and one-bit memory, spells end only via misperception, making duration geometric with parameter ε. The combination implies higher stationary welfare among normals relative to the no-message benchmark (Propositions 1, 2, and 4).

Third, we show how even a small probability of publicity transforms the dynamics. When costly messages can leak, a public record forms. We construct two absorbing equilibria that coexist on a nonempty parameter region for any positive leak probability: a peace trap, where a leaked success coordinates on permanent peace, and a conflict trap, where a leaked failure coordinates on permanent distrust (Theorem 2). Outside absorption, leaked failures make receivers tougher on the conflict path and strictly increase the conditional duration of conflict spells on a nonempty parameter set, so welfare can go either way relative to private back-channels (Propositions 3 and 5).

The mechanisms are intuitive. An old normal who previously (mis)perceived conflict faces the strongest temptation to preempt. Reassurance is valuable only if it moves the young's beliefs enough to elicit restraint; if the cost is too low, bad types mimic and the signal loses bite; if too high, no one signals. In equilibrium the receiver is kept exactly at a cutoff after a signal, making the sender just indifferent about paying the cost—this lowers onset along the peace path. Once conflict starts, private signals do not alter the continuation because the environment provides no new public evidence. Publicity changes this by turning costly signals into public reputation events: a leaked success makes future receivers confident and locks in peace; a leaked failure makes them skeptical and hardens stalemates.

The policy implications are immediate. If audiences or institutions can coordinate on forgiveness after a single leaked success, public pledges dominate: they accelerate absorption into peace. Absent such coordination, private back-channels are preferable: they reduce onset without lengthening stalemates. Which regime is desirable therefore hinges on how public audiences process rare, salient signals of reconciliation.

2 Related Literature

Our paper connects three strands: dynamic models of conflict with persistent heterogeneity and misperception, models of communication in security dilemmas (both costly signals and cheap talk), and the growing literature contrasting private back-channels with public diplomacy and audience costs.

We build most directly on Acemoglu and Wolitzky (2014), who develop an overlapping-generations (OLG) model in which agents from two groups repeatedly meet in a stag–hunt environment under uncertainty about "bad" types and noisy private perceptions of past play. Their core insight is that private misperception plus persistent heterogeneity generate self-fulfilling cycles of distrust and conflict. Acemoglu and Wolitzky (2024) survey and extend this logic in a broader imperfect-information framework, emphasizing how private (mis)perceptions create divergent posteriors that can entrench conflict. Relative to this baseline, we introduce pre-play communication into the OLG security dilemma and characterize a stationary equilibrium in which *one-sided* costly reassurance is used selectively after alarming private histories. The novelty is twofold: first, we show that for an intermediate range of costs, normal senders *mix* only after bad private signals, with receiver cutoffs held at the margin by endogenous mimicking from bad types; second, we analyze how even small probabilities of publicity (leaks) transform the state into public reputation, creating absorption in peace or conflict traps. Neither selective one-sided signaling nor the leak-driven bifurcation appears in the existing OLG models of conflict.

Our communication environment relates to classic reassurance and signaling in international relations. Kydd (2000) shows how benign actors can separate via costly reassurance when trustworthy types are more willing to take risks for peace. We share the reassurance logic but place it in a dynamic OLG setting with private misperceptions and history-contingent use: normal senders reassure *only* after alarming private histories, keeping beliefs marginal by tolerating some mimicking. In contrast to audience-cost and hand-tying approaches, our signals are sunk-cost assurances that do not mechanically create domestic penalties but work by shifting the receiver's posterior. We also connect to the taxonomy in Quek (2021), who distinguishes multiple costly-signaling mechanisms; our baseline operates through sunk costs, while the leak extension effectively adds a public, hand-tying component.

A second anchor is the coordination-game approach to the security dilemma with private information and pre-play talk. Baliga and Sjöström (2004) study arms races with a small probability of dominant (always-arm) types and show how cheap talk can substantially improve outcomes, even driving the arms-race probability toward zero when dominance is very rare. We complement this by letting communication be *costly* and one-sided, embedded in an OLG setting with private misperception rather than two-sided private types alone. Cheap talk in our baseline is uninformative under standard assumptions; we show that adding an arbitrarily small single-crossing perturbation to message costs restores partial informativeness, thereby providing a disciplined route to cheap-talk extensions while preserving tractability. More generally, our equilibrium characterization links directly to the coordination-game logic while delivering new comparative statics for onset hazards and, under publicity, for conditional duration.

A third, and increasingly active, strand contrasts public signals, audience costs, and private back-channels. The audience-cost tradition following Fearon (1994, 1995) and Schultz (2001) emphasizes why public threats can be credible. Yet a parallel literature shows why secrecy can dominate: in Kurizaki (2007) private diplomacy can be *efficient* because going public can impose punitive audience costs on the *target*, making peaceful concession harder. Consistent with this, Katagiri and Min (2019) use a document-based analysis of the Berlin Crisis to show that private statements are often more credible to elites than public ones. We formalize a complementary mechanism: with purely private signals, reassurance reduces the *onset* hazard among normal–normal matches without lengthening conflict once it starts (duration remains governed by private misperception). By contrast, even small leak probabilities create public records that coordinate beliefs into either a peace trap or a conflict trap; leaked failures make receivers tougher and strictly lengthen conditional duration for a nonempty parameter region. Our private-vs-public comparison therefore synthesizes and sharpens the secrecy result in an OLG conflict environment with endogenous reputation dynamics.

Our paper also relates to broader work on crisis bargaining, signaling, and mediation. Classic models explain war as arising from information problems and commitment failures (Fearon, 1995; Leventoglu and Tarar, 2008; Tarar and Leventoglu, 2009). Cheap talk can help in some multidimensional settings or with particular structure (e.g., Trager, 2011), but communication is generally constrained by incentives (Crawford and Sobel, 1982). Mediation models show when biased mediators can credibly transmit information (Kydd, 2003). We differ by focusing on *direct*, history-contingent reassurance inside an OLG security dilemma with misperception and by allowing small publicity shocks to convert private persuasion into public coordination devices that reshape long-run states. On the repeated/OLG side, our emphasis on limited memory and stationary strategies is closest to the organizational OLG perspective of Lagunoff (2001), but our addition of signaling and leak-driven public records yields distinct comparative statics for both onset and duration.

In terms of contribution, we offer: (i) the first integration of selective, one-sided, *private* costly reassurance into the Acemoglu–Wolitzky OLG security-dilemma framework, with an explicit mixed-signaling equilibrium and closed-form belief cutoffs; (ii) a sharp separation between *onset* and *duration* in the private model (onset falls, conditional duration is unchanged), providing clean welfare implications for normal–normal encounters; (iii) a publicity extension with leak probability that generates a tractable reputation state and a generic bifurcation into peace and conflict traps, with testable predictions that leaked failures lengthen stalemates; and (iv) a disciplined cheap-talk extension via an ϵ-cost single-crossing perturbation that restores partial informativeness without sacrificing tractability. These features, taken together, differentiate our analysis from prior OLG conflict models, reassurance models with purely static separation, and audience-cost/secrecy accounts that lack a dynamic, leak-driven public-state mapping.

The rest of the paper proceeds as follows. Section 3 presents the environment, payoffs, limited memory, and the one-sided signaling protocol. Section 4 characterizes the mixed-signaling equilibrium in the private model, analyzes onset, duration, and welfare, and then turns to publicity, absorption, and duration. Section 6 concludes. The appendix contains full proofs, robustness checks, and a minimal cheap-talk extension that restores informativeness under an arbitrarily small single-crossing perturbation to message costs.

3 Model

3.1 Population, types, and horizon

Time is discrete, $t = 0, 1, 2, \ldots$. There are two groups, $\mathcal{G} = \{\mathrm{A}, \mathrm{B}\}$. In each t, a pair interacts: an *old* member of one group meets a *young* member of the other. Roles alternate deterministically (even t: old from A meets young from B; odd t: vice versa). Each individual lives for two periods: young at t and old at $t+1$.

Each group $G \in \mathcal{G}$ has a *persistent* type $\theta_G \in \{\mathrm{n}, \mathrm{b}\}$ ("normal" or "bad"), drawn from the common prior $\mu \in (0, 1)$ and fixed forever; types are independent across groups. Agents know their own group's type but not the opponent's.

3.2 Stage-game actions and payoffs

In any interaction, both players simultaneously choose $a \in \{\mathrm{N}, \mathrm{C}\}$, interpreted as *no conflict* and *conflict*. For a normal player, payoffs are the canonical security-dilemma (stag-hunt) normalization:[1]

	N (opp.)	C (opp.)
N (self)	1	$-\ell$
C (self)	g	0

with parameters $\ell > 0$ (second-strike disadvantage) and $g \in (0, 1)$ (first-strike advantage). Bad types are behavioral: they choose $a = \mathrm{C}$ with probability 1.

Each agent realizes the stage payoff when young and when old; agents maximize the sum of their two stage payoffs net of any signaling costs.

3.3 Private monitoring and limited memory

When young at date t, the agent privately observes a noisy signal $\tilde{a}_t^{\mathrm{opp}} \in \{\mathrm{N}, \mathrm{C}\}$ of the *old opponent's* action with symmetric error $\varepsilon \in [0, 1/2)$:

$$\mathbb{P}(\tilde{a}_t^{\mathrm{opp}} = a_t^{\mathrm{opp}}) = 1 - \varepsilon, \qquad \mathbb{P}(\tilde{a}_t^{\mathrm{opp}} \neq a_t^{\mathrm{opp}}) = \varepsilon.$$

This is the *only* history carried into the old period at $t+1$. We denote the resulting private state by $h \in \{g, b\}$, where $h = g$ iff $\tilde{a}_t^{\mathrm{opp}} = \mathrm{N}$ and $h = b$ iff $\tilde{a}_t^{\mathrm{opp}} = \mathrm{C}$.

3.4 One-sided private signaling protocol (no leaks)

Within a period t:

(i) **Matching.** Old from group G meets young from group $-G$ (roles alternate by t).

[1] Any 2×2 coordination game with a Pareto dominant (N, N) can be affinely renormalized to $\left[\begin{smallmatrix} 1 & -\ell \\ g & 0 \end{smallmatrix}\right]$ with $g \in (0, 1)$ and $\ell > 0$. This pins comparative statics on g and ℓ rather than absolute utilities.

(ii) **Message.** The old may send a *private binary message* $m \in \{0,1\}$ to the young. Sending $m = 1$ costs $k > 0$; $m = 0$ is free. The message is observed only by the receiver; there are no leaks (exogenous publicity prob. $q = 0$).

(iii) **Actions.** After observing m, both simultaneously choose $a \in \{\mathrm{N}, \mathrm{C}\}$.

(iv) **Payoffs and observation.** Stage payoffs are realized. The young privately observes a noisy signal of the old's action as in Section 3.3 and carries $h \in \{g, b\}$ into their old period.

Bad types always play $a = \mathrm{C}$;[2] they may still choose $m = 1$ (mimicking).

3.5 Strategies, beliefs, and equilibrium

A *strategy* specifies:

Old (sender), normal type $\theta_G = \mathrm{n}$. (i) A messaging rule $\sigma^{G,\mathrm{n}}_{\mathrm{msg}} : \{g, b\} \to \Delta(\{0,1\})$; (ii) an action rule $\sigma^{G,\mathrm{n}}_{\mathrm{old}} : \{g, b\} \times \{0,1\} \to \Delta(\{\mathrm{N}, \mathrm{C}\})$.

Old (sender), bad type $\theta_G = \mathrm{b}$. A messaging rule $\sigma^{G,\mathrm{b}}_{\mathrm{msg}} \in \Delta(\{0,1\})$; action is degenerate: $a = \mathrm{C}$.

Young (receiver). An action rule $\sigma^{G}_{\mathrm{young}} : \{0,1\} \to \Delta(\{\mathrm{N}, \mathrm{C}\})$.

Beliefs. Let $\beta_{-G}(\theta \mid m)$ denote the receiver's posterior that the sender is type $\theta \in \{\mathrm{n}, \mathrm{b}\}$ after m. Posteriors are by Bayes' rule on-path using prior μ and equilibrium messaging frequencies; off-path beliefs are arbitrary but fixed.[3]

Definition 1. A *stationary*[4] *perfect Bayesian equilibrium* (SPBE) is a profile

$$\left(\sigma^{G,\mathrm{n}}_{\mathrm{msg}}, \sigma^{G,\mathrm{n}}_{\mathrm{old}}, \sigma^{G,\mathrm{b}}_{\mathrm{msg}}, \sigma^{G}_{\mathrm{young}}\right)_{G \in \mathcal{G}}$$

and beliefs $\left(\beta_G(\cdot \mid m)\right)$ such that:

(a) **Sequential rationality.** Given beliefs and opponents' stationary strategies, each normal old's messaging and action rules maximize expected payoff (net of cost k) for each private state $h \in \{g, b\}$; each young's action rule maximizes expected payoff given $\beta(\cdot \mid m)$; bad types' messaging is optimal given $a = \mathrm{C}$.

(b) **Belief consistency.** On-path beliefs use Bayes' rule; off-path beliefs and strategies satisfy (a).

(c) **Stationarity.** Strategies depend only on role, type, and h (for old) and on m (for young), not on calendar time.

[2] A microfoundation is a one-shot Prisoner's Dilemma for the bad type: C strictly dominates N for any opponent action. See Appendix B.1.

[3] We use beliefs that keep $m = 1$ at the cutoff where needed and assign pessimistic beliefs after off-path $m = 0$ at conflicting histories; see Appendix B.4.

[4] Stationarity means strategies depend only on role, type, and (for old) the private state h, not calendar time. This matches the OLG pairing and the i.i.d. private histories.

Table 1: Notation and primitives

Symbol	Meaning
$\mathcal{G} = \{\text{A}, \text{B}\}$	Groups; in period t old from one group meets young from the other
$\theta \in \{\text{n}, \text{b}\}$	Group type: normal vs bad (bad plays C)
$a \in \{\text{N}, \text{C}\}$	Actions: no conflict / conflict
$g,\ \ell$	First-strike advantage and second-strike loss (normal's payoffs)
ε	Private misperception rate of the opponent's action when young
$h \in \{g, b\}$	Private state carried to old age: perceived good/bad past action
μ	Prior that the opponent group is bad; $\mu_g < \mu_b$ are posteriors by h
k	Cost of sending the private costly message $m = 1$
q	Leak probability (publicity) that turns messages/actions public
$\bar{p} = \frac{\ell}{1+\ell-g}$	Receiver's belief cutoff on "opponent plays N"
$\bar{q} = \bar{p}$	Sender's action cutoff on "receiver plays N"
$k_\star = \frac{g\ell}{1+\ell-g}$	Policy threshold for sender's post-signal action
$R \in \{\emptyset, \mathsf{S}, \mathsf{F}\}$	Public record: none / leaked success / leaked failure

Parameters. $\Theta = (\mu, \varepsilon, k, g, \ell)$ with $\mu \in (0,1)$, $\varepsilon \in [0, 1/2)$, $k > 0$, $g \in (0,1)$, $\ell > 0$.

4 Results

This section develops the core comparative statics and equilibrium characterization. We first set up thresholds and notation that will be used throughout. We then establish a receiver cutoff lemma, show how one–sided costly signaling can implement posteriors at the margin[5] via strategic mimicking, and map best responses on both sides. The main theorem identifies a stationary mixed–signaling equilibrium in which normal olds send reassurance only after alarming private histories and the receiver's reaction is pinned by a simple cost–indifference condition. We then analyze how signaling affects conflict onset and duration, before turning to publicity (leaks) and welfare.

4.1 Thresholds and notation

Before solving behavior, we collect the key objects that collapse choices to one–dimensional cutoffs. The receiver's willingness to cooperate is summarized by a belief cutoff, which in turn induces a cooperation–probability cutoff for the sender's best reply. These cutoffs yield a compact expression for the sender's gross payoff as a function of the receiver's cooperation probability, and they isolate the two policy pivot points that drive our comparative statics.

Define the canonical cutoff[6] for best responses:

$$\bar{p} \;\equiv\; \frac{\ell}{1+\ell-g} \in (0,1).$$

[5] We keep $\mathbb{P}(\theta = \text{n} \mid m = 1) = \bar{p}$ by choosing the bad sender's mix p_b proportional to the normal's s at $h = b$. This is the classic "keep me indifferent" device adapted to one-sided signaling.

[6] The cutoff $\bar{p} = \ell/(1+\ell-g)$ equates the receiver's expected payoffs from N and C given a belief that the sender plays N with probability p.

If a player believes the opponent will play N with probability p, a normal prefers N to C iff $p \geq \bar{p}$. Let μ_g, μ_b denote a normal old sender's posterior (from their private state $h \in \{g, b\}$) that the opponent group is bad, with $0 < \mu_g < \mu_b < 1$. For shorthand, set[7]

$$\bar{q} \equiv \bar{p}, \qquad k_\star \equiv g\,\bar{q} = \frac{g\,\ell}{1 + \ell - g}.$$

4.2 Receiver cutoffs and best responses

The receiver's side of the game is governed by a single inequality: for a given belief that the sender will play N, when is cooperation (playing N) optimal? The next lemma formalizes this cutoff and makes clear why our later constructions keep posteriors exactly at the margin after a costly signal. Intuitively, we will engineer beliefs so that the receiver is just indifferent after a signal and (weakly) pessimistic otherwise.[8]

Lemma 1. *For any $p \in [0, 1]$ equal to the probability that the sender will play* N*, a normal receiver is indifferent between* N *and* C *iff $p = \bar{p}$; prefers* N *iff $p > \bar{p}$; and prefers* C *iff $p < \bar{p}$.*

Proof. The payoff difference is $p \cdot 1 + (1 - p)(-\ell) - pg = p(1 + \ell - g) - \ell$, zero iff $p = \bar{p}$. □

4.3 Implementing posteriors with one-sided signaling

Costly reassurance is only useful if it moves the receiver's belief across the cooperation cutoff. Because bad types may mimic, we cannot rely on literal separation. This subsection shows that, with one–sided messages and appropriate mixing by bad senders, we can pin the posterior after a costly signal *exactly* at the cutoff while keeping the no–signal posterior weakly below it. This implementability result is the workhorse behind the mixed–signaling equilibrium.

Let $\pi_b \in (0, 1)$ be the stationary probability that a normal old is in state $h = b$; set $\pi_g = 1 - \pi_b$. Let $s \in [0, 1]$ be the probability a *normal* old with $h = b$ sends $m = 1$, while normals with $h = g$ never send; let $p_b \in [0, 1]$ be the probability a *bad* old sends $m = 1$. Bayes' rule gives

$$\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \frac{(1 - \mu)\,\pi_b\,s}{(1 - \mu)\,\pi_b\,s + \mu\,p_b}, \tag{1}$$

$$\mathbb{P}(\theta = \mathrm{n} \mid m = 0) = \frac{(1 - \mu)\big[\pi_g + \pi_b(1 - s)\big]}{(1 - \mu)\big[\pi_g + \pi_b(1 - s)\big] + \mu(1 - p_b)}. \tag{2}$$

Lemma 2. *Fix $\bar{p} \in (0, 1)$, $\mu \in (0, 1)$, and (π_b, π_g) with $\pi_b > 0$. There exists $(s, p_b) \in (0, 1] \times [0, 1)$ such that*

$$\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \bar{p} \quad \textit{and} \quad \mathbb{P}(\theta = \mathrm{n} \mid m = 0) \leq \bar{p}.$$

[7] Because the sender's tradeoff is $q(1 + \ell) - \ell$ vs qg, the best reply switches exactly at $q = \bar{q} = \bar{p}$.

[8] We break ties in favor of the equilibrium mixing probabilities (any $q_1 \in [0, 1]$ at the cutoff is optimal). This convention is innocuous and standard in cutoff constructions.

In particular, setting $p_b = \kappa s$ *with*

$$\kappa \;\equiv\; \frac{1-\bar p}{\bar p}\cdot\frac{1-\mu}{\mu}\cdot\pi_b,$$

pins $\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \bar p$ *for any* $s > 0$, *and choosing*

$$s \;\in\; \left(1 - \frac{1}{\pi_b}\left[\tfrac{\bar p\mu}{(1-\bar p)(1-\mu)} - \pi_g\right],\; 1\right]$$

ensures $\mathbb{P}(\theta = \mathrm{n} \mid m = 0) \le \bar p$.

4.4 Sender/receiver best responses

Given the cutoff posteriors, the receiver can rationally mix after a costly signal and decline to cooperate otherwise. On the sender side, optimal actions reduce to a simple piecewise–linear payoff in the receiver's cooperation probability. We derive this mapping and use it to express the sender's indifference condition as a closed–form relation between the signal cost and the receiver's response.

Let q_1 (resp. q_0) be the probability a *normal* receiver plays N after $m = 1$ (resp. $m = 0$). If posteriors equal $\bar p$ after $m = 1$ and do not exceed $\bar p$ after $m = 0$, any $(q_1, q_0) \in [0, 1]^2$ is optimal for the normal receiver. A normal old sender's action best response after message-i is: play N iff $q_i \ge \bar q$, else play C. Define the sender's gross payoff vs. a normal receiver as

$$V(q) \;\equiv\; \max\{\, q(1+\ell) - \ell,\;\; qg \,\},$$

and $V(q) \equiv 0$ vs. a bad receiver (who plays C).

4.5 Mixed signaling after an alarming private history

We now assemble the ingredients into the main equilibrium. When costs are too high, nobody signals; when costs are too low, bad types mimic and the signal loses bite. For an intermediate band of costs, a normal old *mixes* on sending a costly reassurance precisely after an alarming private history, while never signaling after a benign one. The receiver's cooperation after a signal is chosen to keep the sender indifferent, yielding explicit formulas and a clean switch in the sender's post–signal action at a single threshold.

Theorem 1. *Fix primitives* $\Theta = (\mu, \varepsilon, k, g, \ell)$ *and a belief pair* (μ_g, μ_b) *with* $0 < \mu_g < \mu_b < 1$. *Suppose Lemma 2 holds. Then for any*

$$k \;\in\; \left(0,\; 1 - \mu_b\right],$$

there exists a stationary PBE with:

(i) ***State-contingent one-sided messaging.*** *Normal old with* $h = g$ *sends* $m = 0$*; normal old with* $h = b$ *mixes, sending* $m = 1$ *with prob.* $s \in (0, 1)$*; bad old mixes with prob.* $p_b = \kappa s$ *(from Lemma 2).*

(ii) ***Receiver posteriors at cutoffs.*** *$\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \bar{p}$ and $\mathbb{P}(\theta = \mathrm{n} \mid m = 0) \le \bar{p}$.*

(iii) ***Receiver randomization.*** *The normal receiver's cooperation probability after $m = 1$ is*

$$(q_1, q_0) \;=\; \begin{cases} \left(\dfrac{k}{(1-\mu_b)g},\ 0\right) & \text{if } k \in \big(0,\ (1-\mu_b)k_\star\big], \\ \left(\dfrac{k/(1-\mu_b)+\ell}{1+\ell},\ 0\right) & \text{if } k \in \big[(1-\mu_b)k_\star,\ 1-\mu_b\big], \end{cases}$$

so the sender's action after $m = 1$ is C in the first range and N in the second.

(iv) ***Sender indifference at $h = b$.*** *A normal old with $h = b$ satisfies*

$$(1-\mu_b)\left[V(q_1) - V(q_0)\right] \;=\; k,$$

and therefore mixes; a normal old with $h = g$ strictly prefers not sending.

Proof sketch. Use Lemma 2 to keep posteriors at the cutoff after $m = 1$ and at/below after $m = 0$. Set $q_0 = 0$ so $V(q_0) = 0$. For $k \in (0, 1-\mu_b]$, pick q_1 so $(1-\mu_b)V(q_1) = k$. If $q_1 \le \bar{q}$, then $V(q_1) = q_1 g$ so $q_1 = k/[(1-\mu_b)g] \in (0, \bar{q}]$, feasible iff $k \le (1-\mu_b)k_\star$. If $q_1 \ge \bar{q}$, then $V(q_1) = q_1(1+\ell) - \ell$, hence $q_1 = (k/(1-\mu_b)+\ell)/(1+\ell) \in [\bar{q}, 1]$, requiring $k \ge (1-\mu_b)k_\star$. Receivers are indifferent (cutoff posteriors), senders best-respond at the action stage by definition of $V(\cdot)$. Messaging: the $h = b$ normal gains $(1-\mu_b)[V(q_1) - V(q_0)]$ when facing a normal receiver and pays k, so is indifferent; $h = g$ normals strictly prefer $m = 0$; set bad $p_b = \kappa s$ to maintain the cutoff posterior after $m = 1$. □

Corollary 1. *In Theorem 1:*

(a) q_1 increases in k piecewise: $q_1 = k/[(1-\mu_b)g]$ for $k \le (1-\mu_b)k_\star$ and $q_1 = (k/(1-\mu_b) + \ell)/(1+\ell)$ for $k \ge (1-\mu_b)k_\star$.

(b) Higher mistrust (larger μ_b) shrinks the cost interval $(0, 1-\mu_b]$ supporting mixed signaling.

(c) If $k > 1-\mu_b$, no normal type signals in any SPBE of this form; if $k < (1-\mu_b)k_\star$, post-signal play is C by the sender; if $k > (1-\mu_b)k_\star$, post-signal play switches to N by the sender.

Figure 1 maps the equilibrium policy regions in signal cost k: the mixing band is $k \in (0, 1-\mu_b]$, with a sender action switch at $k = (1-\mu_b)k_\star$ where $k_\star = \frac{g\ell}{1+\ell-g}$; outside the band, signaling is unused.

4.6 Onset vs. duration of conflict

Signaling affects conflict on two margins that need not move together:[9] the hazard that a conflict spell *starts*, and the expected *length* of a spell once it has begun. We quantify both. In the private model, reassurance lowers the onset hazard along the peace path, while

[9] Onset depends on pre-play persuasion; duration in the private model depends only on misperceiving C as N with probability ε each period, yielding geometric spells.

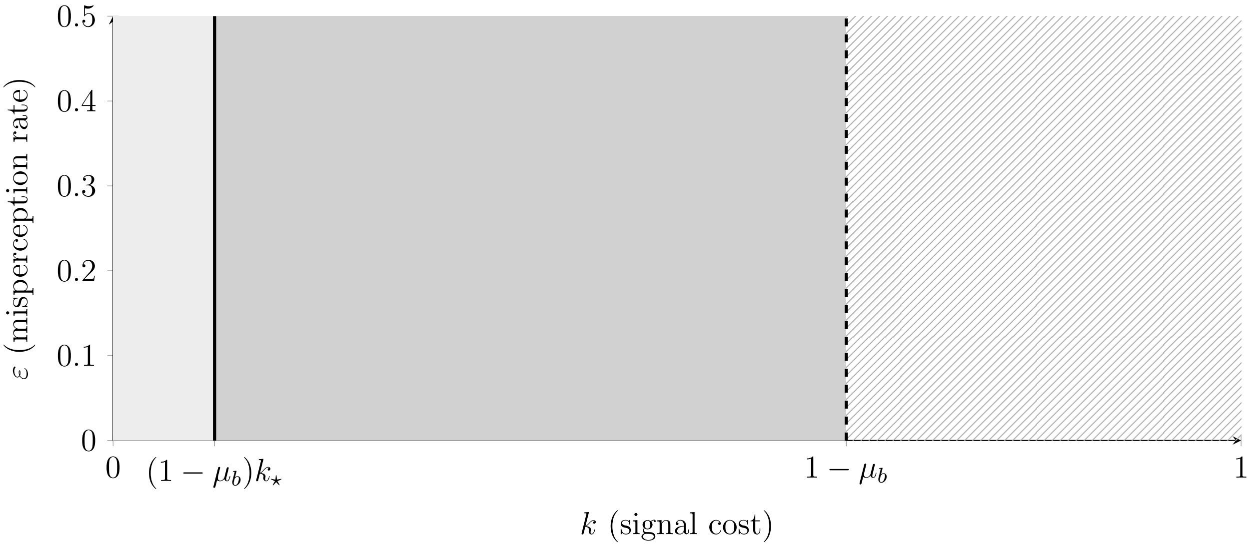

Figure 1: Policy regions as a function of signal cost k (horizontal) and misperception rate ε (vertical). For any fixed (μ_b, g, ℓ), the mixing band is $k \in \big(0,\, 1-\mu_b\big]$; within it, the sender's *post-signal* action switches from C to N at $k = (1-\mu_b)k_\star$ with $k_\star = \frac{g\ell}{1+\ell-g}$.

conditional duration remains tied to misperception. This separation will help us interpret welfare and will also frame the effects of publicity.

We define a period outcome as *peace* if both choose N and as *conflict* otherwise (i.e., at least one chooses C). A *conflict spell* is a (maximal) run of consecutive periods with conflict outcomes. We study (i) the *onset hazard* of a conflict spell, and (ii) the *expected duration* of a conflict spell, both conditional on the interacting agents being normal types; bad types trivially generate conflict in any period they are active.

Let $\pi_b \in (0,1)$ denote the stationary probability that a normal *old* is in the alarming private state $h = b$ when matched with a normal *young*; let $\pi_g = 1 - \pi_b$. Let $\mu_b \in (0,1)$ be the normal old's posterior (from $h = b$) that the opponent group is bad. Throughout this subsection we compare the SPBE constructed in Theorem 1 (one-sided signaling) to the no-signaling benchmark in which the messaging stage is absent (equivalently, $k = \infty$).

Proposition 1. *Fix primitives $\Theta = (\mu, \varepsilon, k, g, \ell)$ and beliefs (μ_g, μ_b) with $0 < \mu_g < \mu_b < 1$. Consider (i) the SPBE from Theorem 1 with one-sided private signaling and cost $k \in (0, 1-\mu_b]$, and (ii) the no-signaling SPBE (same primitives, same monitoring, no message stage).*

Let λ^{on}_{sig} and λ^{on}_{nosig} denote the per-period probability of conflict onset *conditional on both players being normal types and on the old's private state being $h = b$ (the only state from which onsets arise in the stationary environment).*[10]

Then

$$\lambda^{on}_{sig} \;\leq\; \lambda^{on}_{nosig},$$

[10]Formally, we condition on a pre-period information set in which the current old is normal and has $h = b$, the current young is normal, and there is no additional public history. The unconditional onset hazard among normals equals π_b times this conditional probability.

with strict *inequality on a nonempty parameter region (in particular, whenever the equilibrium uses $s > 0$ and either (a) $k > (1 - \mu_b)k_\star$ so the sender sometimes plays* N *after $m = 1$, or (b) $k \le (1 - \mu_b)k_\star$ and the receiver plays* N *after $m = 1$ with positive probability).*

Proof. In the no-message SPBE, a normal old with $h = b$ best-responds by C (the opponent normal receiver's belief that the old will play N is below $\bar{p}$), so the outcome is conflict with probability 1. Hence $\lambda^{\text{on}}_{\text{nosig}} = 1$.

In the signaling SPBE of Theorem 1, a normal old with $h = b$ sends $m = 1$ with prob. $s \in (0, 1)$ and $m = 0$ otherwise. After $m = 0$, the receiver's posterior does not exceed $\bar{p}$ by construction and the outcome is conflict with prob. 1 (sender best-responds by C). After $m = 1$, the receiver is indifferent (posterior equals $\bar{p}$) and plays N with prob. $q_1 \in (0, 1]$; moreover, if $k > (1-\mu_b)k_\star$, the sender best-responds by N. Therefore the probability of peace (both N) conditional on $h = b$ and $m = 1$ is at least $\mathbf{1}\{k > (1-\mu_b)k_\star\}\cdot 1 + \mathbf{1}\{k \le (1-\mu_b)k_\star\}\cdot 0$, and the probability of avoiding conflict (i.e., at least one N) is at least $q_1 > 0$. Hence

$$\lambda^{\text{on}}_{\text{sig}} \;=\; (1 - s)\cdot 1 \;+\; s \cdot \mathbb{P}(\text{conflict} \mid m = 1) \;\le\; 1,$$

with strict inequality whenever $s > 0$ and either the sender plays N after $m = 1$ (high-cost regime) or the receiver chooses N with positive probability after $m = 1$ (low-cost regime). Multiplying by π_b yields the same inequality for the unconditional onset hazard among normals. □

Consistent with Proposition 1, Figure 2 shows that the conditional onset hazard among normals drops only within the high-cost portion of the mixing band, where the receiver responds with $q_1(k) = \frac{k/(1-\mu_b)+\ell}{1+\ell}$ and the sender best-responds by N.

Proposition 2. *Consider a conflict spell that* has just started *in a period when both players are normal types. In the no-signaling SPBE, the sender (old) plays* C *with probability* 1 *whenever in the alarming state $h = b$, and the young's private observation of the old's action equals* C *with probability $1 - \varepsilon$ (and* N *with probability ε). Hence the next period's old (the current young) transitions to $h = b$ with probability $1 - \varepsilon$ and to $h = g$ with probability ε.*

In the signaling SPBE of Theorem 1, any conflict spell that starts among normals must, by construction, arise from the low-cost regime *$k \le (1 - \mu_b)k_\star$, in which the sender best-responds by* C *after $m = 1$ and also plays* C *after $m = 0$. Therefore the same transition applies: the next period's old has $h = b$ with probability $1 - \varepsilon$ and $h = g$ with probability ε.*

Consequently, in both environments the length D (in periods) of a conflict spell among normal types is geometric with success parameter ε, so

$$\mathbb{E}[D] \;=\; \frac{1}{\varepsilon}.$$

Proof. Immediate from the transition logic described: in each subsequent period of the spell, the old is in state $h = b$ and plays C, and the young observes the old's action through the same binary noise; the spell ends when the young misperceives C as N, which occurs with probability ε independently each period. □

Remark 1. Proposition 2 shows that in the *baseline* model with one-bit private memory and private messages, the *conditional* duration of a conflict spell among normals is unaffected

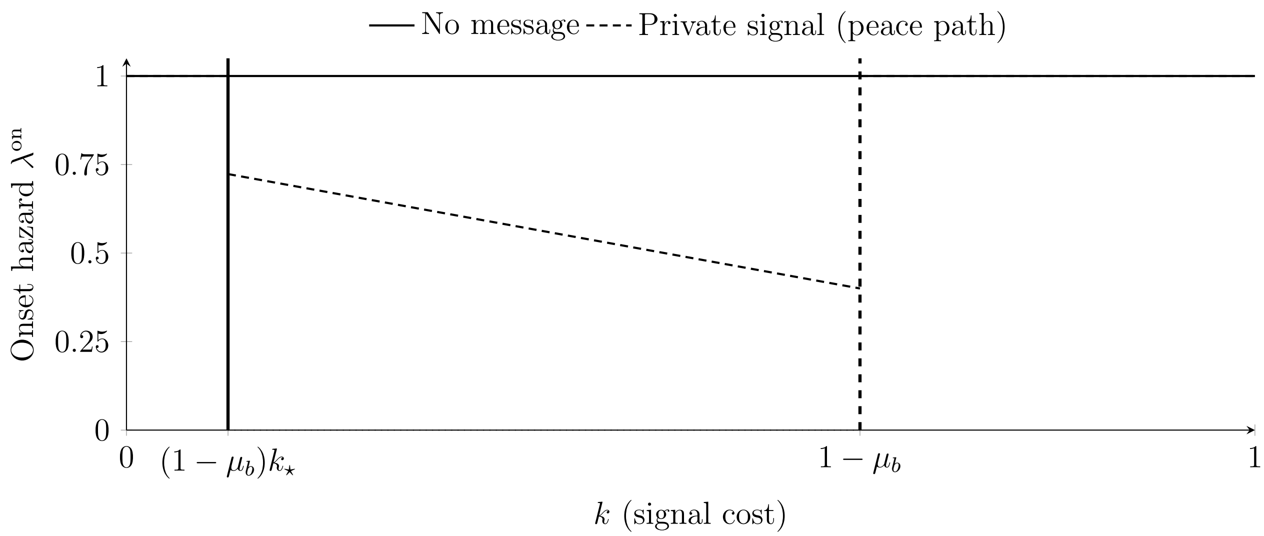

Figure 2: Conditional onset hazard among normal–normal matches with alarming private state ($h = b$). Without messages the hazard is 1. With private signaling, onset falls only in the high-cost mixing range where the sender plays N after a signal: $\lambda^{\text{on}}(k) = 1 - s \cdot q_1(k)$ with $q_1(k) = \frac{k/(1-\mu_b)+\ell}{1+\ell}$ and illustrative $s = 0.6$. For $k > 1 - \mu_b$ signaling is unused. Under publicity, once a leaked *success* occurs the process absorbs into peace (hazard 0 thereafter).

by signaling. The reason is selection: signaling reduces the onset hazard (Proposition 1), and any spell that still starts must come from the low-cost regime where the sender plays C regardless of messages, leaving the continuation identical to the no-message case.

As we show in the publicity extension (Section 5), even a small probability of public visibility of messages or actions generically breaks this equivalence: failed reconciliation attempts become part of the (public) record, making receivers more cautious in responding to future signals on the conflict path and *strictly increasing* expected duration on a nonempty parameter set.

Corollary 2. *Let Λ_{sig} and Λ_{nosig} denote the stationary fractions of periods with conflict (among normals) in the signaling and no-signaling SPBEs, respectively. Then*

$$\Lambda_{sig} \le \Lambda_{nosig},$$

with strict inequality whenever the conditions for Proposition 1's strict inequality hold and $\pi_b \in (0,1)$. The result follows because (i) the onset hazard falls (Proposition 1), while (ii) the conditional expected duration of a conflict spell is $\mathbb{E}[D] = 1/\varepsilon$ in both *environments (Proposition 2), and a renewal argument yields $\Lambda = \lambda^{on} \cdot \mathbb{E}[D] \cdot \pi_b$ up to a normalization constant.*

5 Extensions

We extend the baseline in two directions. First, we allow a small probability that costly messages leak and become public, generating a reputation state that future agents observe.

Second, we discuss communication without costs (cheap talk) and the minimal departures that restore informativeness. The publicity extension will introduce genuine path dependence—absorbing "peace" and "conflict" traps—while the cheap–talk discussion clarifies why costs (even tiny and asymmetric) matter.

5.1 Publicity extension: leaks and public record

Publicity turns private persuasion into a public coordination device. We enrich the timing so that a costly message can leak with exogenous probability, and leaked outcomes are recorded. The resulting public record shapes future receivers' cutoffs and creates two absorbing regimes:[11] after a leaked success, full cooperation becomes self–confirming; after a leaked failure, distrust does. We formalize these regimes and show how publicity can lengthen stalemates once conflict has started.

We modify the timing in Section 3.4 as follows. After the old chooses a private message $m \in \{0, 1\}$, Nature draws a leak $\mathsf{L} \in \{0, 1\}$ with $\mathbb{P}(\mathsf{L} = 1) = q \in (0, 1]$, independently across periods. If $\mathsf{L} = 1$ and $m = 1$, then the fact that a *costly* message was sent becomes public *before* actions; if $m = 0$, nothing becomes public. After actions, if $\mathsf{L} = 1$ and $m = 1$, the realized *action profile* is publicly observed as well (so third parties learn whether the costly peace signal was followed by peace or by conflict). Otherwise, no public information is released.

Let the *public record* (reputation) state be $R \in \{\emptyset, \mathsf{S}, \mathsf{F}\}$, carried from one period to the next and observed by all current players *before* messaging: $R = \emptyset$ denotes no public information; $R = \mathsf{S}$ ("success") denotes that in the most recent leak with $m = 1$ both players chose N (a *successful* reconciliation); $R = \mathsf{F}$ ("failure") denotes that in the most recent leak with $m = 1$ the period outcome was conflict (a *failed* reconciliation). The law of motion is:

$$R_{t+1} = \begin{cases} \mathsf{S} & \text{if } \mathsf{L}_t = 1,\ m_t = 1,\ \text{and } (a_t^{\text{old}}, a_t^{\text{young}}) = (\mathrm{N}, \mathrm{N}), \\ \mathsf{F} & \text{if } \mathsf{L}_t = 1,\ m_t = 1,\ \text{and } (a_t^{\text{old}}, a_t^{\text{young}}) \neq (\mathrm{N}, \mathrm{N}), \\ R_t & \text{otherwise.} \end{cases}$$

Strategies may now condition on (h, R) for old senders and on (R, m) for receivers. Belief consistency applies to both private posteriors (as before) and to public beliefs induced by the record R.

Benchmark conventions. For $q = 0$ we recover the baseline (private) model. When $q > 0$ but $R = \emptyset$, the construction of Theorem 1 can be replicated on the *peace path* (no public events yet). The novelty arises once $R \in \{\mathsf{S}, \mathsf{F}\}$.

5.1.1 Bifurcation in the leak probability

Public leaks make costly signals common knowledge. Intuitively, after a *public failure* ($R = \mathsf{F}$), receivers become stricter (interpreting future signals as cheap or manipulative), whereas after

[11] For any $q > 0$ and a nonempty parameter region (e.g., $k \leq 1 - \mu_b$), we construct SPBEs absorbing at $R = \mathsf{S}$ (peace) and at $R = \mathsf{F}$ (conflict). See Theorem 2.

a *public success* ($R = \mathsf{S}$), receivers become more receptive. The following result formalizes a regime switch.

Theorem 2. *Fix primitives $\Theta = (\mu, \varepsilon, k, g, \ell)$ and consider the publicity extension with any leak probability $q \in (0, 1]$. There exists a nonempty parameter region (e.g., all $k \leq 1 - \mu_b$) on which the following two stationary PBEs both exist:*

(i) ***Peace trap.*** *Starting from $R = \emptyset$ (or $R = \mathsf{S}$), costly messages are used on the peace path and, upon the first leak followed by (N, N), the public state absorbs at $R = \mathsf{S}$; thereafter players choose* N *forever and the long-run conflict frequency among normals is* 0*.*

(ii) ***Conflict trap.*** *Starting from $R = \emptyset$, if a leak occurs that is followed by conflict (which happens with positive probability in the cutoff construction), the public state absorbs at $R = \mathsf{F}$; thereafter receivers play* C *and senders best-respond with* C*, so the long-run conflict frequency among normals is* 1*.*

For $q = 0$ the baseline private model (Section 3) is recovered with the mixed-signaling equilibrium of Theorem 1; for any $q > 0$ both absorbing equilibria above coexist on a nonempty parameter region.

Proof sketch. The constructions in the main text already define two candidates: (a) a peace-trap SPBE that sets $q_1(R = \emptyset) = 1$, triggers (N, N) upon leak, and then plays N forever at $R = \mathsf{S}$; and (b) a conflict-trap SPBE that coincides with the cutoff construction at $R = \emptyset$ and enforces $q_1(R = \mathsf{F}) = q_0(R = \mathsf{F}) = 0$ after any leaked failure. For any $q > 0$, leaks arrive with positive probability each period; under (a) the first leak followed by peace absorbs at $R = \mathsf{S}$; under (b) any leaked failure absorbs at $R = \mathsf{F}$. Sequential rationality and belief consistency follow exactly as in the original proof (receivers' cutoffs on R are best responses; senders' messaging/actions are optimal given those responses). Nonemptiness of the parameter region is witnessed by $k \leq 1 - \mu_b$. □

Figure 3 summarizes the public-record dynamics: from $R = \emptyset$, a leaked success absorbs at $R = \mathsf{S}$ (permanent peace) whereas a leaked failure absorbs at $R = \mathsf{F}$ (permanent conflict), and both absorbing outcomes coexist for any $q > 0$.

5.1.2 Conditional duration with publicity

Public record of *failed* reconciliation ($R = \mathsf{F}$) hardens receivers on the conflict path and can strictly increase the *conditional* expected duration of conflict spells among normal types, even for arbitrarily small q.

Proposition 3. *Fix $\Theta = (\mu, \varepsilon, k, g, \ell)$ and consider the publicity extension with any $q \in (0, 1]$. Focus on the SPBE that coincides with Theorem 1 on the peace path when $R = \emptyset$ (so costly messages are used with mixing probability $s > 0$ at $h = b$). Then there exists a nonempty set of parameters (including all (μ, ε, k) with $s > 0$ and $k \leq (1 - \mu_b)k_\star$) such that the conditional expected duration $\mathbb{E}[D_q]$ of a conflict spell among normals satisfies*

$$\mathbb{E}[D_q] \; > \; \mathbb{E}[D_0] \; = \; \frac{1}{\varepsilon}.$$

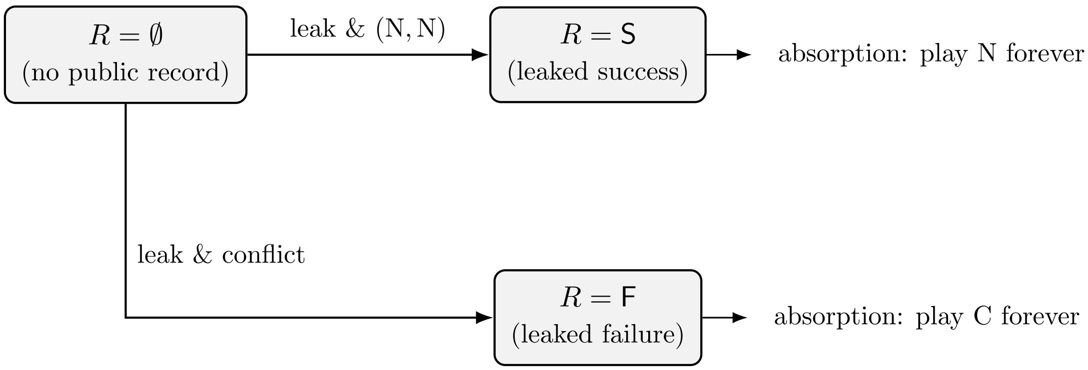

Figure 3: Publicity (leaks) turns private reassurance into public reputation. From $R = \emptyset$, a leaked *success* absorbs at $R = \mathsf{S}$ (permanent peace), while a leaked *failure* absorbs at $R = \mathsf{F}$ (permanent conflict). Both absorbing outcomes coexist for any $q > 0$ on a nonempty parameter region.

Proof sketch. With $q = 0$, Proposition 2 established $\mathbb{E}[D_0] = 1/\varepsilon$. When $q > 0$, during a conflict spell that starts at $h = b$, the sender mixes on $m = 1$ with prob. $s > 0$ and a leak occurs with prob. q. With probability $sq > 0$, a leak happens *before* the spell ends and, because in the low-cost regime the sender best-responds by C after $m = 1$, the leaked period records a *failure* ($R \leftarrow \mathsf{F}$). Under $R = \mathsf{F}$, receivers set $q_1(R = \mathsf{F}) = 0$ (or, more generally, reduce cooperation sufficiently) so that subsequent costly messages on the conflict path do not produce immediate exits. Resolving the spell then requires a rarer sequence (e.g., a further public success or multiple private misperceptions), strictly raising the expected length relative to the geometric baseline. A coupling argument yields $\mathbb{E}[D_q] \geq \mathbb{E}[D_0] + \delta$ for some $\delta > 0$ whenever $sq > 0$ and the $R = \mathsf{F}$ response is sufficiently strict, implying $\mathbb{E}[D_q] > \mathbb{E}[D_0]$ on the stated region. □

Remark 2. Publicity helps *before* conflict (it accelerates belief alignment toward peace when a success leaks) but can hinder *after* onset (a leaked failure freezes beliefs against reconciliation), thereby lengthening stalemates. Together with Theorem 2, this delivers the qualitative dichotomy in your conjecture: under sufficiently high publicity, play is absorbed either in a *peace trap* or a *conflict trap*; for lower publicity, cycles persist but leaked failures make conditional spells longer.

5.2 Welfare and policy

We evaluate welfare at the level of normal–normal matches and compare stationary outcomes across environments.[12] The private model lowers the fraction of periods in conflict without changing conditional duration, thus strictly improving welfare relative to no messages. With publicity, welfare becomes selection–dependent: absorption at a peace trap dominates, while absorption at a conflict trap backfires; outside absorption, onset reductions trade off against

[12]Total welfare scales the normal–normal component by $(1-\mu)^2$; meetings with a bad type are payoff-invariant to messaging, so they cancel in cross-environment comparisons.

longer spells. We translate these findings into simple guidance on when to prefer back–channels over public pledges.

We evaluate social welfare at the interaction level. Let $U^{\mathrm{n}}(a, a')$ denote a normal player's one-shot payoff from Section 3.2; define stage welfare among two normals by

$$W_{\mathrm{NN}}(a, a') \;\equiv\; U^{\mathrm{n}}(a, a') \;+\; U^{\mathrm{n}}(a', a),$$

so $W_{\mathrm{NN}}(\mathrm{N}, \mathrm{N}) = 2$, $W_{\mathrm{NN}}(\mathrm{C}, \mathrm{C}) = 0$, and $W_{\mathrm{NN}}(\mathrm{N}, \mathrm{C}) = W_{\mathrm{NN}}(\mathrm{C}, \mathrm{N}) = g - \ell \le 0$. When at least one party is bad, conflict occurs and we take the realized welfare to be $W_{\mathrm{NB}} \le 0$ if exactly one is bad and $W_{\mathrm{BB}} = 0$ if both are bad.[13]

We compare stationary expected welfare among *normal-normal* matches across environments. Let Λ be the stationary fraction of periods with conflict among normals; equivalently Λ is the product of the onset hazard and the expected spell length divided by the expected cycle length in the renewal decomposition. Then stationary expected welfare among normals is

$$\bar{W}_{\mathrm{NN}} \;=\; (1 - \Lambda) \cdot 2 \;+\; \Lambda \cdot \bar{W}_{\mathrm{NN}}^{\mathrm{conf}}, \qquad \text{where} \quad \bar{W}_{\mathrm{NN}}^{\mathrm{conf}} \in \{0,\, g - \ell\} \le 0,$$

depending on the within-spell composition of (C, C) versus (N, C) and (C, N).[14]

Proposition 4. *In the private model ($q = 0$), for any $k \in (0, 1 - \mu_b]$ with $s > 0$ in Theorem 1, stationary expected welfare among normals strictly increases relative to the no-message benchmark:*

$$\bar{W}_{\mathrm{NN}}^{sig} \;>\; \bar{W}_{\mathrm{NN}}^{nosig}.$$

Proof. By Corollary 2, $\Lambda_{\mathrm{sig}} \le \Lambda_{\mathrm{nosig}}$ with strict inequality whenever $s > 0$, and by Proposition 2 the conditional expected spell length is identical across the two environments. Since $\bar{W}_{\mathrm{NN}}^{\mathrm{conf}} \le 0$ is the same under both (conflict spells are (C, C)), reducing Λ strictly raises $\bar{W}_{\mathrm{NN}}$. □

Proposition 5. *In the publicity extension ($q > 0$), the welfare ranking depends on the selected SPBE:*

(a) In the peace trap *of Theorem 2, $\Lambda = 0$ and $\bar{W}_{\mathrm{NN}} = 2$, which strictly dominates both the private model and the no-message benchmark.*

(b) In the conflict trap*, $\Lambda = 1$ and $\bar{W}_{\mathrm{NN}} = 0$, which is weakly dominated by both the private model and the no-message benchmark.*

(c) In the SPBE that coincides with Theorem 1 on the peace path (for any small $q > 0$), Λ is ambiguous*: onset falls, but the conditional expected duration among normals strictly increases on a nonempty region. Consequently, for sufficiently small q,*

$$\bar{W}_{\mathrm{NN}}^{public} \;\lesseqgtr\; \bar{W}_{\mathrm{NN}}^{private},$$

and each inequality can occur depending on parameters (the sign is governed by the balance between onset reduction and duration lengthening).

[13] These normalizations are without loss for our comparisons: bad types' choices are unaffected by messaging, so any constant welfare level assigned to those meetings cancels in differences across environments.

[14] Under our constructions, conflict spells among normals involve (C, C) with probability 1 on the private path and after $R = \mathsf{F}$ publicly; hence $\bar{W}_{\mathrm{NN}}^{\mathrm{conf}} = 0$ in our main comparisons.

Proof. (a) and (b) are immediate from $\Lambda \in \{0, 1\}$ under the absorbing records $R \in \{\mathsf{S}, \mathsf{F}\}$. For (c), by Proposition 1 the per-period onset hazard drops when messages are allowed; by Proposition 3 the conditional expected duration strictly rises on the stated region. A renewal decomposition writes $\Lambda = \lambda^{\text{on}} \cdot \mathbb{E}[D]/\mathbb{E}[\text{cycle length}]$; the denominator can be normalized across the two environments in our stationary comparison, so the sign of the welfare difference reduces to the sign of $\Delta\lambda^{\text{on}} \cdot \mathbb{E}[D_0] + \lambda^{\text{on}}_{\text{sig}} \cdot \Delta\mathbb{E}[D]$, which is ambiguous. □

Corollary 3. *When the public can coordinate on the* peace trap *(e.g., institutions select cooperative responses after a first leaked success), raising publicity q weakly increases welfare and strictly so until absorption at peace. Absent such selection, a planner who values shorter stalemates should prefer* private back-channels *(low q): they strictly reduce onset without lengthening conditional spells. Public pledges (high q) are preferable when either (i) audience institutions make $R = \mathsf{S}$ focal after one success, or (ii) onset reductions are large enough to dominate the duration effect.*

Remark 3. If we weight by population shares, total stationary welfare equals

$$\bar{W} \;=\; (1-\mu)^2\, \bar{W}_{\text{NN}} \;+\; 2\mu(1-\mu)\, W_{\text{NB}} \;+\; \mu^2\, W_{\text{BB}},$$

where the latter two terms are unchanged across our environments (bad types' behavior is invariant to messaging). Thus all welfare comparisons above lift to the full population by the prefactor $(1-\mu)^2$.

5.3 Cheap talk: impossibility under the baseline and a minimal fix

Could zero–cost messages substitute for costly reassurance? In our baseline, the answer is no: because all sender types prefer higher receiver cooperation, cheap talk cannot be informative.[15] We prove this impossibility for one– and two–sided finite protocols and then show how an arbitrarily small, single–crossing cost asymmetry restores partial informativeness in a disciplined way.

We now study a one-round, one-sided *cheap talk* variant of Section 3.4: the old can send $m \in \{D, H\}$ at zero cost; messages are private; actions follow simultaneously; types and payoffs are as in the baseline (bad types always play C).

5.3.1 Baseline impossibility

Let $q_D, q_H \in [0, 1]$ denote the normal receiver's cooperation probabilities (play N) after messages D and H, respectively. A normal old of private state $h \in \{g, b\}$ who sends message $m \in \{D, H\}$ obtains expected *gross* value $(1-\mu_h)V(q_m)$, where $V(q) = \max\{q(1+\ell) - \ell,\ qg\}$ is increasing in q. A bad old's expected value from m equals $(1-\mu)\, g\, q_m$, also increasing in q. Hence *all sender types strictly prefer higher q_m*, so their *message* preferences are perfectly aligned.

[15] All sender types strictly prefer transcripts with higher receiver cooperation, so zero-cost messages cannot separate. An ϵ-cost single-crossing perturbation restores partial informativeness (Section 5.3).

Proposition 6. *In the one-round, one-sided private cheap talk game with baseline primitives (bad types always play* C*; no message costs), every stationary PBE is* babbling*: messages are uninformative and the receiver's* (q_D, q_H) *does not depend on the distribution of sender messages. In particular, any equilibrium outcome coincides with the no-message benchmark for some* (q_D, q_H) *best responding to the prior.*

Proof. Since all sender types' payoffs are strictly increasing in q_m, any equilibrium in which $q_D \neq q_H$ makes all types strictly prefer the message with the larger q_m; thus no nondegenerate separation or partial separation is incentive compatible. Therefore only pooling (babbling) can be sustained, with (q_D, q_H) independent of the message distribution and chosen as the receiver's best response to the prior, as in the no-message benchmark. □

5.3.2 A minimal perturbation: ϵ-costly cheap talk with single-crossing

Introduce a *tiny, message-dependent* cost for sending D, denoted c_h for a normal old with private state $h \in \{g, b\}$ and c_b^{bad} for a bad old. Assume *single-crossing*:

$$0 \;\le\; c_g \;<\; c_b \;\le\; c_b^{\text{bad}}, \qquad \text{with all } c\text{'s arbitrarily small.}$$

Intuitively, "dovish" messaging is slightly less costly for truly reassuring states ($h = g$) and *most* costly for bad senders. The receiver still maps messages to (q_D, q_H).

Proposition 7. *Fix any small costs satisfying* $0 \le c_g < c_b \le c_b^{\text{bad}}$*. There exists a stationary PBE with* partial separation *in which*

normal old with $h = g$ *sends* D*,* *normal old with* $h = b$ *sends* H*,* *bad old sends* H*,*

and the receiver plays

$$q_H \;=\; 0, \qquad q_D \;=\; \min\left\{ \bar{q}, \;\; \frac{c_b}{(1-\mu_b)\,g} \right\}.$$

In particular, if $c_b \le (1-\mu_b)\,g\,\bar{q}$*, then* $q_D = c_b/[(1-\mu_b)g] \in (0, \bar{q}]$ *and the receiver sometimes cooperates only after* D*, making* D *informative.*

Proof. Receiver: given the sender strategy above, after H the posterior weight on types that would ever choose N is zero, so $q_H = 0$ is optimal; after D, pick $q_D \in (0, \bar{q}]$ so the sender (if normal) would best-respond by C, keeping $V(q_D) = gq_D$ and receiver indifferent at the cutoff posterior.

Sender IC: - For $h = b$ normal, the gain from deviating $H \to D$ equals $(1-\mu_b)gq_D - c_b \le 0$ by choice of q_D. - For the bad old, the gain equals $(1-\mu)gq_D - c_b^{\text{bad}} \le (1-\mu_b)gq_D - c_b \le 0$ using $c_b^{\text{bad}} \ge c_b$ and $\mu \ge \mu_b$. - For $h = g$ normal, the gain from $H \to D$ equals $(1-\mu_g)gq_D - c_g \ge (1-\mu_b)gq_D - c_b \ge 0$ since $\mu_g < \mu_b$ and $c_g < c_b$, so D is (weakly) preferred. Off-path beliefs are not needed as both messages occur. □

Remark 4. Allowing a short *second* cheap-talk message by the receiver (e.g., a binary "ready to cooperate/not") creates further scope for informativeness even with $c_g = c_b = 0$: the receiver can threaten $q = 0$ unless the sender first sends H and the receiver then confirms readiness, generating the misaligned incentives needed for partial revelation (Baliga and Sjöström, 2004).

6 Conclusions

This paper integrates pre–play communication into an overlapping–generations model of the security dilemma with persistent heterogeneity and thin, private memory. In the baseline with private messages, we characterize a stationary mixed–signaling equilibrium in which normal olds send a costly reassurance only after alarming private histories, while never signaling after benign ones. The construction keeps receiver beliefs exactly at a cutoff via endogenous mimicking by bad types and pins the sender's mixing through a simple cost–indifference condition. Allowing such reassurance strictly reduces the hazard that conflict *starts* among normal–normal meetings, while conditional *duration* is unchanged in the private model. With a small probability of publicity (leaks), costly signals create a public record and the environment bifurcates: play generically absorbs in either a peace trap or a conflict trap, and even outside absorption leaked failures harden receivers and lengthen stalemates on a nonempty region. These results yield clear welfare comparisons and simple policy guidance about when to prefer private back–channels over public pledges.

The mechanism is parsimonious. Costly reassurance is used precisely where it matters most—after privately alarming histories—yet remains only marginally persuasive because bad types mimic just enough to keep the receiver at indifference. The sender's mixing equates the expected persuasion gain with the cost, generating explicit formulas for the receiver's cooperation after a signal and a sharp switch in the sender's post–signal action at a single threshold in the cost parameter. Publicity changes the nature of evidence: a leaked success coordinates on cooperation and locks in peace, whereas a leaked failure coordinates on distrust and lengthens the conflict path.

The policy implications are immediate. When audiences or institutions can coordinate on forgiveness after a single leaked success, making assurances *public* is desirable: it accelerates absorption into peace. Where such coordination is unlikely or polarizing, *private* back–channels dominate: they reduce the onset of conflict without lengthening stalemates. Our framework clarifies why back–channels and public pledges can perform so differently even when they transmit the same underlying intent.

Several limitations point to tractable extensions. First, richer memory and public learning (e.g., multi–bit or serially correlated private histories; finite–lived public records) can be accommodated and preserve our qualitative results while generating new comparative statics. Second, heterogeneity in signaling costs across groups or over time could proxy for domestic constraints or leadership turnover, allowing us to study asymmetric reassurance and credibility. Third, two–sided communication may matter when receivers can credibly pre–commit to reciprocation; our zero–cost result shows why pure cheap talk babbles, while an ϵ–cost single–crossing perturbation restores partial informativeness in a disciplined way. Fourth, mediators or institutions that screen messages before they become public could endogenize the leak probability and shape the basins of attraction of peace and conflict traps.

The model delivers testable predictions. Back–channel reassurance should reduce the incidence of *new* conflict spells (onset) across otherwise similar dyads. Publicized *failed* overtures should be followed by longer conditional durations and slower de–escalation, while publicized *successful* overtures should be followed by durable peace. Event–study designs around documented leaks, or variation in media openness or audience polarization (proxying for effective q), offer empirical leverage. Historical episodes featuring parallel private and public tracks provide qualitative cross–checks.

Taken together, our analysis disentangles how communication shapes conflict on the two margins that matter—onset and duration—and explains when publicity transforms reassurance from a private persuasion device into a public coordination device. The framework is deliberately simple but flexible enough to incorporate richer memory, institutions, and asymmetric environments. We hope it clarifies why some peace talks avert war while others, when exposed, can entrench it—and how to design communication protocols that favor the former.

A Proofs for Section 4

A.1 Proof of Lemma 2

Fix $\bar{p} \in (0,1)$, $\mu \in (0,1)$, and (π_b, π_g) with $\pi_b > 0$, $\pi_g = 1 - \pi_b$. Let $s \in (0,1]$ be the probability that a *normal* old with $h = b$ sends $m = 1$ (normals with $h = g$ send $m = 0$), and let $p_b \in [0,1]$ be the probability a *bad* old sends $m = 1$. By Bayes' rule,

$$\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \frac{(1-\mu)\,\pi_b\, s}{(1-\mu)\,\pi_b\, s + \mu\, p_b},$$
$$\mathbb{P}(\theta = \mathrm{n} \mid m = 0) = \frac{(1-\mu)\,[\pi_g + \pi_b(1-s)]}{(1-\mu)\,[\pi_g + \pi_b(1-s)] + \mu(1-p_b)}.$$

Step 1 (hit the cutoff after $m = 1$). Define

$$\kappa \;\equiv\; \frac{1-\bar{p}}{\bar{p}} \cdot \frac{1-\mu}{\mu} \cdot \pi_b \in (0, \infty).$$

Setting $p_b = \kappa s$ yields $\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \bar{p}$ for any $s > 0$ (direct algebra). Validity requires $p_b \le 1$, i.e. $s \le s^{\max} := 1/\kappa$.

Step 2 (push $m = 0$ to the conflict side). With $p_b = \kappa s$,

$$\mathbb{P}(\theta = \mathrm{n} \mid m = 0) \le \bar{p} \iff (1-\mu)\,[\pi_g + \pi_b(1-s)] \;\le\; \frac{\bar{p}}{1-\bar{p}}\,\mu.$$

Rearranging gives

$$s \;\ge\; s^{\min} \;\equiv\; 1 - \frac{1}{\pi_b}\left[\frac{\bar{p}\,\mu}{(1-\bar{p})(1-\mu)} - \pi_g\right].$$

Hence any $s \in [s^{\min}, s^{\max}]$ with $s^{\min} < s^{\max}$ and $s^{\max} \le 1$ is feasible, and choosing $s \in (s^{\min}, \min\{1, s^{\max}\}]$ completes the construction.

Step 3 (nonemptiness of the interval). Since $\pi_b > 0$ and $(\mu, \bar{p})$ are interior, there is a nonempty parameter region with $s^{\min} < \min\{1, s^{\max}\}$.[16] This proves the lemma. □

[16] A simple sufficient condition is $\pi_b > \frac{\bar{p}\,\mu}{(1-\bar{p})(1-\mu)} - \pi_g$, which holds for all interior $(\mu, \bar{p})$ if π_b is not too small. If this fails at knife-edge values, one can allow *vanishing* mixing by normals with $h = g$ (i.e., send $m = 1$ with tiny probability $\eta > 0$) to enlarge the feasible set; all results are robust to this $\eta \downarrow 0$.

A.2 Proof of Theorem 1

Fix $\Theta = (\mu, \varepsilon, k, g, \ell)$ and (μ_g, μ_b) with $0 < \mu_g < \mu_b < 1$. We construct a stationary PBE.

Strategies.

- *Messaging.* A normal old with $h = g$ sends $m = 0$. A normal old with $h = b$ mixes: sends $m = 1$ with $s \in (0, 1)$ (to be chosen below). A bad old mixes with probability $p_b = \kappa s$ where κ is as in Lemma 2.

- *Receiver actions.* After $m = 1$, a normal receiver plays N with probability q_1; after $m = 0$, plays N with probability $q_0 = 0$. We set

$$q_1 \;=\; \begin{cases} \frac{k}{(1-\mu_b)g}, & \text{if } k \in \big(0, (1-\mu_b)k_\star\big], \\ \frac{k/(1-\mu_b)+\ell}{1+\ell}, & \text{if } k \in \big[(1-\mu_b)k_\star,\, 1-\mu_b\big], \end{cases} \qquad \text{where} \quad k_\star = \frac{g\ell}{1+\ell-g},\ \bar{q} = \frac{\ell}{1+\ell-g}.$$

- *Sender actions.* A normal old best-responds at the action stage: if $q_i \geq \bar{q}$, play N; else play C. A bad receiver always plays C; bad senders always play C.

Beliefs. Receivers use Bayes' rule on-path: choose $s \in (s^{\min}, \min\{1, s^{\max}\}]$ from Lemma 2 and $p_b = \kappa s$, so that

$$\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \bar{p}, \qquad \mathbb{P}(\theta = \mathrm{n} \mid m = 0) \leq \bar{p}.$$

Off-path beliefs are irrelevant because $m \in \{0, 1\}$ both occur with positive probability in equilibrium.

Sequential rationality (receivers). By Lemma 1, when $\mathbb{P}(\theta = \mathrm{n} \mid m = 1) = \bar{p}$ the normal receiver is indifferent and any $q_1 \in [0, 1]$ is optimal; when $\mathbb{P}(\theta = \mathrm{n} \mid m = 0) \leq \bar{p}$, any $q_0 \in [0, \bar{q}]$ (in particular $q_0 = 0$) is optimal.

Sequential rationality (senders: action stage). By construction,

$$V(q) \equiv \max\{\, q(1+\ell) - \ell,\ qg\,\} = \begin{cases} qg, & q \leq \bar{q}, \\ q(1+\ell) - \ell, & q \geq \bar{q}. \end{cases}$$

Thus after $m = 0$ and $q_0 = 0$, a normal sender plays C. After $m = 1$, in the low-cost range $q_1 \leq \bar{q}$ so the sender plays C; in the high-cost range $q_1 \geq \bar{q}$ so the sender plays N.

Sequential rationality (senders: messaging stage). Fix $h = b$. The sender's expected *gross* gain from sending (vs. not sending) equals $(1 - \mu_b)[V(q_1) - V(q_0)]$, since against a bad receiver the payoff is 0 regardless. Our choice of q_1 makes

$$(1 - \mu_b)\, V(q_1) \;=\; k, \qquad V(q_0) = 0,$$

hence the $h = b$ normal is indifferent and mixes as prescribed. For $h = g$, the same calculation gives the gross gain $(1 - \mu_g)\, V(q_1) < (1 - \mu_b)\, V(q_1) = k$ (since $\mu_g < \mu_b$ and $V(q_1) \geq 0$), so the net gain is negative and the $h = g$ normal strictly prefers not sending.

Feasibility of (s, p_b). Choose $s \in (s^{\min}, \min\{1, s^{\max}\}]$ with $s^{\max} = 1/\kappa$ from Lemma 2; this guarantees $p_b = \kappa s \in (0, 1]$ and $\mathbb{P}(\theta = \mathrm{n} \mid m = 0) \le \bar{p}$. Since both messages occur with positive probability, Bayes' rule applies on-path. Stationarity holds by construction. This completes the proof. □

A.3 Proof of Proposition 1

Work with the conditional information set: both players are normal types and the old's private state is $h = b$.

No-signaling benchmark. Without messages, the receiver's belief that the old will play N is below $\bar{p}$ under $h = b$; hence the old's best response is C and the outcome is conflict with probability 1. Thus $\lambda^{\mathrm{on}}_{\mathrm{nosig}} = 1$.

With signaling. In the SPBE of Theorem 1, the $h = b$ normal sends $m = 1$ with probability $s \in (0, 1)$, otherwise $m = 0$. After $m = 0$ the constructed posterior is at/below $\bar{p}$ and the outcome is conflict with probability 1. After $m = 1$, the receiver plays N with probability $q_1 > 0$ (indifference at the cutoff), and when $k > (1 - \mu_b)k_\star$ the sender best-responds by N. Therefore

$$\lambda^{\mathrm{on}}_{\mathrm{sig}} = (1 - s) \cdot 1 + s \cdot \mathbb{P}(\text{conflict} \mid m = 1) \; < \; 1$$

whenever $s > 0$ and either the sender plays N with positive probability after $m = 1$ (high-cost regime) or the receiver does (low-cost regime with $q_1 > 0$). Weak inequality holds in general. □

A.4 Proof of Proposition 2

Condition on a conflict spell that has just started among normal types. In the no-message SPBE, the old plays C and the young observes C with probability $1 - \varepsilon$ (and N with probability ε). Hence the next period's old (today's young) has $h = b$ with probability $1 - \varepsilon$ and $h = g$ with probability ε; the spell ends precisely when the young misperceives C as N, an event of probability ε each period, independently across periods. Therefore the spell length D is geometric with parameter ε and $\mathbb{E}[D] = 1/\varepsilon$.

In the signaling SPBE from Theorem 1, any spell that starts among normals must occur in the low-cost regime $k \le (1 - \mu_b)k_\star$, where the sender plays C regardless of m. Thus the same transition applies and the same geometric argument yields $\mathbb{E}[D] = 1/\varepsilon$. □

A.5 Proof of Theorem 2

We construct two stationary PBE candidates in the publicity extension (leak probability $q \in (0, 1]$) and verify sequential rationality and belief consistency.

Public state and updates. The public record $R \in \{\emptyset, \mathsf{S}, \mathsf{F}\}$ evolves as in Section 5.1: if a costly message leaks, the period outcome is recorded as S if (N, N) and F otherwise; absent a leak, R persists. Agents observe R at the start of the period.

Peace-trap equilibrium (absorption at S). Define strategies as follows.

- If $R = \emptyset$: normal old *with any* h sends $m = 1$ with probability 1; receiver plays as in Theorem 1 but set $q_1(R = \emptyset) = 1$, $q_0(R = \emptyset) = 0$; the normal sender best-responds by N after $m = 1$.
- If $R = \mathsf{S}$: no one sends (messages off), both play N.
- If $R = \mathsf{F}$: strategies not reached under this equilibrium (they can be set arbitrarily consistent with sequential rationality, e.g. as in the conflict trap below).

Incentives at $R = \emptyset$. The normal old's gain from sending is $(1-\mu_h)\cdot 1 - k$, where $\mu_h \in \{\mu_g, \mu_b\}$ is the posterior on the receiver being bad; this is $\geq 1 - \mu_b - k$. For any $k \leq 1 - \mu_b$ sending is optimal for both $h = g$ and $h = b$. (Against a bad receiver the payoff is 0 regardless.) The receiver's action $q_1 = 1$ is optimal given the posterior at/above the cutoff after $m = 1$; $q_0 = 0$ is optimal after $m = 0$. *Dynamics.* Because the old sends with probability 1 at $R = \emptyset$, a leak occurs with probability q each period; at the first leak, (N, N) is realized and R switches to S, after which both play N forever. Hence starting from $R = \emptyset$, $\{R = \mathsf{S}\}$ is reached with probability 1 for any $q > 0$ and is absorbing.

Conflict-trap equilibrium (absorption at F). Define strategies as follows.

- If $R = \mathsf{F}$: receivers set $q_1(R = \mathsf{F}) = q_0(R = \mathsf{F}) = 0$; normal senders best-respond by C at the action stage; messaging is off-path and can be set arbitrarily (e.g., no one sends).
- If $R = \emptyset$: replicate the cutoff construction of Theorem 1: normals with $h = b$ mix on sending with prob. $s \in (0, 1)$, $p_b = \kappa s$, posteriors at the cutoff after $m = 1$, and $q_1 \in (0, 1)$, $q_0 = 0$.

Incentives. At $R = \mathsf{F}$ the receiver's beliefs (supported by the public failure) justify $q_1 = q_0 = 0$; then the sender's best response is C, and messaging is useless (strictly dominated if $k > 0$). At $R = \emptyset$ the cutoff construction is as before. *Dynamics.* If a leak occurs at $R = \emptyset$ and is followed by conflict (which has positive probability under the cutoff construction), the state becomes F and remains there forever. Thus $\{R = \mathsf{F}\}$ is absorbing.

Bifurcation statement. The two candidates demonstrate coexistence of a peace-trap and a conflict-trap SPBE on the nonempty region $\{(\mu, \varepsilon, k, g, \ell) : k \leq 1 - \mu_b\}$ for *any* $q > 0$. Therefore any $q^\star \in (0, 1)$ suffices for the theorem's existence claim. □

Remark 5. The proof shows a stronger statement: for any $q > 0$, both absorbing equilibria exist on a nonempty parameter region (e.g., $k \leq 1 - \mu_b$). The $q^\star$ threshold is therefore weakly less than any given number in $(0, 1)$ on that region.

A.6 Proof of Proposition 3

Fix Θ and any $q > 0$. Consider the SPBE coinciding with Theorem 1 at $R = \emptyset$ so that a normal old with $h = b$ mixes on sending with probability $s > 0$, and suppose $k \leq (1 - \mu_b)k_\star$ so the sender best-responds by C after $m = 1$.

Start a conflict spell among normals. Under $q = 0$ (private model), Proposition 2 yields $\mathbb{E}[D_0] = 1/\varepsilon$. Under $q > 0$, in each period of the spell there is an event {leak of $m = 1$} with probability $sq > 0$. When this event occurs, the public record is updated to $R = \mathsf{F}$ (since the sender plays C in the low-cost regime and peace fails). Under $R = \mathsf{F}$, receivers adopt $q_1(R = \mathsf{F}) = q_0(R = \mathsf{F}) = 0$ (or, more generally, sufficiently low cooperation probabilities) so that future costly messages on the conflict path do not induce immediate exits. Thus, conditional on such a leak before the baseline geometric exit time, the spell length strictly exceeds the baseline.

Define two processes on a common probability space: (i) the baseline ($q = 0$) spell with stopping time T_0 (first C $\to$ N misperception), and (ii) the publicity ($q > 0$) spell with stopping time T_q, which equals T_0 unless a failure leak occurs first, in which case further exits are delayed by the $R = \mathsf{F}$ response. Then $\mathbb{P}(T_q > T_0) \geq sq(1 - \varepsilon) > 0$ and $T_q \geq T_0$ a.s. Hence $\mathbb{E}[T_q] > \mathbb{E}[T_0] = 1/\varepsilon$. □

B Robustness and Microfoundations

B.1 A. Microfounding the bad type via a Prisoner's Dilemma

Let the "bad" type's one-shot payoff against opponent action $a' \in \{\mathrm{N}, \mathrm{C}\}$ be a PD: $T > R > P > S$ with the usual ordering and R normalized to 0 for C–C comparison. Let N correspond to "cooperate" and C to "defect." Then for any belief over a', C strictly dominates N. Map parameters to our normalization by identifying the normal's table $\{1, -\ell;\, g, 0\}$ with $R = 1$, $S = -\ell$, $T = g$, $P = 0$. The bad type's payoff need not equal the normal's; only the dominance of C is required.

Lemma 3. *If a type has PD payoffs with* C *strictly dominant, then in any SPBE of the extended game the type plays* C *with probability* 1 *at the action stage, independent of messages and beliefs. Hence the behavioral "bad = always* C*" assumption is microfounded.*

Proof. Immediate from strict dominance in the stage game; messaging cannot alter best replies at the action node. □

B.2 Alternative monitoring and memory

B.2.1 Asymmetric misperception

Let false positive and false negative rates be $(\varepsilon_{\mathrm{N}\to\mathrm{C}}, \varepsilon_{\mathrm{C}\to\mathrm{N}}) \in [0, 1)^2$ with $\varepsilon_{\mathrm{N}\to\mathrm{C}} \neq \varepsilon_{\mathrm{C}\to\mathrm{N}}$. Define the private state by $h = b$ iff the young *perceives* C, i.e., $h = b$ arises after true (C) with probability $1 - \varepsilon_{\mathrm{C}\to\mathrm{N}}$ and after true (N) with probability $\varepsilon_{\mathrm{N}\to\mathrm{C}}$.

Proposition 8. *All baseline results (Theorem 1, Propositions 1–2, Theorem 2, Proposition 3) continue to hold if we replace (μ_g, μ_b) by the induced belief pair under $(\varepsilon_{\mathrm{N}\to\mathrm{C}}, \varepsilon_{\mathrm{C}\to\mathrm{N}})$, provided $\mu_b > \mu_g$. The duration result becomes geometric with parameter $\varepsilon_{\mathrm{C}\to\mathrm{N}}$ (spells end when* C *is misperceived as* N*).*

Proof. Our construction uses only (i) $\mu_b > \mu_g$ and (ii) independence of private histories across generations; both survive under asymmetry. In duration arguments, termination requires misperceiving C as N, hence parameter $\varepsilon_{\mathrm{C}\to\mathrm{N}}$. □

B.2.2 Multi-bit or serially correlated memory

Let an agent carry a K-bit private summary of the last K perceived actions (or a Markov summary with serial correlation across perceptions). Suppose the summary induces a total preorder of "alarm" levels with a highest-alarm state $h^{\max}$ and a lowest $h^{\min}$.

Proposition 9. *If the induced posteriors satisfy $\mu(h^{\max}) > \mu(h) > \mu(h^{\min})$ for any interior state h, then there exists a stationary PBE with the same qualitative structure as Theorem 1: the normal old mixes on sending only on a nonempty subset of high-alarm states (including $h^{\max}$), and never sends at low alarm (including $h^{\min}$). Comparative statics in k retain the same piecewise-linear $q_1(k)$ mapping.*

Proof sketch. Collapse the state space to $\{\text{low}, \text{high}\}$ by thresholding at an alarm level $\hat{h}$; apply Lemma 2 to the induced $(\pi_{\text{high}}, \pi_{\text{low}})$ and beliefs. Refinements preserve the construction. □

B.3 Uniqueness within the stationary cutoff class

Define the *cutoff class* of stationary PBEs as those where (i) $\mathbb{P}(\theta = \text{n} \mid m = 1) = \bar{p}$ and $\mathbb{P}(\theta = \text{n} \mid m = 0) \leq \bar{p}$; (ii) the receiver uses (q_1, q_0) with $q_0 \leq \bar{q}$; (iii) the sender's action is a best reply given (q_1, q_0).

Proposition 10. *Fix (μ_g, μ_b) and $k \in (0, 1 - \mu_b]$. Within the cutoff class, the receiver's response $q_1(k)$ that yields sender indifference is unique and given by*

$$q_1(k) = \begin{cases} \dfrac{k}{(1-\mu_b)g}, & \text{if } k \leq (1-\mu_b)k_\star, \\ \dfrac{k/(1-\mu_b) + \ell}{1+\ell}, & \text{if } k \geq (1-\mu_b)k_\star, \end{cases} \qquad \text{with } q_0 = 0.$$

Consequently, the sender's post-signal action is uniquely determined in each regime.

Proof. Indifference requires $(1 - \mu_b)V(q_1) = k$. Since $V(\cdot)$ is strictly convex, piecewise linear with a unique kink at $\bar{q}$, q_1 is uniquely pinned below or above $\bar{q}$ accordingly. Setting $q_0 = 0$ maximizes the sender's desire to send; any $q_0 > 0$ would violate indifference unless q_1 is reduced, which would contradict receiver optimality at the cutoff. □

B.4 Off-path beliefs and message support

In the private baseline both messages occur with positive probability in equilibrium. If one wishes to shrink support (e.g., set $s = 1$ at $h = b$ or $p_b = 0$ at bounds), adopt the following beliefs: after an off-path $m = 1$ at histories where it should not occur (e.g., $h = g$), let the receiver place posterior $\mathbb{P}(\theta = \text{n} | m = 1) = \bar{p}$ (tie-breaking in favor of indifference) and choose q_1 as in Theorem 1; after an off-path $m = 0$ where it should not occur, let the receiver infer $\mathbb{P}(\theta = \text{n} | m = 0) = 0$ and best-respond by C. These beliefs satisfy Bayes consistency wherever applicable and maintain sequential rationality.

In the publicity model, if a costly message is observed at $R = \mathsf{F}$ off-path in the conflict-trap equilibrium, assign $\mathbb{P}(\theta = \text{n} | m = 1, R = \mathsf{F}) = 0$ so that $q_1(R = \mathsf{F}) = 0$ remains optimal; symmetric beliefs support the peace trap.

B.5 Finite-lived public record (leaks with forgetting)

Modify the publicity extension so that after each period the public record $R \in \{\emptyset, \mathsf{S}, \mathsf{F}\}$ *forgets* with probability $\rho \in (0,1)$, reverting to $\emptyset$. Leaks occur with probability $q \in (0,1]$ as before.

Proposition 11. *For any $q > 0$ and $\rho \in (0,1)$ there exist stationary PBEs that mirror the peace- and conflict-trap constructions: $R = \mathsf{S}$ and $R = \mathsf{F}$ become* quasi-absorbing *with expected sojourn times of order $(1-\rho)^{-1}$. In each equilibrium the long-run fraction of periods in peace (respectively conflict) can be made arbitrarily close to 1 by taking ρ small enough.*

Proof sketch. Replicate the trap strategies. With probability $1-\rho$ the record persists, so the chain spends geometric-length blocks in $\{\mathsf{S}\}$ or $\{\mathsf{F}\}$. Upon reversion to $\emptyset$, the next leak reselects the basin according to the original construction. Stationary distributions place mass $1 - O(\rho)$ on the selected trap. □

Remark 6. The duration result in Proposition 3 extends: leaked failures still harden receivers on the conflict path, raising conditional duration by at least a constant that does not vanish as $\rho \downarrow 0$.

C Two-Sided, Zero-Cost Cheap Talk

We allow pre-play *cheap talk* with zero message costs and private messages. The sender is the old, the receiver is the young. Bad types still play C at the action stage.

C.1 Protocols

One-round two-sided talk. The sender first sends $m_1 \in \{D, H\}$; the receiver then sends $m_2 \in \{r, n\}$; finally both choose actions simultaneously. Messages are costless and non-binding; there is no public record here ($q = 0$).

Multi-round. A finite number $T \geq 2$ of alternating zero-cost messages is allowed before the action stage. Strategies and beliefs can condition on the full private transcript.

In either case, let $q(m_1, m_2, \ldots)$ denote the receiver's (normal type) cooperation probability Pr(N) as a function of the transcript; the sender's (normal) best action after a transcript with cooperation probability q is N iff $q \geq \bar{q}$, else C, where $\bar{q} = \frac{\ell}{1+\ell-g}$ as in the main text.

C.2 Impossibility of informative equilibria

The key observation is that *all* sender types—bad, normal with $h = g$, and normal with $h = b$—strictly prefer transcripts that induce *higher* receiver cooperation. Formally, for a normal sender in state $h \in \{g, b\}$ the gross value after a transcript inducing q is $(1-\mu_h)\, V(q)$ with $V(q) = \max\{q(1+\ell) - \ell,\ qg\}$, strictly increasing in q on $[0,1]$; for a bad sender it is $(1-\mu)\, g\, q$, also strictly increasing in q.

Proposition 12. *In the two-sided, zero-cost cheap-talk protocols above (one round or any finite number of rounds), every stationary PBE is babbling: the receiver's action distribution* *depends only on the prior (and equilibrium mixing), not on the realized messages. In particular, no partially separating equilibrium exists.*

Proof. Suppose, towards a contradiction, that some equilibrium induces different cooperation probabilities across on-path transcripts, say $q(\tau_1) \neq q(\tau_2)$. Since every sender type's payoff is strictly increasing in q, every type strictly prefers the transcript with the larger q. Then no sender type can be willing to follow a strategy that assigns positive probability to the transcript with the smaller q, contradicting on-pathness unless $q(\tau_1) = q(\tau_2)$. Iterating this argument over all on-path transcripts implies $q(\tau)$ is constant on-path. Given zero costs and no public record, off-path beliefs cannot overturn this alignment. Hence messages are uninformative, and the receiver best responds to the prior, as in the no-message benchmark. □

Lemma 4. *If preferences over transcripts are strictly monotone in the receiver's cooperation probability for all sender types (as here), then allowing any finite number of additional zero-cost message rounds cannot restore informativeness: Proposition 12 holds for all finite* $T \geq 1$.

Proof. Adding rounds only enlarges the set of transcripts; since all types weakly prefer higher-q transcripts, any profile that attempts to create distinct posteriors across transcripts induces profitable deviations toward the highest-q transcript. The same contradiction as in the proof of Proposition 12 applies. □

References

Acemoglu, D. and A. Wolitzky (2014). Cycles of conflict: An economic model. *American Economic Review 104*(4), 1350–1367.

Acemoglu, D. and A. Wolitzky (2024). Mistrust, misperception, and misunderstanding: Imperfect information and conflict dynamics. In *Handbook of the Economics of Conflict*, Volume 1. Elsevier. Chapter 2.

Baliga, S. and T. Sjöström (2004). Arms races and negotiations. *Review of Economic Studies 71*(2), 351–369.

Crawford, V. P. and J. Sobel (1982). Strategic information transmission. *Econometrica 50*(6), 1431–1451.

Fearon, J. D. (1994). Domestic political audiences and the escalation of international disputes. *American Political Science Review 88*(3), 577–592.

Fearon, J. D. (1995). Rationalist explanations for war. *International Organization 49*(3), 379–414.

Katagiri, Y. and B. W. Min (2019). The credibility of public and private signals: A document-based approach. *American Political Science Review 113*(1), 156–172.

Kurizaki, S. (2007). Efficient secrecy: Public versus private threats in crisis diplomacy. *American Political Science Review 101*(3), 543–558.

Kydd, A. H. (2000). Trust, reassurance, and cooperation. *International Organization 54*(2), 325–357.

Kydd, A. H. (2003). Which side are you on? bias, credibility, and mediation. *American Journal of Political Science 47*(4), 597–611.

Lagunoff, R. (2001). Organizations and overlapping generations games: Memory, communication, and altruism. Earlier version of Lagunoff and Matsui (2004), *Review of Economic Design*.

Leventoglu, B. and A. Tarar (2008). Does private information lead to delay or war in crisis bargaining? *International Studies Quarterly 52*(3), 533–553.

Quek, K. (2021). Four costly signaling mechanisms. *American Political Science Review 115*(2), 537–549.

Schultz, K. A. (2001). Looking for audience costs. *Journal of Conflict Resolution 45*(1), 32–60.

Tarar, A. and B. Leventoglu (2009). Public commitment in crisis bargaining. *International Studies Quarterly 53*(3), 817–839.

Trager, R. F. (2011). Multi-dimensional diplomacy. *International Organization 65*(3), 469–506.