Assurance in Coercive Bargaining

Shusuke Ioku[†] Shuhei Kurizaki[‡]

September 16, 2026

Abstract. Political actors can fail to secure concessions despite credible threats when targets fear that yielding could invite further harm. We show that this assurance problem can complicate signaling through coercive demands, making it harder for targets to distinguish restraint from plans to attack. Our bargaining model allows concessions to improve the demanding state's military position before it decides whether to attack. This gives states planning an attack a reason to imitate the modest demands of states seeking peace. Greater military strength can make reassuring demand patterns harder to sustain and permit additional ways to conceal an intended attack. These changes can reduce the concessions the demanding state obtains peacefully and increase the risk of war, imposing collective costs when disputes that could have been settled peacefully instead lead to fighting. Lack of assurance can exacerbate information problems in interstate politics, undermining coercive bargaining and making peace harder to secure.

[†]Ph.D. Candidate, Department of Political Science, University of Rochester.

[‡]Associate Professor, Faculty of Political Science and Economics, Waseda University.

Introduction

Political actors use threats to obtain concessions, but targets may resist even credible threats when compliance does not guarantee safety (Schelling 1966). During the Iran nuclear negotiations, Iranian officials feared that an agreement would leave sanctions in place or invite new demands. The Obama administration sought to address these concerns by separating nuclear sanctions from other disputes and managing actors who might obstruct an agreement. Iran accepted the agreement in 2015. In 2018, the Trump administration withdrew and reimposed sanctions to seek further concessions, and Iran subsequently defied the demands and expanded uranium enrichment (Pauly 2024). These events illustrate the assurance problem: persuading the target that compliance will end the pressure. We examine how this assurance problem shapes the demands states make and what targets can learn from them.

We develop a parsimonious model of conflict by relaxing the standard bargaining assumption that compliance ends the crisis. The demanding state knows its own cost of war but the target does not. It makes a demand that the target accepts or resists. If the target resists, the state can fight or back down. If the target complies, the state can keep the settlement or attack, and the concession improves its chance of victory. We compare this setting with an otherwise identical model in which compliance guarantees peace.

Our main result is that the assurance problem itself can make demands harder to interpret, leaving the target less able to distinguish restraint from plans to attack. When concessions improve military prospects, a state planning an attack can imitate one seeking a peaceful settlement to obtain compliance. A modest demand can therefore come from a state that wants to avoid fighting or from one that wants a concession before fighting. In one such pattern, a larger demand can be more reassuring because the state making it is willing to stop if the target complies. Other patterns allow a middle demand to conceal plans to attack even when both smaller and larger demands assure peace.

These changes in signaling also restrict what bargaining can achieve. Allowing demands and responses to adjust cannot recover the best outcomes for peaceful concessions and welfare available when compliance guarantees peace, or its lowest war risk. The losses can arise in different ways. A state planning an attack can gain by choosing a modest demand that is more likely to be accepted and then attacking after receiving the concession. The target can therefore concede more often while the demanding state obtains less through peaceful settlement. If fear of attack instead leads the target to concede less, the demanding state can lose as a peaceful settlement becomes less valuable. Smaller concessions also make war more attractive to states that would otherwise stop.

We contribute to three literatures. First, we provide a parsimonious framework that isolates how assurance affects coercive bargaining. The assurance literature has developed largely through qualitative arguments and case studies. Pauly (2024) explains how separating disputes and managing other actors can help states honor their assurances. Cebul, Dafoe, and Monteiro (2021) use a survey experiment to examine how reputations for restraint affect the perceived credibility of assurances. Following Paine and Tyson (2020), we adopt an experimental approach to formal modeling: we compare otherwise identical games that differ only in whether the coercer can attack after compliance. This framework allows us to trace the effects of lacking assurance through bargaining behavior to concessions, war, and welfare.

Second, we contribute to research on what demands reveal about intentions. Kydd and McManus (2017) show how an aggressive state can imitate the larger demand and assurance of a state willing to settle; the smallest demand remains reassuring in their example. Slantchev (2010) shows that strong states can imitate weak states' modest demands to preserve the element of surprise if refusal leads to war. Acceptance ends the crisis in his model. Here, deception seeks compliance itself: the concession improves the coercer's position for a subsequent attack. This motive can disrupt the ordering of demands as signals of restraint. A smaller demand can be less reassuring than a larger one, and a middle demand can conceal plans to attack even when smaller and larger demands assure peace.

Third, we contribute to research on the limits of military strength in coercive bargaining. Sechser (2010) studies signaling by a target whose war cost is private: resistance can build a reputation for resolve that induces smaller demands in future disputes. Our model reverses these roles, placing private war costs with the coercer and making its demand the signal. When compliance leaves the target exposed to attack, greater strength can enable additional patterns of deception that are unavailable at lower strength. Lack of assurance can also weaken the effect of military strength on coercive ability: increases in strength can yield smaller gains in the average peaceful concession obtained when the target complies.

The Framework

Two states bargain over a disputed good valued at 1. The demanding state, or *coercer*, privately observes its cost of war; the target knows only its distribution. We compare games that differ only in whether the coercer can attack after the target complies. Figure 1 previews the structure of the two games.

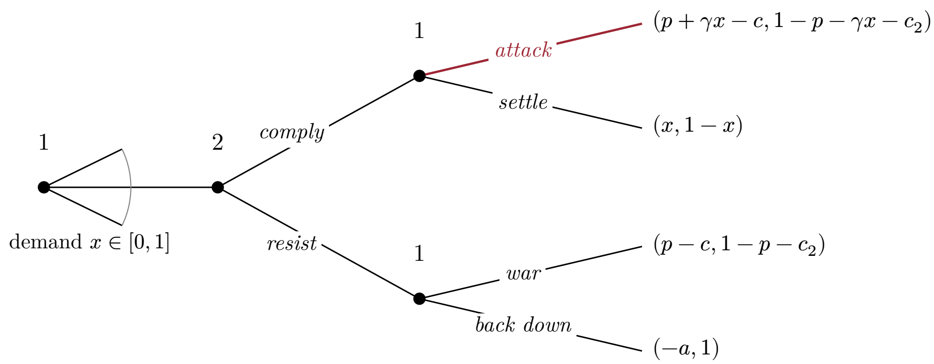

Figure 1: Coercive bargaining games with and without attack after compliance

Note: Player 1 is the coercer and player 2 the target; the coercer's payoff appears first. The tree begins after the coercer learns its war cost c. The target observes the demand but cannot observe this cost, so it cannot distinguish between its decision nodes following the same demand. Under protected compliance, yielding ends the crisis in settlement. Under exposed compliance, the red branch permits attack after yielding, and a concession x raises the coercer's probability of victory from p to $p + \gamma x$.

Actions and payoffs

The coercer chooses how much of the disputed good to demand. Write this demand as $x \in [0, 1]$. The target observes the demand and either complies, transferring that amount, or resists. If the target resists, the coercer chooses between war and backing down.

War is costly for both states. The coercer privately knows its cost, c, which is uniformly distributed on $[0, \overline{c}]$. This privately known cost defines the coercer's *type*. The target's cost, $c_2 \geq 0$, is known to both states. If the coercer's initial probability of victory is $p \in (0, 1)$, war after resistance gives the two states payoffs

$$(p - c,\, 1 - p - c_2)$$

Backing down imposes an *audience cost*, a political penalty for retreating from a threat, of $a > 0$. The target keeps the good, yielding $(-a, 1)$. A coercer with a sufficiently low war cost fights after resistance; one with a higher cost backs down. We call these types *resolved* and *weak*, respectively.

We compare two games that differ in what happens after compliance. Under the game with *protected compliance*, the transfer ends the crisis, with payoffs $(x, 1 - x)$. Under the game with *exposed compliance*, the coercer can keep this settlement or attack. Receiving the concession first improves its probability of victory by γx, so an attack gives

$$(p + \gamma x - c,\, 1 - p - \gamma x - c_2) \tag{1}$$

The parameter γ measures how much military advantage a concession provides. It can represent the effect of demobilization, surrendered terrain, or access to strategically valuable territory. At Munich in 1938, Czechoslovakia ceded territory containing fortifications that the agreement required it to leave intact.[1] An attack puts the entire disputed good at stake again. The concession improves the coercer's chance of winning but is not added separately to its war payoff.

[1] Munich Agreement, 29 September 1938, arts. 1–4; Wilson to the Secretary of State, 6 October 1938, FRUS 1938, I, doc. 698.

We treat this military advantage as irreversible within the game. Actual agreements can address this danger through phased implementation and military safeguards. The 1979 Egypt–Israel peace treaty paired phased Israeli withdrawal from Sinai with limits on military deployments and provision for outside monitoring.[2]

Only coercers with sufficiently low war costs prefer attacking to keeping the settlement. We call these types *belligerent* at the demand in question. They are willing to fight after resistance as well as after compliance. A resolved coercer need not be belligerent: it can fight if refused and settle if accepted. Attack after compliance carries no separate reputational penalty in the model. Costly promises and reputations for restraint provide other ways to alter this decision (Kydd and McManus 2017; Cebul, Dafoe, and Monteiro 2021).

We impose two sets of parameter restrictions. The first focuses the analysis on partial concessions while preserving the target's initial uncertainty about whether the coercer would fight if refused.

Assumption 1 (Interior bargaining): The parameters satisfy

$$p + c_2 < 1, \quad p + a < \overline{c},$$
$$\overline{c} > 2p + c_2, \quad (p + c_2)(p + a) < p\overline{c}.$$

The first condition makes the target prefer war after resistance to surrendering the entire good. The second ensures that both resolved and weak coercers are possible. The last two restrict bargaining over a single accepted demand to partial concessions. The third makes the target's expected gain from accepting this shared demand rather than resisting decrease as the amount demanded rises. The fourth keeps the target's expected loss from resistance before observing a demand below p. Together, they keep each game's largest such demand strictly between zero and p.

[2] Treaty of Peace, 26 March 1979, Annex I, arts. I–II and VI, reproduced by the American Presidency Project.

The second assumption allows concessions to improve military prospects while making settlement more attractive relative to attack. Each unit conceded adds one unit to settlement value but only $\gamma \in (0, 1)$ to the expected return from attack. Larger concessions can therefore encourage restraint even as they improve the coercer's military position. We also require victory to remain uncertain after the largest possible concession.

Assumption 2 (Concessions and military power): The military effect of concessions satisfies

$$0 < \gamma < 1, \quad p + \gamma < 1.$$

If attacking and keeping the concession give the coercer the same payoff, we assume it keeps the concession and stops. If fighting and backing down after resistance give it the same payoff, we assume it backs down. For any given demand, each equality holds at a single war-cost value, which has zero probability under the continuous distribution. As we show below, different demands can give entire groups of coercers the same payoff.

Kydd and McManus (2017) let the target choose how much to concede after observing the coercer's goal and any threat or assurance. We retain their link between concessions and military power, but let the coercer choose a demand that the target accepts or resists. Allowing the coercer's private war cost to vary continuously lets us identify which types change demands and which abandon settlement when the target concedes less.

Equilibrium and scope

We solve for perfect Bayesian equilibria satisfying the Intuitive Criterion (Cho and Kreps 1987). After observing a demand, the target revises its assessment of the coercer's war cost. Each state chooses the action that gives it the highest expected payoff at every stage, given the other state's behavior and the target's assessment. A coercer with a given war cost chooses one demand without randomizing across demands. The target may sometimes comply and sometimes resist only when these choices give it the same expected payoff. The

Intuitive Criterion limits what the target can infer from a demand that no coercer would make in equilibrium. Supporting Information (SI) A.2 specifies these restrictions and the beliefs that sustain each equilibrium.

We use *pooling* to cover equilibria with one demand that receives compliance, including partial pooling in which some coercers instead make demands that the target always resists. In *two-demand bargaining*, two demands for nonzero concessions each attract a group of coercers and are accepted at least sometimes. Each demand must be chosen with nonzero probability, so a demand chosen only at an isolated war-cost value does not qualify. We characterize every equilibrium with these two accepted demands, allowing their sizes and acceptance probabilities to vary. Here, *reassurance* means that compliance ends in settlement with certainty; it describes behavior rather than a separate message.

Protected compliance cannot sustain a third demand meeting these conditions (Proposition 2). Bargaining with one or two accepted demands therefore covers all demand counts that both games can sustain while producing concessions. We also characterize three-demand bargaining under exposed compliance and allow the number of accepted demands to vary when comparing the best outcomes each game can sustain.

Measuring the outcomes of coercive bargaining

Compliance can produce either settlement or an attack. Obtaining agreement therefore differs from obtaining a peaceful concession. We keep that distinction when measuring the outcomes of coercive bargaining.

> Definition 1 (Coercive bargaining outcomes): *Coercive ability* is the average peaceful concession per instance of compliance, counting compliance followed by attack as zero. The *expected peaceful transfer* averages peaceful concessions over all crises, including those in which the target resists. *War risk* is the probability

of fighting, whether after resistance or compliance. *Conflict cost* is the expected sum of both states' war costs. *Total welfare* is the sum of their expected payoffs.

The two concession measures answer different questions. Coercive ability describes what compliance delivers on average; the expected peaceful transfer also accounts for how often compliance occurs. When the target never complies, the peaceful transfer is zero and coercive ability is undefined. SI A.1 gives the formulas for these measures.

What Demands Reveal

We begin by examining what the size of a demand reveals about the coercer's willingness to stop after compliance. A small demand can reflect opposing motives. A weak coercer seeks a concession that the target is likely to accept because resistance would force it to back down. A belligerent coercer can also value acceptance when the concession improves its position before an attack. Coercers with intermediate war costs can instead prefer a larger demand, risking resistance but settling if the target complies.

In the protected compliance game, a low demand attracts coercers who place more value on avoiding resistance. Among coercers using the two accepted demands, those with lower war costs choose the larger one and those with higher costs choose the smaller one. We call this ordering *monotone sorting*. This ordering can persist in the exposed compliance game, even if some coercers making the higher demand attack after acceptance. Allowing attack also permits belligerent and weak types to choose the same low demand while coercers with intermediate war costs demand more. We call this *feigned restraint*. These are the only two patterns when each coercer prefers its chosen demand to the alternative, except at the war-cost values that divide the groups. We first establish when each pattern is possible and whether acceptance of either demand can guarantee settlement.

Proposition 1 (Lack of assurance changes what demands can reveal): Under Assumptions 1–2, both games allow pooling in which the target complies at least sometimes.

- In the *protected compliance game*, two demands can distinguish coercers with lower and higher war costs while both groups settle after compliance. Feigned restraint is impossible.
- In the *exposed compliance game*, two demands can distinguish coercers with lower and higher war costs while both groups settle after compliance if and only if $\gamma \leq \frac{c_2}{p+c_2}$. Feigned restraint is always possible.

When both demand groups settle after compliance, the larger demand is $p + c_2$. A higher γ makes the concession more valuable as preparation for an attack. A higher p permits a larger concession, which increases the military advantage obtained before an attack. Greater initial strength can therefore undermine settlement after compliance even when γ is unchanged. In Figure 2, equilibria in which both groups settle exist on and below the curve; feigned restraint remains possible on both sides.

Consider feigned restraint in which each coercer prefers its chosen demand to the other demand, except at the war-cost values that divide the groups. Belligerent types choose the low demand to obtain its higher chance of acceptance and then attack from a stronger position. Weak types choose the same demand to avoid resistance. Coercers with intermediate war costs make the high demand, fight if refused, and settle if accepted. A larger demand is therefore more reassuring in this pattern: acceptance ends the crisis, whereas acceptance of the low demand can precede attack.

The distinction between these motives also matters for what the target learns. Compare this pattern with an equilibrium in the protected compliance game, using the same demands and acceptance probabilities. In the exposed compliance game, belligerent types switch from the high demand to the low one. Their choice becomes indistinguishable from the choice of weak types, so demands reveal less about whether the coercer would fight after

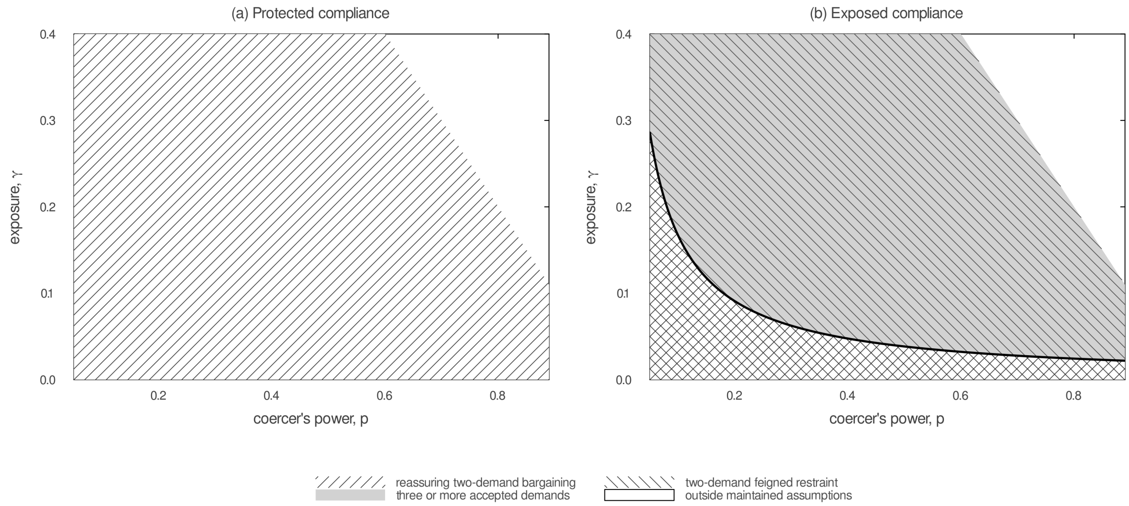

Figure 2: Lack of assurance permits additional demand signals and limits reassurance

Note: Hatching identifies parameter combinations that allow two different demands to receive compliance. In panel (a), two-demand equilibria in which all coercers settle after compliance exist wherever the assumptions hold; feigned restraint and bargaining with three or more accepted demands are impossible. In panel (b), feigned restraint is possible throughout this range of parameters; two-demand equilibria in which all coercers settle after compliance exist on and below the black curve. The gray area strictly above the curve also allows three or more demands to receive compliance. Each demand counted here asks for a nonzero concession, is chosen with nonzero probability, and is accepted at least sometimes. Proposition 2 establishes this existence region for three or more demands; SI D.3 gives the proof for larger finite demand counts. Pooling is possible wherever the assumptions hold. The white wedge violates $p + \gamma < 1$. Fixed parameters are $c_2 = 0.02$, $a = 0.30$, and $\overline{c} = 2.00$.

resistance. They do not become less informative about every aspect of war costs: the high demand now identifies coercers with intermediate costs who previously shared their demand with the most belligerent types. Lemma 4 in SI E.2 establishes this precise information comparison.

This pattern and its outcome comparisons do not depend on war costs being uniformly distributed (SI E.1). They remain possible under any continuous distribution that assigns some probability to every interval in the assumed cost range. A different distribution may require different demand sizes or acceptance probabilities. The pooling results and the outcome bounds below continue to assume uniformly distributed war costs.

Signaling through Additional Demands

The assurance problem can also change how many demands bargaining sustains. We next examine whether the target can face more than two accepted demands and what their relative sizes reveal about restraint.

Proposition 2 (Lack of assurance can complicate signaling through demands): Under Assumptions 1–2, suppose each coercer chooses one demand without randomizing. Consider equilibria with finitely many demands receiving compliance. Each such demand asks for a nonzero concession and is chosen with nonzero probability.

- In the *protected compliance game*, no equilibrium can use three or more such demands.
- In the *exposed compliance game*, an equilibrium with three or more such demands exists if and only if $\gamma > \frac{c_2}{p+c_2}$. When this condition holds, there are equilibria with exactly three demands in which only the middle demand carries a risk of attack after compliance: acceptance of either a smaller or a larger demand always ends in settlement.

The threshold for the military effect of concessions is the same as in Proposition 1.[3]

In the protected compliance game, every accepted demand above the smallest comes from coercers who would fight if refused and settle if accepted. A larger demand must be accepted less often; otherwise these coercers would all prefer it. To accept each higher demand only sometimes, the target must receive the same expected payoff from accepting it as from resisting. Resistance leads to war after all these demands, while acceptance ends the crisis. Only one concession amount can therefore make acceptance and resistance equally attractive. In the exposed compliance game, accepting a higher demand can instead lead

[3]SI D.1 proves that protected compliance permits at most two accepted demands. SI D.2 characterizes all three-demand equilibria under exposed compliance; SI D.3 establishes the condition for the finite class with three or more accepted demands. War costs remain uniformly distributed.

to either attack or settlement. The risk of attack can differ across demands, allowing the target to be willing to accept different amounts.

Consider a three-demand equilibrium in which each coercer prefers its chosen demand to both alternatives, except at the war-cost values dividing the groups. Belligerent coercers then use the middle demand to obtain a concession before attacking. They could imitate coercers making the largest demand, but its lower chance of acceptance more than offsets the greater military advantage of receiving it. Given their low war costs, seeking the largest demand and settling if accepted also gives them a lower payoff than making the middle demand and attacking after compliance. Weak coercers use the smallest demand to avoid resistance, while resolved coercers seeking settlement choose among all three. The smallest and largest demands therefore assure peace, while a demand between them can conceal a willingness to attack.[4]

The Consequences of Assurance

Comparing the consequences of assurance across the two games is challenging because neither game has a unique equilibrium. Different demand sizes, groups of coercers, and acceptance probabilities can all be consistent with the states' incentives.

We begin by comparing the outcomes each game can sustain while allowing bargaining terms to differ freely across games. This asks whether states in the exposed compliance game can achieve the favorable outcomes available in the protected compliance game, without specifying which equilibrium occurs. We then compare families of pooling and two-demand equilibria to trace the changes in behavior behind these limits.

[4]The full characterization also covers groups of attackers for whom the low and middle demands give the same payoff, or for whom the middle and highest demands do so. In the latter case, some attackers can choose the highest demand alongside coercers who settle. All three patterns become possible above the same threshold; see SI D.2.

What Bargaining Can Achieve

We allow demand sizes, the groups choosing each demand, acceptance probabilities, and the number of accepted demands to differ across games. We compare what each game can achieve in peaceful concessions, war, conflict costs, the coercer's expected payoff, and total welfare.

To make this comparison, we retain uniformly distributed war costs and require each coercer to choose one demand without randomizing. We allow any finite number of demands that the target accepts at least sometimes. Each must ask for a nonzero concession and be chosen with nonzero probability. Coercers who obtain the same payoff from several demands need not all choose the same one. The comparison therefore includes all three-demand equilibria characterized in SI D.2.

> PROPOSITION 3 (Lack of assurance restricts attainable outcomes): Under Assumptions 1–2 and the equilibrium conditions above:
>
> - *Peaceful concessions.* With positive compliance, both concession measures in the exposed compliance game remain strictly below their highest attainable values in the protected compliance game.
> - *War and conflict costs.* The lowest attainable war risk and expected conflict costs rise; their highest values are unchanged.
> - *Payoffs and welfare.* The highest attainable expected payoff for the coercer and total welfare fall; their lowest values are unchanged.

In each game, one pooling equilibrium achieves the lowest war risk and conflict costs and the highest coercer payoff and total welfare: the target always accepts the largest pooling demand. Compared with this equilibrium in the same game, every equilibrium with two or more accepted demands has higher war risk and conflict costs, a lower expected payoff for the coercer, and lower total welfare. Two-demand equilibria can nevertheless come arbitrarily close to these best outcomes.

To see why additional accepted demands cannot improve these outcomes, first suppose the target always accepts the smallest demand. Some coercers who would fight after resistance choose larger demands that risk refusal. War then includes some coercers with higher costs, whereas the best pooling equilibrium concentrates war among coercers with the lowest costs. For a given overall probability of war, spreading the fighting across coercers with higher costs raises its expected cost.

These extra costs also limit the concessions the target will accept. The target can always resist, so an accepted demand must give it at least the expected payoff from resistance. Together with the costs of fighting, this requirement keeps the smallest accepted demand below the largest pooling demand. Settlement is therefore less valuable to the coercer. If the target sometimes refuses even the smallest demand, resistance creates an additional route to war. Such equilibria also cannot improve on the best pooling outcomes.

The limit on peaceful concessions follows from the same requirement that the target be willing to accept. War reduces the target's payoff, leaving less room for peaceful transfers that it prefers to resistance. Because the exposed compliance game has a higher minimum war risk, no adjustment of demands and responses can recover the protected game's largest expected peaceful transfer. The peaceful concession obtained per instance of compliance also remains strictly below its protected maximum. SI F.3 proves both results; neither requires matched bargaining terms.

At the equilibria with the lowest war risk and highest total welfare, the coercer bears the welfare loss. The target receives the same expected payoff as it would by resisting every demand. That payoff is unchanged across games, so the additional expected war costs reduce the coercer's highest attainable expected payoff by the same amount. This is an average across coercers with different war costs; some coercers may still gain.

Both games also sustain equilibria in which the target resists every demand. These deliver the highest war risk and conflict costs and the lowest coercer payoff and total welfare. Between these common worst outcomes and each game's best outcomes, every intermediate value of each measure is attainable. The exposed compliance game therefore loses the most favorable part of each range. For the two concession measures, every exposed value is also attainable under protection, while a range of high protected values is unavailable under exposure. SI F gives the ranges and proofs.[5]

How Changes in Behavior Affect Bargaining Outcomes

Our objective in the focused comparisons is to identify how lack of assurance reduces peaceful concessions, increases war, and changes the two states' payoffs. The pooling comparison examines what happens when the target responds to the prospect of attack by accepting only a smaller concession. The comparison with two demands asks whether peaceful concessions can fall and war can rise even with the same demands and the same probabilities of acceptance. It isolates the consequences of belligerent coercers imitating restraint and attacking after acceptance.

For the *pooling comparison*, we choose each game's largest demand that can receive compliance in a pooling equilibrium. At a given probability of acceptance, this gives the coercer its highest expected payoff among pooling equilibria. Both games support every positive probability of accepting their respective demands, and we compare them at the same probability. This separates the effect of less valuable peace terms from changes in how often the target accepts them. Demand sizes can differ across games. In the protected compliance game, coercers who prefer war make demands that are always refused; the remaining coercers share the accepted demand. In the exposed compliance game, all coercers share the accepted demand. Belligerent coercers also prefer it because the concession improves their military position and they can still attack afterward.[6]

[5] Each range describes one measure at a time; it does not identify which combinations are jointly attainable. These comparisons do not rank independently selected equilibria across games; see SI B.3 and SI F.

[6] SI F establishes the payoff comparison within each game; SI B.2 proves the outcome comparisons across games.

The *comparison with two demands* asks whether lack of assurance can reduce peaceful concessions and increase war even when the target accepts the same demands with the same probabilities. We therefore hold each demand's size and acceptance probability fixed across games. Initial military capabilities, war costs, and the target's initial information are also unchanged. The only change in the underlying game is that the coercer can attack after compliance.

This comparison pairs monotone sorting in the protected compliance game with feigned restraint in the exposed compliance game. In both games, the target always accepts the low demand and accepts the high demand only sometimes. In the protected compliance game, belligerent coercers choose the high demand and settle if it is accepted. In the exposed compliance game, they switch to the low demand and attack after acceptance. Other coercers retain their demand choices and their decisions about war and settlement. The outcome differences therefore come entirely from belligerent coercers' switch and subsequent attack.

Keeping these bargaining terms fixed is consistent with equilibrium in both games. For every equilibrium in the exposed compliance game used here, SI E.1 constructs an equilibrium in the protected compliance game with the same demands and acceptance probabilities. The target's response remains optimal given the different coercers choosing each demand, and each coercer prefers its chosen demand except at the war-cost values dividing the groups. Coercers choosing the high demand settle after acceptance and fight after resistance in both games, so the target is indifferent between accepting and resisting the same high demand.[7]

Both comparisons cover whole families of equilibria. The pooling results hold for every common positive acceptance probability, and the two-demand results hold throughout the class just described. The latter also survive replacing the uniform distribution of war costs

[7]The target strictly prefers accepting the low demand in the constructed equilibrium in the protected compliance game. Matching a positive probability of refusing the low demand would therefore violate the target's incentives. SI E.1 proves why this comparison requires certain acceptance of the low demand.

with other continuous distributions (SI E.1). These comparisons explain how changes in concessions and demand choices produce losses from lack of assurance.

Peaceful Concessions

We first trace the loss of peaceful concessions. In the pooling comparison, lack of assurance reduces the largest concession the target will accept. In the protected compliance game, the most belligerent types can make a demand the target always resists and then fight. The remaining types make the largest acceptable demand and settle when the target complies. In the exposed compliance game, belligerent types gain by making the same demand as coercers seeking settlement: compliance improves their military position before they fight. North Korea voiced this assurance concern in May 2018 when it rejected nuclear disarmament before compensation. Its statement invoked the fate of Libya and Iraq to warn that yielding could leave a state vulnerable to attack (Pauly 2025).[8]

In the model, the target responds to this vulnerability by accepting only a smaller demand. That defensive response draws additional coercers into war. A smaller concession reduces the value of settlement more than it reduces the return from attacking. Some types that would previously have settled now prefer to fight. War therefore expands for two reasons: concessions directly improve the coercer's military position, and anticipated exploitation makes the target willing to accept only less valuable peace terms. The second effect reinforces the first.

In the comparison with two demands, belligerent coercers' imitation of restraint eliminates settlements that previously produced peaceful concessions. In the protected compliance game, belligerent coercers choose the high demand and settle if it is accepted. In the exposed compliance game, they switch to the low demand, which is always accepted, and then attack. Their switch makes compliance more frequent, but their subsequent attacks

[8]Kim Kye Gwan's statement of 16 May 2018, reported in the Pyongyang Times, 17 May 2018, archived by KCNA Watch; contemporaneous quotations appear in Anna Fifield's Washington Post report of 16 May 2018.

eliminate the peaceful concessions they previously obtained. Both the average peaceful concession per instance of compliance and the peaceful transfer averaged over all crises therefore fall.

> PROPOSITION 4 (Lack of assurance reduces peaceful concessions): Comparing the exposed compliance game with the protected compliance game gives the following results:
>
> - In the *pooling comparison*, the expected peaceful transfer and the peaceful concession obtained per instance of compliance both fall.
> - In the *comparison with two demands*, both measures fall even though compliance becomes more frequent.

The pooling comparison also lets us examine how military power affects peaceful concessions. A greater military advantage from concessions reduces the largest demand the target will accept and leads more coercers to attack instead of settling. Initial military strength affects bargaining before any concession occurs. Greater initial strength can still increase peaceful concessions, but the increase can be smaller in the exposed compliance game than in the protected compliance game. SI B.2 gives the conditions for this result.

War and Its Costs

The loss of peaceful concessions matters beyond the amount transferred: some agreements that would have ended peacefully now lead to war. We therefore examine both the probability of war and the costs borne by the two states. In the pooling comparison, the exposed compliance game produces additional wars among coercers who settle after acceptance in the protected compliance game. These coercers fight after resistance in both games. The increase in war therefore comes entirely from lost settlements after compliance.

In the comparison with two demands, the additional wars involve the belligerent coercers who switch from the high demand to the low demand. They sometimes secure a peaceful settlement in the protected compliance game but always fight in the exposed

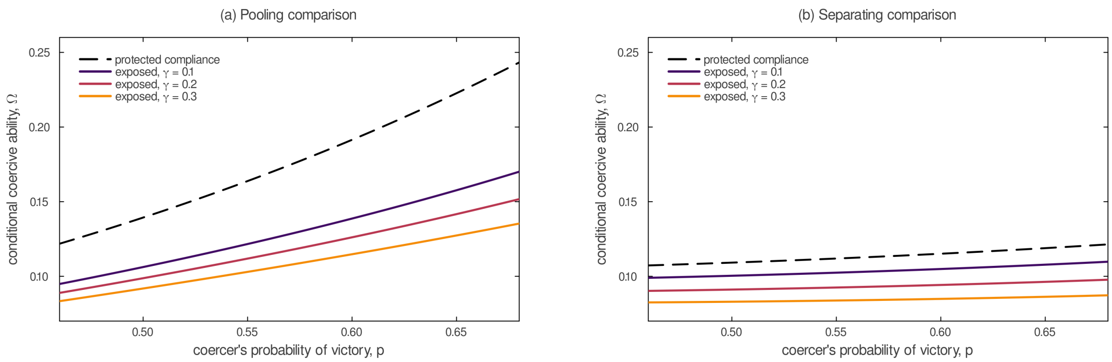

Figure 3: Lack of assurance reduces peaceful concessions

Note: Both panels show peaceful concessions per instance of compliance as the coercer's initial probability of victory increases. The black dashed line represents the protected compliance game. The solid lines represent the exposed compliance game, with the military effect of concessions set to 0.10, 0.20, and 0.30 from dark to light. In panel (a), the target always accepts each game's largest pooling demand; an accepted demand followed by attack counts as no peaceful concession. Panel (b) uses the comparison with two demands: the low demand asks for 10 percent of the good and is always accepted, while the high demand is accepted 8 percent of the time. In both panels, the target's war cost is 0.02, the audience cost is 0.30, and the upper bound on the coercer's war cost is 2.

compliance game. Other coercers retain their demand choices and decisions about war. The collective cost thus depends on both how many settlements are lost and the war costs of the coercers involved.

Proposition 5 (Lack of assurance increases war and conflict costs): Comparing the exposed compliance game with the protected compliance game gives the following results:

- In the *pooling comparison*, the probability of war and the expected conflict costs borne by the two states strictly increase.
- In the *comparison with two demands*, both measures strictly increase. The additional wars are exactly the peaceful settlements forgone when belligerent types switch to the low demand and attack.

For the comparison with two demands, we can count the lost settlements directly. Multiply the share of coercers who switch demands by the probability that their high demand is

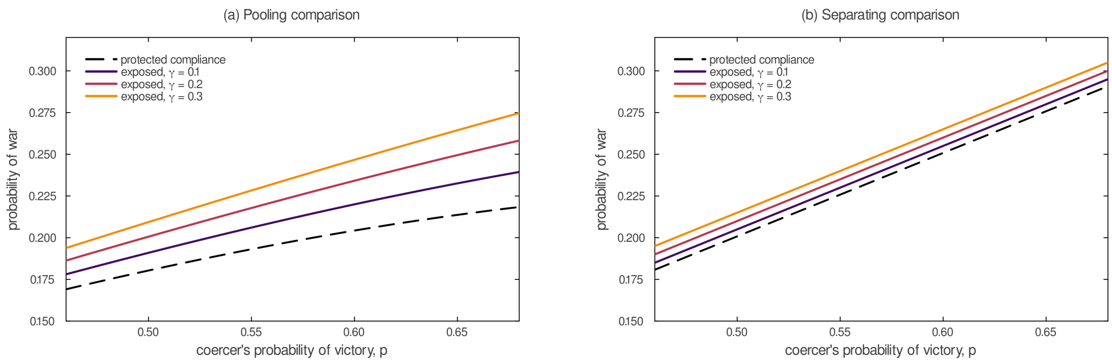

Figure 4: Lack of assurance increases the probability of war

Note: Both panels show the probability of war as the coercer's initial probability of victory increases. Parameters and line meanings are the same as in Figure 3. In panel (a), the target always accepts each game's largest pooling demand. Panel (b) uses the comparison with two demands: belligerent types switch to the low demand and attack after receiving the concession. War includes fighting after either resistance or compliance.

accepted in the protected compliance game. This gives the increase in war probability, because those accepted high demands would have ended peacefully. Summing both states' costs of these additional wars gives the increase in expected conflict costs.

We can also ask whether more frequent acceptance eliminates the assurance problem's effect on war. In the pooling comparison, raising the common acceptance probability reduces war in both games. War nevertheless remains more likely in the exposed compliance game. Panel (a) of Figure 4 shows the target always accepting each game's largest pooling demand. These are the pooling equilibria with the lowest war risk in their respective games.

Who Gains and Who Loses

The additional wars impose collective costs, but the coercer may still benefit from being able to attack after compliance. We therefore examine how the welfare loss is divided between the two states.

In the pooling comparison, the coercer bears the entire welfare loss. Weak coercers face the same probability of resistance and the same penalty for backing down in both games, so expected political costs do not change. At each game's largest pooling demand, the target receives the same expected payoff from acceptance as from resistance. The same coercers would fight after resistance in either game, so the target's expected payoff before it observes a demand is unchanged across games. The coercer therefore loses the full amount of the decline in total welfare.

In the comparison with two demands, the coercer instead benefits from lack of assurance. Each demand has the same size and acceptance probability in both games, so a coercer can obtain at least its payoff in the protected compliance game by keeping its previous choice. Belligerent coercers do better by switching to the low demand and attacking after acceptance. The target loses both the amount gained by the coercer and the additional war costs borne by the two states.

Proposition 6 (Lack of assurance can benefit or harm the coercer while reducing total welfare): In both comparisons, total welfare is lower in the exposed compliance game than in the protected compliance game.

- In the *pooling comparison*, the target's expected payoff is unchanged and the coercer's expected payoff falls.
- In the *comparison with two demands*, the coercer's expected payoff rises, while the target's payoff falls by the coercer's gain plus both states' additional expected war costs.

The opposite effects on the coercer's payoff reflect different responses by the target. In the comparison with two demands, each demand retains its size and acceptance probability, so adding the option to attack cannot reduce any coercer's payoff. In the pooling comparison, the target accepts only a smaller concession at the same probability, making settlement less valuable. The resulting loss is an average across coercers with different war costs; some coercers may still gain despite the overall decline.[9]

[9] Other models show how improved limited-attack capabilities can hurt their holder (Schram 2022) or how switching to a cheaper, weaker campaign can reduce concessions (Spaniel and İdrisoğlu 2024). Their games end when an agreement is accepted; ours adds an attack decision after compliance.

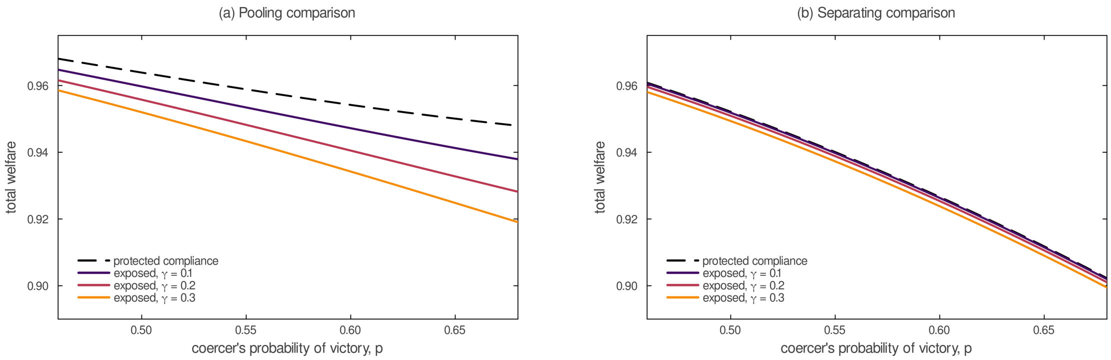

Figure 5: Lack of assurance lowers total welfare

Note: Both panels show the sum of the states' expected payoffs as the coercer's initial probability of victory increases. Parameters and line meanings are the same as in Figure 3. In panel (a), the target always accepts each game's largest pooling demand. Panel (b) uses the comparison with two demands. In both comparisons, the additional wars reduce total welfare.

Equilibrium Scope

The outcome bounds cover more equilibria than we characterize fully. We now clarify which equilibria each result covers and how the choice of comparison affects its interpretation.

SI C.1 and SI C.2 characterize all equilibria with two accepted demands; SI D.2 characterizes all equilibria with three in the exposed compliance game. In each case, coercers choose demands without randomizing, and each accepted demand asks for a nonzero concession and is chosen with nonzero probability. The characterizations include groups of coercers who receive the same payoff from different demands.

The bounds in Proposition 3 also cover equilibria with more than three accepted demands if their number is finite and they satisfy these conditions. We do not provide a full characterization of those larger equilibria. The bounds exclude equilibria with accepted demands for no concession, infinitely many accepted demands, or coercers who randomize across demands.

To isolate the military effect of concessions, a final comparison retains the attack option but removes any military advantage from compliance. If the target accepts each game's largest pooling demand with the same probability, the games produce the same expected peaceful transfer, war risk, conflict costs, and payoffs. Compliance can still occur more often in the exposed compliance game: some coercers obtain a concession before attacking, whereas their counterparts fight after resistance in the protected compliance game. SI G establishes this result.

Conclusion

The assurance problem can make demands harder to interpret. A coercer planning an attack can imitate the modest demand of one seeking peace because the concession improves its military prospects. Demands therefore need not provide a simple guide to restraint: a larger demand can be more reassuring than a smaller one. Lack of assurance also permits equilibria with three or more accepted demands. In some of these equilibria, the coercer may attack after compliance with a middle demand. Acceptance of either a smaller or a larger demand always ends in settlement.

The same assurance problem limits what bargaining can achieve. Within the equilibrium class analyzed here, we allow demand sizes, acceptance probabilities, and the number of accepted demands to differ freely across games. Lack of assurance puts both measures of peaceful concessions below their protected maxima, raises the lowest attainable war risk and conflict costs, and lowers the highest attainable coercer payoff and total welfare. These results concern the outcomes each game can sustain and do not require matching bargaining terms across games.

The paired comparisons explain how these losses arise and who bears them. Who gains depends on the target's response. In the pooling comparison, the target accepts a smaller demand with the same probability in the exposed compliance game as in the protected compliance game. The smaller concession makes settlement less attractive, and the coercer's expected payoff falls. The comparison with two demands isolates belligerent coercers' imitation of restraint. These coercers gain by switching to the low demand and attacking after receiving the concession, while other coercers' behavior stays the same. The target loses both their gain and the additional war costs borne by the two states. In both comparisons, the amount obtained through peaceful settlement falls and total welfare declines.

To assess the signaling mechanism empirically, we need evidence about how coercers choose demands. Feigned restraint requires belligerent coercers to choose the same modest demand as coercers seeking to avoid resistance. It also requires comparison with coercers who ask for more and stop after acceptance. An attack after a concession alone cannot establish this pattern.

References

Blackwell, David. 1951. "Comparison of Experiments." *Proceedings of the Second Berkeley Symposium on Mathematical Statistics and Probability*: 93–102. doi:10.1525/9780520411586-009.

Cebul, Matthew D., Allan Dafoe, and Nuno P. Monteiro. 2021. "Coercion and the Credibility of Assurances." *The Journal of Politics* 83(3): 975–91. doi:10.1086/711132.

Cho, In-Koo, and David M. Kreps. 1987. "Signaling Games and Stable Equilibria." *The Quarterly Journal of Economics* 102(2): 179–221. doi:10.2307/1885060.

Kydd, Andrew H., and Roseanne W. McManus. 2017. "Threats and Assurances in Crisis Bargaining." *Journal of Conflict Resolution* 61(2): 325–48. doi:10.1177/0022002715576571.

Paine, Jack, and Scott A. Tyson. 2020. "Uses and Abuses of Formal Models in Political Science" eds. Dirk Berg-Schlosser, Bertrand Badie, and Leonardo Morlino. *The SAGE Handbook of Political Science: A Global Perspective*: 188–202. doi:10.4135/9781529714333.n14.

Pauly, Reid B. C. 2024. "Damned If They Do, Damned If They Don't: The Assurance Dilemma in International Coercion." *International Security* 49(1): 91–132. doi:10.1162/isec_a_00488.

Pauly, Reid B. C. 2025. *The Art of Coercion*. Cornell University Press. doi:10.1353/book.131106.

Schelling, Thomas C. 1966. *Arms and Influence*. New Haven, CT: Yale University Press. doi:10.12987/9780300253481.

Schram, Peter. 2022. "When Capabilities Backfire: How Improved Hassling Capabilities Produce Worse Outcomes." *The Journal of Politics* 84(4): 2246–60. doi:10.1086/719637.

Sechser, Todd S. 2010. "Goliath's Curse: Coercive Threats and Asymmetric Power." *International Organization* 64(4): 627–60. doi:10.1017/s0020818310000214.

Slantchev, Branislav L. 2010. "Feigning Weakness." *International Organization* 64(3): 357–88. doi:10.1017/s002081831000010x.

Spaniel, William, and Işıl İdrisoğlu. 2024. "Endogenous Military Strategy and Crisis Bargaining." *Conflict Management and Peace Science* 41(4): 365–91. doi:10.1177/07388942231180499.

Supporting Information

Contents

A. Notation and Equilibrium Refinement

Throughout the SI, P and E denote the protected and exposed compliance games. Write $h = p + c_2$, $r = p + a$, and $R(c) = \max\{p - c, -a\}$. Assumptions 1–2 require

$$h < 1, \quad r < \bar{c}, \quad \bar{c} > 2p + c_2, \quad hr < p\bar{c}, \quad 0 < \gamma < 1, \quad p + \gamma < 1.$$

The war-cost distribution F is uniform on $[0, \bar{c}]$ unless stated otherwise. Comparing continuation payoffs, the coercer fights after resistance exactly when $p - c > -a$, or $c < r$. It attacks after exposed compliance exactly when $p + \gamma x - c > x$, or $c < t_E(x) = p - (1 - \gamma)x$. At a tie it chooses peace. Write $t_P(x) = p - x$ for the protected type indifferent between settlement and ordinary war.

A.1. Outcome measures

Let C denote compliance, S settlement after compliance, W war after either response, and D the demand. The measures in Definition 1 are

$$P_C = \Pr(C), \quad T = E[D\mathbb{1}[S]], \quad \Omega = \frac{T}{P_C}, \quad q = \Pr(W), \quad K = E[(c + c_2)\mathbb{1}[W]].$$

The ratio Ω requires $P_C > 0$. Define

$$I(z) = \int_0^z (c + c_2)\,\mathrm{d}F(c) = \frac{z^2 + 2c_2 z}{2\bar{c}} \quad \text{under uniformity.}$$

The coercer's expected payoff is U, the target's is U_2, and total welfare is $V = U + U_2$. Conflict cost K excludes the cost of backing down; V includes it. Outcome differences, denoted Δ, subtract protected from exposed values.

A.2. The Intuitive Criterion and Off-Path Beliefs

Let $v_j(c)$ denote type c's equilibrium payoff in game $j \in \{P, E\}$. Its payoff from demand d when the target complies with probability λ is

$$u_P(d, c, \lambda) = (1 - \lambda)R(c) + \lambda d,$$

$$u_E(d, c, \lambda) = (1 - \lambda)R(c) + \lambda \max\{d, p + \gamma d - c\}.$$

Continuation play follows the cutoff rules above. Let $\Lambda_j(d)$ be the set of compliance probabilities that are optimal for the target under some posterior on $[0,\bar{c}]$. After an unused demand d, the types not ruled out by equilibrium dominance are

$$\mathcal{T}_j(d) = \left\{ c \in [0,\bar{c}] : \max_{\lambda \in \Lambda_j(d)} u_j(d,c,\lambda) \geq v_j(c) \right\}.$$

The Intuitive Criterion rejects an equilibrium outcome if some type in $\mathcal{T}_j(d)$ strictly gains from d under every target best response to beliefs supported on $\mathcal{T}_j(d)$ (Cho and Kreps 1987). If this set is empty, the criterion imposes no restriction at d. The following construction supplies one admissible belief and response that deters every type at each unused demand.

LEMMA 1 (Completion under the Intuitive Criterion): Under Assumptions 1–2, an on-path profile with pure demand strategies extends to a PBE satisfying the Intuitive Criterion if its posteriors obey Bayes' rule, its target responses are optimal, no type prefers another used demand, and $v_j(c) \geq R(c)$ for every type. The completion leaves all on-path demands, responses, and payoffs unchanged.

Proof. Fix a game and suppress its subscript. Because $r < \bar{c}$, weak types exist. Choose a demand s used by a weak type and let λ_s be its compliance probability. All weak types face identical payoffs from each demand, so sender optimality gives them the common equilibrium payoff

$$w = \lambda_s s - (1-\lambda_s)a \leq \lambda_s s.$$

Every type can imitate this demand. Its resistance payoff is at least $-a$, and its compliance payoff is at least s in either game. Consequently, $v(c) \geq w$. Under exposure, a resolved type can also attack after imitating s, so

$$v(c) \geq R(c) + \lambda_s \gamma s \quad \text{for } c < r.$$

These bounds, together with $v(c) \geq R(c)$, suffice to specify all unused demands.

First consider $0 \leq d \leq h$ with $d \geq w$. Compliance is a target best response to belief on a resolved type that would settle after compliance. Such a type exists in both games: under exposure, $\max\{0, t_E(d)\} < r$. Thus acceptance is among the responses used to test equilibrium dominance. It gives every weak type $d \geq w$, so weak types belong to $\mathcal{T}(d)$. Put posterior probability one on a type $c > r$ and resist. The target obtains 1 from resistance and $1 - d$ from compliance, making resistance strictly optimal for $d > 0$ and weakly optimal at zero. Every sender then receives $R(c) \leq v(c)$.

Next consider $0 \leq d < w$. Settlement after compliance gives $d < w \leq v(c)$. If exposure instead induces attack, then $c < t_E(d) \leq p < r$, and the compliance payoff is

$$p + \gamma d - c = R(c) + \gamma d < R(c) + \lambda_s \gamma s \leq v(c).$$

Resistance also gives at most $v(c)$. Hence no pure or mixed target response makes the deviation profitable. If $\mathcal{T}(d)$ is nonempty, put belief on a surviving type and choose a target best response. If it is empty, use a weak-type belief and resist. Either choice is sequentially rational and deters every sender. This step can require compliance with an unused demand, but that demand remains unchosen because no type gains from it.

Finally, let $d > h$. Against a resolved settler, compliance costs $d > h$ rather than the war loss h. Against a weak type, it costs $d > 0$ rather than zero. Against an exposed attacker, it adds $\gamma d > 0$ to the war loss. Resistance is therefore strictly optimal under every posterior, so $\Lambda_j(d) = \{0\}$. The surviving types, if any, satisfy $v(c) = R(c)$. Put belief on any such type, or use any posterior if no type survives, and resist. Again, no sender gains.

The cases exhaust $[0, 1]$, including zero demands, payoff ties, and arbitrary on-path compliance probabilities. The prescribed beliefs meet the criterion whenever it restricts beliefs, and target responses deter all deviations. Bayes' rule and optimality on path are unchanged, while the cutoff rules supply optimal continuation play after either target action. This completes the equilibrium. $\square$

B. Pooling Equilibria and Comparisons

Pooling has one positive on-path demand accepted with positive probability; any other used demand is always resisted. Write $g_P = 0$, $g_E = \gamma$, and $t_j(x) = p - \big(1 - g_j\big)x$. The following argument establishes pooling existence in Proposition 1 and the maximal demands used in the comparisons.

B.1. Equilibrium construction and maximal demands

Proof. In the protected compliance game, types with $R(c) < x$ prefer an accepted demand x to resistance, and those with $R(c) > x$ prefer resistance. An accepted demand $x \geq p$ would therefore attract the entire prior, apart from a possible endpoint. Target optimality would require $x \leq hF(r) = h\frac{r}{\bar{c}} < p$, a contradiction. Thus $x < p$, and the accepted pool is $[p - x, \bar{c}]$. Its target acceptance surplus, multiplied by the posterior denominator and $\bar{c}$, is

$$h(r - t_P(x)) - x(\bar{c} - t_P(x)) = ha - (\bar{c} - 2p - c_2)x - x^2.$$

In the exposed compliance game, every type strictly gains from compliance with a positive demand. An attacker gains $\gamma x > 0$ over resistance; a resolved settler has $x > p - c$; a weak type has $x > -a$. Thus every type joins the accepted pool. Compliance costs the target at least x, so again $x \leq hF(r) < p$. Its acceptance surplus, multiplied by $\bar{c}$, is

$$hr - (h + \gamma x)t_E(x) - x(\bar{c} - t_E(x)) = ha - [\bar{c} - (1 - \gamma)(2p + c_2)]x - (1 - \gamma)^2 x^2.$$

Both expressions can be written as

$$\Psi_j(x) = hr - \bar{c}x - t_j(x)\left[t_j(x) + c_2\right] = ha - \left[\bar{c} - \left(1 - g_j\right)(2p + c_2)\right]x - \left(1 - g_j\right)^2 x^2.$$

Under Assumptions 1–2, $\Psi_j(0) = ha > 0$, $\Psi_j(p) < 0$, and $\Psi'_{j(x)} < 0$ for $x \geq 0$. Each game therefore has a unique maximal pooling demand $x_j \in (0, p)$, determined by

$$x_P^2 + (\bar{c} - 2p - c_2)x_P - ha = 0,$$
$$(1 - \gamma)^2 x_E^2 + [\bar{c} - (1 - \gamma)(2p + c_2)]x_E - ha = 0.$$

At x_P, the negative exposed surplus is

$$-\Psi_E(x_P) = \gamma x_P[2p + c_2 - (2 - \gamma)x_P] > 0.$$

Consequently, writing $t_j = t_j\left(x_j\right)$,

$$0 < x_E < x_P < p, \quad t_E - t_P = \gamma x_E + x_P - x_E > 0, \quad 0 < t_P < t_E < r < \bar{c}.$$

For any $0 < x \leq x_P$, assign protected types below $p - x$ to demand 1 and all others to x. The target resists 1 because $h < 1$ and its posterior contains only resolved types. For any $0 < x \leq x_E$, assign every exposed type to x. In either construction, the target accepts x certainly if $x < x_j$ and with any probability $\lambda_j \in (0, 1]$ if $x = x_j$. The payoff comparisons above verify sender optimality and $v_j(c) \geq R(c)$; Bayes' rule gives the stated posteriors. Lemma 1 completes both constructions. No larger accepted pooling demand is feasible because its acceptance surplus is negative. $\square$

B.2. Outcome comparisons

The formulas below concern the maximal demands. Define the common resistance outcomes

$$q_R = F(r), \quad K_R = I(r), \quad U_R = -a + \frac{r^2}{2\bar{c}}, \quad U_{2,R} = 1 - hF(r), \quad V_R = 1 - I(r) - a[1 - F(r)].$$

Directly counting settlement, war, and backing down gives, for $j \in \{P, E\}$,

$$\Omega_P = x_P, \quad \Omega_E = x_E[1 - F(t_E)], \quad T_j(\lambda_j) = \lambda_j x_j \left[1 - F(t_j)\right],$$
$$q_j(\lambda_j) = \lambda_j F(t_j) + (1 - \lambda_j) q_R, \quad K_j(\lambda_j) = \lambda_j I(t_j) + (1 - \lambda_j) K_R,$$
$$U_j(\lambda_j) = \lambda_j \left[x_j + \frac{t_j^2}{2\bar{c}}\right] + (1 - \lambda_j) U_R,$$
$$U_{2,j}(\lambda_j) = U_{2,R}, \quad V_j(\lambda_j) = 1 - K_j(\lambda_j) - (1 - \lambda_j) a[1 - F(r)].$$

The payoff formula integrates $\max\{x_j, p + g_j x_j - c\}$ under certain compliance and then averages with $R(c)$ when resistance occurs. In the protected game, types below t_P receive their war payoff regardless of the pool's response. The target's payoff equals $U_{2,R}$ because it is indifferent at every accepted maximal pool and resists any remaining types.

Pooling parts of Proposition 4 through Proposition 6.

Proof. Set $\lambda_P = \lambda_E = \lambda > 0$. Since $x_E < x_P$ and $t_E > t_P$,

$$\Omega_E < x_E < x_P = \Omega_P, \quad T_E(\lambda) < T_P(\lambda),$$
$$\Delta q = \lambda[F(t_E) - F(t_P)] > 0, \quad \Delta K = \lambda[I(t_E) - I(t_P)] > 0.$$

Expected backing-down costs and the target's payoff are identical across games, so

$$\Delta V = \Delta U = -\lambda[I(t_E) - I(t_P)] < 0, \quad \Delta U_2 = 0.$$

These inequalities prove all three pooling comparisons for every common positive acceptance probability. $\square$

The military effects discussed in the main text follow by implicit differentiation. Put $s = 1-\gamma$, $D_P = 2x_P + \overline{c} - 2p - c_2$, and $D_E = 2s^2 x_E + \overline{c} - s(2p + c_2)$. Then

$$\frac{\partial x_E}{\partial \gamma} = -\frac{x_E[2p + c_2 - 2sx_E]}{D_E} < 0, \quad \frac{\partial t_E}{\partial \gamma} = x_E - s\frac{\partial x_E}{\partial \gamma} > x_E,$$

$$\frac{\partial x_P}{\partial p} = \frac{2x_P + a}{D_P}, \quad \frac{\partial x_E}{\partial p} = \frac{2sx_E + a}{D_E}.$$

Thus the smaller concession reinforces the direct military increase in the attack cutoff. For the comparison of initial strength, let $y = sx_E$ and $H = 2p + c_2$. Substituting $x_E = \frac{y}{s}$ into the exposed root equation gives $y^2 + \left(\frac{\overline{c}}{s} - H\right)y - ha = 0$. Its larger linear coefficient implies $0 < y < x_P$. Subtracting the demand derivatives gives

$$\frac{\partial x_P}{\partial p} - \frac{\partial x_E}{\partial p} = \frac{2(\overline{c} - H - a)(x_P - y) + \gamma(2x_P + a)(H - 2y)}{D_P D_E} > 0$$

for $p < \overline{p} = \frac{\overline{c} - c_2 - a}{2}$. In this range,

$$\frac{\partial t_E}{\partial p} = \frac{\overline{c} - s(H + a)}{D_E} > 0, \quad \frac{\partial \Omega_E}{\partial p} = \frac{\partial x_E}{\partial p}[1 - F(t_E)] - \frac{x_E}{\overline{c}}\frac{\partial t_E}{\partial p} < \frac{\partial \Omega_P}{\partial p}.$$

The exposed derivative is positive as well: substitution of its root equation yields

$$\overline{c}s\frac{\partial \Omega_E}{\partial p} = \frac{a(\overline{c} - H) + 2y(\overline{c} - p + a) + 3y^2}{2y + \frac{\overline{c}}{s} - H} > 0.$$

Hence greater strength raises peaceful concessions less under exposure throughout this range.

B.3. Independent Compliance Rates at Maximal Pooling Demands

Matching compliance probabilities is consequential because the Intuitive Criterion permits their independent selection. For example, let $p = 0.60$, $c_2 = 0.02$, $a = 0.30$, $\overline{c} = 2$, and $\gamma = 0.10$, which satisfy Assumptions 1–2. The maximal demands are $x_P \approx 0.19146$ and $x_E \approx 0.17782$. Choosing $\lambda_P = 0.10$ and $\lambda_E = 1$ gives

$$q_P(\lambda_P) \approx 0.42543 > 0.21998 \approx q_E(\lambda_E), \quad T_P(\lambda_P) \approx 0.01524 < 0.13870 \approx T_E(\lambda_E)$$

$$V_P(\lambda_P) \approx 0.65657 < 0.94721 \approx V_E(\lambda_E)$$

Both equilibria satisfy the Intuitive Criterion by Lemma 1. The common-response comparisons therefore identify assurance effects conditional on the matched response, not an ordering of all equilibria satisfying the criterion.

C. Two-Demand Equilibria

This section gives the complete characterization with pure demand strategies and two strictly positive demands, each chosen by a positive mass of types and receiving positive compliance. Let $0 < x < y \leq 1$, and let $\alpha > 0$ and $\beta > 0$ denote the target's compliance probabilities after y and x. Necessarily, $\alpha < \beta$. To see why, first note that any type assigned to x must obtain a nonnegative compliance gain: it can otherwise demand more than h, induce resistance regardless of the target's belief, and obtain its resistance continuation payoff. Moreover, the compliance gain is strictly increasing in the demand. If $\alpha \geq \beta$, therefore, every type assigned to x would strictly prefer to mimic y. A two-demand equilibrium thus requires $0 < \alpha < \beta \leq 1$. Write $\rho = \frac{\alpha}{\beta} < 1$. For game $j \in \{P, E\}$, define the sender's gain from compliance rather than resistance as

$$G_P(d,c) = \begin{cases} d - p + c & \text{if } c < r \\ d + a & \text{if } c \geq r \end{cases}$$

$$G_E(d,c) = \begin{cases} \max\{\gamma d, d - p + c\} & \text{if } c < r \\ d + a & \text{if } c \geq r. \end{cases}$$

A type chooses y rather than x exactly when

$$\rho G_j(y,c) \geq G_j(x,c) \tag{2}$$

Under protected compliance, compare the gains directly:

$$\rho G_P(y,c) - G_P(x,c) = \begin{cases} \rho y - x + (1-\rho)(p-c) & \text{if } c < r \\ \rho(y+a) - (x+a) & \text{if } c \geq r. \end{cases}$$

This difference is strictly decreasing below r and constant above it. Its nonnegative set therefore has at most one cutoff. The argument also covers $c_2 = 0$ without dividing by a zero compliance gain. Under exposed compliance, both gains are strictly positive. Define $t_d = p - (1-\gamma)d$. Then

$$\frac{G_E(x,c)}{G_E(y,c)} = \begin{cases} \frac{x}{y} & \text{if } c \le t_y \\ \frac{\gamma x}{y-p+c} & \text{if } t_y < c < t_x \\ \frac{x-p+c}{y-p+c} & \text{if } t_x \le c < r \\ \frac{x+a}{y+a} & \text{if } c \ge r. \end{cases}$$

The ratio is first flat, then decreasing, then increasing, and finally flat, on the branches that intersect the type support. This nonmonotonicity permits either one cutoff or two.

Receiver optimality rules out the upper indifference plateau when $x > 0$. If $\rho \ge \frac{x+a}{y+a}$, every resolved type under protection strictly prefers y. Under exposure, every resolved type that would settle after complying with x strictly prefers y. Thus a low-demand pool could contain only weak types, and under exposure also attackers. Compliance costs the target $x > 0$ against a weak type and an additional $\gamma x > 0$ against an attacker. The target would strictly resist. Consequently, every equilibrium in this class has $\rho < \frac{x+a}{y+a}$, and its upper cutoff lies below r.

C.1. Monotone sorting

In a monotone profile, types $c < k$ demand y and types $c \ge k$ demand x. The following conditions are necessary and sufficient.

A monotone two-pool equilibrium in the protected compliance game exists exactly under conditions (i)–(iii):

(i) $y = h$ and $0 < k < r$;

(ii) $\rho = \frac{x-p+k}{h-p+k} = \frac{x-p+k}{k+c_2}$ lies in $(0,1)$; and

(iii) $x(\bar{c} - k) \le h(r - k)$.

A monotone two-pool equilibrium in the exposed compliance game exists exactly under conditions (iv)–(viii). Let $m_y = \max\{0, t_y\}$. Then

(iv) $m_y < k < r$ and $t_x < k$;

(v) $(h - y)(k - m_y) = \gamma y m_y$;

(vi) $\rho = \frac{x-p+k}{y-p+k}$ lies in $(0,1)$;

(vii) $\rho G_E(y, 0) \ge G_E(x, 0)$; and

(viii) $x(\bar{c} - k) \le h(r - k)$.

In either game, a strict final inequality implies $\beta = 1$ and $\alpha = \rho$. At equality, any $\beta \in (0, 1]$ and $\alpha = \rho\beta$ is admissible.

Both demands assure peace in the following exposed subclass. If no type attacks after receiving the high demand, then $m_y = 0$, condition (v) forces $y = h$, and the additional exposed incentive restrictions reduce to

$$\gamma \leq \frac{c_2}{h}, \quad \rho c_2 \geq \gamma x \tag{3}$$

These are exactly the protected profiles that survive exposure without changing their on-path behavior. When $m_y > 0$, by contrast, low-cost types demand $y < h$ and attack after compliance. The target can still mix because settlements by the remaining high-demand types offset those attacks.

Proof. We prove necessity using sender indifference and target optimality, then apply Lemma 1 for sufficiency. Every equilibrium satisfying the criterion is a PBE, so the usual incentive and posterior restrictions remain necessary. Under protected compliance, the gain difference in (2) is strictly decreasing below r. The upper-plateau exclusion above gives $k < r$, and both pools have positive mass, so $k > 0$. Cutoff indifference gives

$$\rho = \frac{x - p + k}{y - p + k}$$

Every type in the high pool fights after resistance and settles after compliance. Since $0 < \alpha < \beta \leq 1$, the target mixes after y, which requires $y = h$. Its cutoff gain $h - p + k = k + c_2$ is strictly positive, so the displayed ratio is well defined. At the low signal, the target's compliance surplus, multiplied by the posterior denominator, is $h(r - k) - x(\bar{c} - k)$, which yields the final inequality. The decreasing gain difference assigns exactly the types below k to y. Types above r strictly prefer x because cutoff indifference implies $(x + a) > \rho(h + a)$. Every on-path gain is nonnegative: $h - p + c = c_2 + c \geq 0$ in the high pool, and $x - p + k > 0$ at the bottom of the low pool. This verifies necessity and sender optimality for sufficiency.

Under exposed compliance, the high-demand pool contains attackers $\left[0, m_y\right)$ and settlers $\left[m_y, k\right)$. The target's unnormalized payoff difference between compliance and resistance after y is

$$(h-y)(k-m_y)-\gamma y m_y$$

Because $0<\alpha<1$, this expression must equal zero. A high pool consisting only of attackers would make it negative, so $m_y<k$. A nondegenerate upper cutoff must lie on the increasing branch of the ratio, giving $t_x<k$ and condition (vi). The ratio is weakly decreasing before t_x and increasing afterward up to r. Thus the bottom type supplies the only additional incentive constraint, condition (vii), while the cutoff supplies the upper boundary. The low signal contains only settlers after compliance, so target optimality again gives $x(\bar{c}-k)\le h(r-k)$.

Conversely, conditions (iv)–(viii) make the target indifferent after y and weakly willing to comply after x. The shape of the ratio and the bottom and cutoff inequalities assign every $c<k$ to y and every $c\ge k$ to x. Types $c<m_y$ attack after high-demand compliance, those in $[m_y,k)$ settle, and every type assigned to x settles because $t_x<k$. Exposed compliance gains are strictly positive, and the protected gains were shown to be nonnegative. Both profiles therefore meet Lemma 1, with $s=x$, $\lambda_s=\beta$, and $w=\beta x-(1-\beta)a$. Its off-path prescription completes both games for every admissible β, including receiver-indifference boundaries. Under protection with $c_2=0$, the zero-cost type has $v_P(0)=R(0)$; belief on that type and compliance with an unused $d<w$ satisfy the criterion without creating a profitable deviation. This proves necessity and sufficiency under the stated solution concept. □

C.2. Feigned restraint

In a two-sided profile, an intermediate group demands y, while attackers and high-cost coercers demand x. The necessary and sufficient conditions include both strict demand choices and indifference among attackers.

With two strictly positive demands, protected compliance has no two-sided two-pool PBE. Exposed compliance has the following two classes.

First, a regular two-sided equilibrium exists if and only if $y=h$, $0<x<p$, and

$$\rho_{\min}<\rho<\min\{\rho_0,\rho_r\}$$

where

$$t_h=\gamma h-c_2,\quad t_x=p-(1-\gamma)x$$

$$\rho_{\min} = \frac{\gamma x}{t_x + c_2}, \quad \rho_0 = \frac{\gamma x}{\max\{\gamma h, c_2\}}, \quad \rho_r = \frac{x+a}{h+a}$$

The cutoffs are

$$b = \frac{\gamma x}{\rho} - c_2, \quad k = \frac{p - x + \rho c_2}{1-\rho}$$

they satisfy $\max\{0, t_h\} < b < t_x < k < r$, and target optimality requires

$$x(\bar{c} - k + \gamma b) \le h(r-k) \tag{4}$$

Second, a tie-supported two-sided equilibrium exists if and only if $0 < x < y < h$, $t_y > 0$, $\rho = \frac{x}{y}$, $k = p$, and

$$b = t_y - \frac{(h-y)\big(p - t_y\big)}{\gamma y}$$

lies in $\big(0, t_y\big)$, while

$$x(\bar{c} - p) + \gamma x b \le ha$$

In a canonical representation, types in $[0, b)$ and $[p, \bar{c}]$ demand x, while types in $[b, p)$ demand y. More generally, any measurable subset of $\big[0, t_y\big]$ with total length b, hence probability $\frac{b}{\bar{c}}$, can form the lower low-demand pool. Boundary equalities give the corresponding weak and degenerate profiles.

In the regular class, low-cost types below b attack after low-demand compliance, types in $[b, k)$ settle after high-demand compliance, and types above k settle after the low demand. If the target's final inequality is strict, $\beta = 1$ and $\alpha = \rho$; at equality, any $\beta \in (0, 1]$ with $\alpha = \rho\beta$ works. In the tie-supported class, all types below t_y are indifferent between the two signals. The specified assignment divides this plateau so that the target is willing to mix after y.

Proof. Under protected compliance, the gain difference $\rho G_P(y, c) - G_P(x, c)$ is strictly decreasing for $c < r$ and constant above r. The high-demand set is its nonnegative set, which cannot be an interior interval. Two-sided sorting is impossible.

Under exposed compliance, the ratio displayed at the start of the two-demand analysis is flat, decreasing, increasing, and flat over its four branches. For a regular two-sided profile, ρ

must cross it once on the decreasing branch and once on the increasing branch. The endpoint restrictions are exactly

$$\rho_{\min} < \rho < \min\{\rho_0, \rho_r\}$$

Solving the two equalities yields b and k and implies $\max\{0, t_h\} < b < t_x < k < r$. Because every high-pool type fights after resistance and settles after compliance, target mixing forces $y = h$. At the low signal, attackers $[0, b)$ impose the extra exposure loss $\gamma x b$, settlers $[k, r)$ avert war, and types above r accept a transfer they could otherwise avoid. The target's unnormalized compliance surplus is therefore

$$h(r-k) - x(\bar{c} - k + \gamma b)$$

which proves the regular condition.

The upper flat branch cannot support indifference between two positive-compliance demands, as established at the start of the two-demand analysis. A positive-mass indifference set can therefore arise only on the bottom flat branch, where the ratio equals $\frac{x}{y}$. Hence $\rho = \frac{x}{y}$ and $t_y > 0$. Above t_x, cutoff indifference gives $k = p$. Let types $[0, b)$ use x and $\left[b, t_y\right]$ use y. Target indifference after y requires

$$(h-y)\left(p - t_y\right) = \gamma y\left(t_y - b\right)$$

which yields the stated b. The target's low-signal compliance condition is $x(\bar{c} - p) + \gamma x b \leq ha$. These restrictions are also sufficient because every type below t_y is indifferent, types in $\left(t_y, p\right)$ strictly prefer y, and types above p prefer x.

In both exposed classes, on-path compliance gains are strictly positive. The on-path incentive and target-optimality conditions therefore meet Lemma 1, with $s = x$, $\lambda_s = \beta$, and $w = \beta x - (1-\beta)a$. Its off-path prescription completes both classes under the Intuitive Criterion, including mixed responses and arbitrary assignments of indifferent bottom types. The protected impossibility follows already from PBE incentives. The strict classes, their boundary equalities, and the positive-mass bottom plateau exhaust the ways a flat-decreasing-increasing-flat ratio can generate a two-sided assignment. $\square$

C.3. Proof of Proposition 1

Proof. Pooling exists throughout the maintained parameter space by SI B. For protected monotone sorting, choose $y = h$, $k = p$, and $\rho = \frac{x}{h}$, where

$$0 < x < \min\left\{h, \frac{ha}{\bar{c} - p}\right\}.$$

The cutoff condition becomes $\rho = \frac{x-p+k}{k+c_2} = \frac{x}{h}$, and the low-demand surplus is $ha - x(\bar{c} - p) > 0$. These satisfy SI C.1. The decreasing protected gain difference rules out feigned restraint.

If both demands assure peace under exposure, the high pool has no attackers. Target mixing then requires $y = h$, and the absence of attackers requires $t_h = \gamma h - c_2 \leq 0$. Conversely, when $\gamma \leq \frac{c_2}{h}$, the preceding protected construction also satisfies the exposed conditions: $t_x < p = k$, and the bottom-type constraint is

$$\rho G_E(h, 0) = \frac{xc_2}{h} \geq \gamma x = G_E(x, 0).$$

This proves the exact threshold for monotone sorting with settlement after either demand.

For feigned restraint, Lemma 3 constructs a regular exposed equilibrium at every admissible parameter vector, including the uniform prior. Finally, the gain difference under protection has one crossing; the exposed gain ratio is flat, decreasing, increasing, and flat. Its sublevel set is a lower or an interior interval. Target optimality excludes upper-plateau ties, and SI C.2 includes all lower-plateau assignments. These two orderings therefore exhaust the two-positive-demand class, including indifferent groups. Lemma 1 completes every construction under the stated refinement. □

D. Additional Accepted Demands

D.1. Protected compliance

This section establishes the protected-compliance part of Proposition 2. The bound follows from sender and target optimality; it does not require an additional equilibrium refinement.

Lemma 2 (Accepted demands under protected compliance): In any PBE in the protected compliance game with pure demand strategies, at most two distinct strictly positive demands can each be chosen by a positive prior mass of types and accepted with positive probability. If there are two, the larger equals $h = p + c_2$ and is chosen only by resolved types, up to sets of prior probability zero. The bound permits other demands that are zero or always resisted.

Proof. Take any two distinct positive accepted demands $x < y$, and write $\lambda_x = \lambda(x)$ and $\lambda_y = \lambda(y)$. A type choosing x must have $G_P(x, c) \geq 0$: otherwise it could demand 1, induce resistance, and obtain a strictly higher payoff. Since $G_P(y, c) = G_P(x, c) + y - x$, the inequality $\lambda_y \geq \lambda_x$ would make y strictly preferable for every type choosing x. Hence $0 < \lambda_y < \lambda_x \leq 1$.

Weak types must strictly prefer x to y. To prove this, suppose instead that $\lambda_y(y + a) \geq \lambda_x(x + a)$. For every resolved type $c < r$, the difference between the gains from the two demands is

$$\begin{aligned} \lambda_y G_P(y, c) - \lambda_x G_P(x, c) = \lambda_y(y + a) - \lambda_x(x + a) \\ +\big(\lambda_y - \lambda_x\big)(c - r) > 0. \end{aligned}$$

Every resolved type would therefore strictly prefer y. The pool choosing x could contain only weak types, against whom the target strictly prefers resistance to conceding $x > 0$. This contradicts positive compliance. Thus weak types strictly prefer the smaller of any two positive accepted demands.

Now suppose three such demands $x < y < z$ existed. The preceding argument excludes weak types from both y and z and gives $0 < \lambda_z < \lambda_y < \lambda_x \leq 1$. The target must therefore mix after both higher demands. Because their coercers are resolved, resistance gives the target $1 - h$, whereas accepting a demand d gives it $1 - d$. Indifference requires $y = h = z$, contradicting distinct demands. The same reasoning implies that the larger of exactly two positive accepted demands equals h and contains only resolved types. Zero and always-resisted demands do not change this pairwise argument. $\square$

D.2. Three-demand equilibria under exposed compliance

This section characterizes equilibria in the exposed compliance game with deterministic demand choices whose entire set of on-path demands consists of exactly three strictly positive amounts. Each demand is chosen by positive prior mass and receives positive compliance. The prior remains uniform, and Assumptions 1–2 are unchanged. An additional used zero demand is outside this class. Every type strictly prefers a positive demand receiving positive compliance to a surely resisted demand, so no additional refused demand can attract positive mass.

Three-demand equilibria have three forms: attackers strictly prefer the middle demand, are indifferent between the low and middle demands, or are indifferent between the middle and high demands. Weak coercers choose the lowest demand, and acceptance becomes less likely as the demand rises. The conditions below characterize every form, including all admissible target responses and assignments of indifferent coercers. Every admissible on-path profile has a completion satisfying the Intuitive Criterion.

The existence condition is

$$t_h = \gamma h - c_2 > 0.$$

By Proposition 1, this is precisely the region in which every equilibrium with two positive accepted demands involves a risk of attack after compliance. Three-demand bargaining therefore becomes possible above the same exposure threshold.

Order the demands and write their acceptance probabilities as

$$0 < x < z < y \leq h, \quad \left(\lambda_x, \lambda_z, \lambda_y\right) = (\beta, \beta\sigma, \beta\sigma\rho),$$

$$0 < \beta \leq 1, \quad 0 < \sigma < 1, \quad 0 < \rho < 1.$$

Thus σ is the middle demand's response relative to the low demand's response, and ρ is the high demand's response relative to the middle demand's. Write $t_d = p - (1-\gamma)d$. The quantities describing sets of attackers below are lengths in the cost support; dividing by $\overline{c}$ gives their prior probabilities.

For each family, let S_x be the target's low-demand compliance surplus multiplied by its positive posterior denominator. The response restriction is

$$S_x \geq 0, \quad (1-\beta)S_x = 0.$$

A strict surplus requires certain low-demand compliance. At equality, every $\beta \in (0,1]$ is admissible, with both higher responses scaled in the same proportion. Both higher responses remain strictly between zero and one. The following conditions apply when $t_h > 0$ and disregard assignments of cutoff types with prior probability zero.

D.2.1. Attackers Strictly Prefer the Middle Demand

Set $y = h$, and define

$$\rho_{\min} = \frac{\gamma z}{h-(1-\gamma)z}, \quad \rho_{\min} < \rho < \frac{z}{h},$$

$$b = \frac{\gamma z}{\rho} - c_2, \quad k = \frac{p - z + \rho c_2}{1-\rho}, \quad \ell = k + \frac{\gamma z b}{h-z}.$$

The remaining conditions are

$$p < \ell < r, \quad \sigma = \frac{x+\ell-p}{z+\ell-p}, \quad S_x = h(r-\ell) - x(\bar{c}-\ell) \geq 0.$$

Together with $0 < x < z < h$, these restrictions imply

$$0 < t_h < b < t_z < k < p < \ell < r, \quad \sigma z > \max\{x, \sigma\rho h\}.$$

The demand assignment is

$$d(c) = \begin{cases} z \text{ if } 0 \leq c < b \\ h \text{ if } b \leq c < k \\ z \text{ if } k \leq c < \ell \\ x \text{ if } \ell \leq c \leq \bar{c}. \end{cases}$$

Types below b attack after compliance. Every other type settles after compliance, and each type strictly prefers its assigned demand except at b, k, and ℓ. The middle demand attracts both the lowest-cost attackers and an interval of resolved settlers. Weak types choose the lowest demand.

D.2.2. Attackers Are Indifferent between the Low and Middle Demands

Set $y = h$, require $\rho_{\min} < \rho < \frac{z}{h}$, and define b and k as in the preceding family. Now impose

$$\sigma = \frac{x}{z}, \quad m_z = \frac{(h-z)(p-k)}{\gamma z}, \quad 0 < m_z \le b, \quad m_x = b - m_z.$$

The low-demand constraint becomes

$$S_x = ha - x(\bar{c} - p + \gamma m_x) \ge 0.$$

The cutoffs satisfy $0 < t_h < b < t_z < k < p < r$. Every type in $[0, b)$ is indifferent between x and z and attacks after either demand is accepted. Choose any measurable subset of $[0, b)$ with total length m_x to demand x; its complement demands z. Types in $[b, k)$ demand h, types in $[k, p)$ demand z, and types in $[p, \bar{c}]$ demand x. All types in these three upper groups settle after compliance.

The boundary $m_x = 0$ is included: all attackers then choose z, although they remain indifferent between the two lower demands. The middle demand must receive a positive mass of attackers. Here

$$x = \sigma z > \sigma \rho h.$$

D.2.3. Attackers Are Indifferent between the Middle and High Demands

Allow $y \le h$ and impose

$$\rho = \frac{z}{y}, \quad m_y = \frac{(h-y)\big(p - t_y\big)}{\gamma y} = \frac{(1-\gamma)(h-y)}{\gamma},$$

$$m_z = t_y - m_y = t_h - \frac{(1-\gamma)^2(h-y)}{\gamma} > 0.$$

The nonnegative length m_y is the high-demand attacker group, and m_z is the middle-demand attacker group. These restrictions imply $0 < t_y < p$. Define

$$\ell = p + \frac{\gamma z m_z}{h - z}, \quad p < \ell < r, \quad \sigma = \frac{x + \ell - p}{z + \ell - p}.$$

The low-demand constraint is

$$S_x = h(r - \ell) - x(\bar{c} - \ell) \ge 0.$$

Every type in $\big[0, t_y\big)$ is indifferent between z and y and attacks after either is accepted. Choose any measurable subset of this interval with total length m_y to demand y; its complement demands z. Types in $\left[t_y, p\right)$ demand y, types in $[p, \ell)$ demand z, and types in $[\ell, \bar{c}]$ demand x. All types in these upper groups settle after compliance.

The boundary $y = h$ gives $m_y = 0$, so only the middle demand receives attackers. When $y < h$, both tied demands receive a positive mass of attackers. In this family,

$$\sigma z = \sigma \rho y > x.$$

D.2.4. Proof of the three-demand characterization

Proof. The protected impossibility in Proposition 2 follows from Lemma 2. For exposure, we establish necessity from target optimality and the sender's gain envelope, verify sufficiency and the Intuitive Criterion, and then prove the exact existence region. The strict first family establishes the additional signaling pattern: the middle demand can carry attack risk while both extreme demands assure settlement. The extension to the full finite class follows in SI D.3.

Demand and response restrictions. Against a resolved settler, a demand $d > h$ costs the target more than resistance. Against a weak type it costs more than backing down, and against an attacker it adds $\gamma d > 0$ to the resistance loss. Thus every posterior induces resistance at $d > h$, and the highest accepted demand satisfies $y \leq h$. The exposed compliance gain is strictly positive and strictly increasing in the demand for every type. If a larger used demand received compliance at least as often as a smaller one, every type assigned to the smaller demand would strictly prefer to mimic the larger one. Hence

$$0 < \lambda_y < \lambda_z < \lambda_x \leq 1.$$

Weak types strictly prefer the smaller of any two positive accepted demands. To see this, suppose $d_1 < d_2$ but $\lambda_2(d_2 + a) \geq \lambda_1(d_1 + a)$. For a resolved type that settles after compliance with d_1,

$$\begin{aligned}\lambda_2 G_E(d_2, c) - \lambda_1 G_E(d_1, c) \geq \lambda_2(d_2 + a) - \lambda_1(d_1 + a) \\ +(\lambda_2 - \lambda_1)(c - r) > 0.\end{aligned}$$

No resolved settler would choose d_1. Its pool could contain only weak types and attackers, both of whom impose a strictly negative compliance surplus at a positive demand. This contradicts positive compliance. Therefore all weak types choose x, and both higher pools contain only resolved types. Both higher target responses are interior. Since $z < y \leq h$, target indifference at z requires a positive mass of attackers and a positive mass of resolved settlers.

The sender's gain envelope. For a resolved type, divide the gain from the three available demands by β. Let

$$M = \max\{x, \sigma z, \sigma\rho y\}.$$

The best available gain is

$$H(c) = \max\{\gamma M, x - p + c, \sigma(z - p + c), \sigma\rho(y - p + c)\}.$$

A type that attacks at its chosen demand must use a demand whose probability-times-demand product equals M. Because the middle pool contains attackers,

$$\sigma z = M \geq \max\{x, \sigma\rho y\}.$$

The three settlement lines have distinct positive slopes. Their intersections with one another or with the constant branch occur at isolated costs. Positive-mass indifference among resolved types can therefore arise only when two maximal constant branches coincide.

The high product is strictly smaller. Suppose $\sigma\rho y < \sigma z$. No attacker chooses y, so target mixing requires $y = h$. The high settlement line crosses the maximal constant at b and the middle settlement line at k. A positive high-demand interval requires

$$b < k \equiv \rho > \frac{\gamma z}{h - (1 - \gamma)z}.$$

The strict product inequality gives $\rho < \frac{z}{h}$, and these inequalities imply $b < t_z < k < p$. The positive middle-attacker mass also requires $b > 0$.

If $x < \sigma z$, the middle and low settlement lines cross at a unique $\ell > p$, satisfying

$$\ell = p + \frac{\sigma z - x}{1 - \sigma}.$$

Weak types' strict preference for x requires $\ell < r$. The envelope assigns z below b, h between b and k, z between k and ℓ, and x above ℓ. This is the first family's assignment. If $x = \sigma z$, the middle and low settlement lines instead cross at p. Both demands share the maximal constant below b, so the attackers there can be divided between them. This gives the second family's assignment.

The high and middle products are equal. Suppose $\sigma\rho y = \sigma z$. Then $\rho = \frac{z}{y}$. If $x < \sigma z$, the high settlement line dominates the middle line below p, while the middle dominates above p until the low line takes over at $\ell > p$. The constant ends at t_y. The middle pool's required attackers imply $t_y > 0$, and weak types' strict preference again gives $\ell < r$. These are precisely the third family's assignment sets.

If all three products are equal, the high settlement line dominates the middle line for $c < p$ and the low line dominates it for $c > p$. The middle demand could contain only attackers, apart from the probability-zero type $c = p$. Its target response would then be resistance. Triple equality is impossible. These cases exhaust all product orders. Ties between nonmaximal products create no further assignments because their constant branches are never optimal, and the weak-type ranking already excludes upper-plateau ties.

Target optimality and the permitted attacker assignments. Uniformity makes the target's surplus depend on the lengths of the assigned attacker and settler sets, including when a tied attacker set is not an interval. In the first family, the high-demand surplus is zero and middle-demand indifference requires

$$(h - z)(\ell - k) - \gamma z b = 0.$$

This determines ℓ. In the second family, middle-demand indifference requires

$$(h - z)(p - k) - \gamma z m_z = 0,$$

which determines the required middle-attacker length and leaves $m_x = b - m_z$. In the third family, high- and middle-demand indifference require

$$(h - y)\big(p - t_y\big) - \gamma y m_y = 0,$$
$$(h - z)(\ell - p) - \gamma z m_z = 0,$$

together with $m_y + m_z = t_y$. These equations give exactly the lengths and cutoff stated above.

The resulting low-demand surplus is $h(r - \ell) - x(\overline{c} - \ell)$ in the first and third families. In the second family it is

$$(h-x)(r-p)-x(\bar{c}-r)-\gamma x m_x = ha - x(\bar{c}-p+\gamma m_x).$$

Thus target optimality is exactly $S_x \geq 0$ and $(1-\beta)S_x = 0$. The relative responses and all sender comparisons are unchanged by the common scale β.

Conversely, each family's conditions give the stated envelope order, positive mass at every demand, and optimal target responses under Bayes' rule. The arbitrary measurable attacker assignments preserve both sender indifference and the required target surplus. Each type strictly exceeds its resistance payoff because some positive demand receives positive compliance. Weak types have the common low-demand payoff

$$w = \beta x - (1-\beta)a.$$

Apply Lemma 1 with $s = x$ and $\lambda_s = \beta$ to complete every such profile. Its beliefs deter all unused demands without changing on-path behavior, including at target-mixing boundaries and with noninterval attacker sets. Necessity holds for every PBE, and sufficiency supplies a completion satisfying the Intuitive Criterion.

The existence region. It remains to establish why all three families require $t_h > 0$ and why this condition suffices. For the first two families, define

$$D(\rho) = k - p + \frac{\gamma z b}{h-z} = -\frac{z-\rho h}{1-\rho} + \frac{\gamma z\left(\gamma \frac{z}{\rho} - c_2\right)}{h-z}.$$

On $\rho_{\min} < \rho < \frac{z}{h}$,

$$D'(\rho) = \frac{h-z}{(1-\rho)^2} - \frac{\gamma^2 z^2}{\rho^2(h-z)} > 0.$$

All denominators are positive. The derivative inequality is equivalent to $\rho > \rho_{\min}$. As ρ approaches the open upper endpoint from below,

$$\lim_{\rho \to \frac{z}{h}} D(\rho) = \frac{\gamma z t_h}{h-z}.$$

The first family requires $D = \ell - p > 0$. The second requires

$$D = \frac{\gamma z(b-m_z)}{h-z} \geq 0.$$

Strict monotonicity and the open upper endpoint therefore imply $t_h > 0$ in both families, including the second family's $m_x = 0$ boundary. In the third family,

$$0 < m_z = t_h - \frac{(1-\gamma)^2(h-y)}{\gamma} \leq t_h,$$

so the same condition is necessary.

For sufficiency, suppose $t_h > 0$. Choose $z > 0$ small enough that $z < h$ and $\gamma z \frac{t_h}{h-z} < a$. Choose $\rho < \frac{z}{h}$ sufficiently close to $\frac{z}{h}$, while exceeding $\rho_{\min}$. Then $0 < D < a$, so the first family's cutoffs satisfy $p < \ell < r$. Choose

$$0 < x < \min\left\{z, \frac{h(r-\ell)}{\bar{c}-\ell}\right\}, \quad \sigma = \frac{x+\ell-p}{z+\ell-p}, \quad \beta = 1.$$

All first-family conditions hold, with strict low-demand acceptance and strict sender preferences away from the cutoffs. The completion above gives an equilibrium satisfying the Intuitive Criterion.

Both tie families are nonempty in the same region. For the second, take ρ close enough to $\frac{z}{h}$ from below that $D > 0$, yielding $0 < m_z < b$, and choose $x > 0$ sufficiently small for strict low acceptance. For the third, choose $y < h$ sufficiently close to h that $m_z > 0$, then choose $z > 0$ small enough that $\ell < r$ and $x > 0$ small enough for strict low acceptance. Both tied attacker groups can have positive mass. No three-demand equilibrium exists at $t_h = 0$. This proves the characterization and the exact existence condition. □

D.3. Higher Demand Counts in the Existence Map

Under the maintained assumptions, uniform prior, and deterministic demand choices, an equilibrium with a finite set of three or more strictly positive accepted demands, each chosen by positive prior mass, exists under exposure if and only if $\gamma h > c_2$. No such equilibrium exists under protection. Thus the gray region in panel (b) of Figure 2 is exactly the region above the curve; the curve itself is excluded. The accepted-demand set in this statement contains no zero demand.

Proof. The protected impossibility follows from Lemma 2 in SI D.1. Exposed sufficiency follows from the three-demand construction in SI D.2.4. For necessity, we first bound the number of positive accepted demands. Acceptance probabilities strictly decrease with demand, and weak coercers strictly prefer the lowest demand. Each higher accepted demand below h must therefore contain both attackers and resolved settlers to make the target indifferent. Because it attracts attackers, its acceptance-weighted amount must maximize $\lambda_d d$. Denote this maximum by M. The resolved settlement gain at any such demand is

$$M + \lambda_d(c - p).$$

Among demands sharing M, only the largest can maximize this expression below p, and only the smallest can maximize it above p. A demand strictly between them cannot attract a positive mass of settlers. Thus at most two demands above the lowest can lie below h. Including the lowest and possibly h gives at most four positive accepted demands.

Necessity for exactly three demands is established in SI D.2.4. If four demands are used, they must be $x < z < y < h$, with h the fourth amount. The middle products satisfy $\lambda_z z = \lambda_y y = M$. Both endpoint products must be strictly smaller; equality at either endpoint would deprive one middle demand of its required settler group. Write $\rho = \frac{\lambda_h}{\lambda_y}$. The highest demand attracts settlers on $[b, k)$, the upper middle demand attracts settlers on $[k, p)$, and the lower middle demand attracts settlers above p, where

$$b = \frac{\gamma y}{\rho} - c_2, \quad k = \frac{p - y + \rho c_2}{1 - \rho}, \quad \frac{\gamma y}{h - (1 - \gamma)y} < \rho < \frac{y}{h}.$$

The two middle demands divide all attackers in $[0, b)$. Let their positive lengths be m_z and m_y, so $m_z + m_y = b$. Target indifference after y requires $(h - y)(p - k) = \gamma y m_y$. Consequently,

$$k - p + \frac{\gamma y b}{h - y} = \frac{\gamma y m_z}{h - y} > 0.$$

The left side is the function $D(\rho)$ in the preceding existence proof, with y replacing z. It is strictly increasing on the displayed open interval and approaches $\gamma \frac{y(\gamma h - c_2)}{h - y}$ at its upper endpoint. Positivity therefore requires $\gamma h > c_2$. Four demands cannot occur on or below the curve, and five or more are impossible. The existing three-demand construction proves existence above it, completing the claim for the finite positive accepted-demand class. $\square$

E. Two-Demand Comparisons and Information

This section allows any atomless, strictly increasing CDF F with full support on $[0, \bar{c}]$. Assumptions 1–2 otherwise remain in force, and $I(z) = \int_0^z (c + c_2)\,\mathrm{d}F(c)$. Demand sizes and responses may adjust across priors. The attainable-outcome bounds in Proposition 3 retain uniformity.

E.1. Regular Sorting under a General Prior

We prove the two-demand parts of Proposition 4 through Proposition 6 directly for a general prior. The uniform-prior results are special cases.

Lemma 3 (Regular feigned restraint and its matched comparisons survive a general prior): A regular exposed profile has $y = h$, $0 < x < h$, $0 < \rho < 1$, and

$$b = \frac{\gamma x}{\rho} - c_2, \quad k = \frac{p - x + \rho c_2}{1 - \rho}, \quad \max\{0, t_h\} < b < t_x < k < r$$

Its remaining acceptance constraint is

$$x[1 - F(k) + \gamma F(b)] \leq h[F(r) - F(k)]$$

These conditions are necessary and sufficient, with $\beta = 1$ at strict inequality and any $\beta \in (0, 1]$ at equality. Such profiles exist for every admissible $\gamma > 0$. At $\beta = 1$, the same demands and responses support a protected monotone counterpart. Both profiles satisfy the Intuitive Criterion, and their outcome differences are

$$\Delta P_C = (1 - \rho)F(b) > 0, \quad \Delta T = -\rho h F(b) < 0$$

$$\Delta q = \rho F(b) > 0, \quad \Delta K = \rho I(b) > 0$$

$$\Delta U = \rho \int_0^b (b - c)\,\mathrm{d}F(c) > 0, \quad \Delta U_2 = -\rho(b + c_2)F(b) < 0, \quad \Delta V = -\rho I(b) < 0$$

Peaceful concessions per compliance, Ω, also fall.

Proof. Demand incentives are pointwise in the cost. The ratio $\frac{G_E(x,c)}{G_E(h,c)}$ is constant, decreasing, increasing, and constant on the branches displayed in the feigned-restraint characterization. Its two strict crossings with ρ determine the same b and k under every prior. The high-demand group consists of resolved settlers. Since the target mixes there, its indifference requires $y = h$. In the low-demand group, attackers below b generate an acceptance loss γx, resolved settlers between k and r generate a gain $h - x$, and weak settlers above r generate a loss x. The unnormalized acceptance surplus is therefore

$$S_F = h[F(r) - F(k)] - x[1 - F(k) + \gamma F(b)]$$

This proves the characterization and the response restrictions.

For existence, choose a fixed $b_0 \in (\max\{0, \gamma h - c_2\}, p)$ and set

$$\rho(x) = \frac{\gamma x}{b_0 + c_2}, \quad k(x) = \frac{p - x + \rho(x)c_2}{1 - \rho(x)}$$

As x tends to zero from above, $\rho(x)$ tends to zero, t_x and $k(x)$ tend to p, and

$$k(x) - t_x = \frac{\rho(x)(t_x - b_0)}{1 - \rho(x)} > 0$$

for all sufficiently small positive x. All strict cutoff restrictions then hold. Continuity and full support give

$$S_F \to h[F(r) - F(p)] > 0$$

so the low demand is surely accepted. No density or hazard-rate restriction is required.

Under protection, the same h, x, ρ have a single sender crossing at k. The low-demand acceptance surplus becomes $S_F + \gamma x F(b) > 0$, while high-demand indifference is unchanged. This proves the matched counterpart at $\beta = 1$. An exposed profile with $\beta < 1$ cannot be matched at the same responses, because the protected target strictly prefers low-demand compliance.

Lemma 1 uses pointwise sender incentives rather than uniformity. Its completion therefore applies to both profiles with posteriors obtained by conditioning F on their demand groups.

Only types below b change behavior in the matched pair. Under protection, they demand h, settle with probability ρ, and fight otherwise. Under exposure, they obtain x certainly and attack. Their changes in compliance, peaceful transfer, and war probability are $1 - \rho$, $-\rho h$, and ρ. Their payoff change is

$$p + \gamma x - c - [p - c + \rho(c + c_2)] = \rho(b - c)$$

using $\gamma x = \rho(b + c_2)$. The target loses γx for each such type. Integrating proves every displayed identity. Since $T_P > 0$, $\Delta T < 0$, and $\Delta P_C > 0$,

$$\Delta\Omega = \frac{\Delta T - \Omega_P \Delta P_C}{P_{C,E}} < 0$$

$\square$

E.2. What Demand Signals Reveal

We next compare what the two demand signals reveal about whether the coercer would fight if refused. This separates information about that binary decision from information about the coercer's exact war cost.

> Lemma 4 (Exposure reduces information about binary resistance resolve): Fix F and a matched regular pair from Lemma 3. For $Z = \mathbb{1}[c < r]$, the exposed demand signal is a strict Blackwell garbling of the protected signal. Relabel a protected high demand as low with probability $\frac{F(b)}{F(k)}$ and leave a protected low demand unchanged. The two demand experiments are Blackwell incomparable about the full cost c.

Proof. Use the comparison of experiments in Blackwell (1951). Let $A = F(k)$, $B = F(b)$, and $\pi = F(r)$, so $0 < B < A < \pi < 1$. Conditional on $Z = 0$, both games always produce the low demand. Conditional on $Z = 1$, their high-demand probabilities are $\frac{A}{\pi}$ and $\frac{A-B}{\pi}$. The stated kernel reduces the first probability to the second without conditioning on the realized binary state. A reverse kernel must keep an exposed low demand low, because it is the only signal in state $Z = 0$. It can therefore produce a high demand in state $Z = 1$ with probability at most $\frac{A-B}{\pi} < \frac{A}{\pi}$. The binary order is strict.

For the full cost, protection partitions the support into $[0, k)$ and $[k, \bar{c}]$, while exposure partitions it into $[b, k)$ and $[0, b) \cup [k, \bar{c}]$. A garbling of the protected high signal cannot send it to low for $c < b$ and to high for $b < c < k$. A reverse garbling of the exposed low signal cannot send it to high for $c < b$ and to low for $c > k$. All three intervals have positive prior mass, so neither garbling exists. Equivalently, protection identifies $\mathbb{1}[c < k]$ perfectly, whereas exposure identifies $\mathbb{1}[b < c < k]$ perfectly; each classification has positive error under the other experiment. $\square$

The binary comparison averages over the fixed prior's costs within each resolve state. Its kernel depends on $\frac{F(b)}{F(k)}$ and is not a common kernel for arbitrary changes in that within-state distribution. Nor does the result rank information about willingness to stop after compliance. That decision depends on the demand and attack technology. The target's actual payoffs also depend on these features, so binary informativeness alone does not sign its equilibrium payoff.

F. Bargaining Adjustment and Attainable Outcomes

This section proves Proposition 3 without matching demands or acceptance probabilities across games. It gives exact marginal ranges for war, conflict costs, the coercer's payoff, and total welfare, and bounds for both peaceful-concession measures. The result maintains the uniform prior and pure demand strategies. The bounds cover every finite set of strictly positive on-path accepted demands, each chosen by positive prior mass, including all three-demand equilibria in Proposition 2. Accepted zero demands are excluded. A marginal range describes one outcome at a time; the ranges do not specify which combinations are jointly attainable.

Use g_j, $t_j(x)$, $\Psi_j(x)$, and the maximal demands x_j from SI B. The favorable pooling endpoints are

$$q_j^* = F\big(t_j\big), \quad K_j^* = I\big(t_j\big), \quad U_j^* = hF(r) - I\big(t_j\big), \quad V_j^* = 1 - I\big(t_j\big)$$

where $t_j = t_j\big(x_j\big)$. The resistance endpoints q_R, K_R, U_R, V_R are defined in SI B.2.

Within the class of pure demand strategies whose on-path accepted demands form a finite set of at least two strictly positive amounts, each chosen by positive prior mass, the exact marginal ranges in game j are

$$q \in \big(q_j^*, q_R\big), \quad K \in \big(K_j^*, K_R\big), \quad U \in \big(U_R, U_j^*\big), \quad V \in \big(V_R, V_j^*\big)$$

Every value in each interval is already attained by a two-demand equilibrium, and no endpoint is attained within the larger finite-demand class. Adding the characterized pooling and all-resisted equilibria closes the four intervals. Because $t_E > t_P$, each exposed interval is a strict subset of its protected counterpart. This is a statement about attainable outcomes, without a rule selecting an equilibrium.

F.1. Equilibria in which every demand is resisted

Both games admit equilibria in which every demand is resisted. Map types injectively into demands in $(h, 1]$ and let the target resist after inferring the type. If compliance would lead to attack, resistance avoids the additional loss γx. If compliance would lead to settlement, a demand above h costs more than ordinary war. Types above r back down after resistance. Every type receives its resistance continuation value, so none gains by mimicking another resisted signal. Apply Lemma 1 with $\lambda_s = 0$ and $w = -a$. After every unused demand, including zero, assign belief to a weak type and resist. At $d \leq h$, weak types could strictly gain from acceptance and survive the criterion. At $d > h$, every type survives because refusal is the only rational response and gives its equilibrium payoff. Thus these equilibria also satisfy the Intuitive Criterion in both games.

F.2. Proof of the attainable-outcome bounds

Proof. Consider any equilibrium in the stated finite class with at least two accepted demands. Under protection, Lemma 2 limits this to two; the high demand h gives all positive-cost resolved types a positive compliance gain, so surely resisted choices have zero prior mass. Under exposure, every type strictly gains from some positive accepted demand, which also excludes surely resisted choices. In both games, acceptance probabilities strictly decrease with demand and weak types strictly prefer the smallest demand, as proved in SI D.1 and SI D.2.4.

Let x be that smallest demand and β its acceptance probability. First take $\beta = 1$. Every weak type obtains x for certain, so backing down never occurs.

Let $v(c)$ be the coercer's equilibrium payoff and $w(c)$ its probability of war. Each demand's continuation payoff is convex and piecewise affine in the cost. Their finite maximum v is absolutely continuous, and sender optimality gives $v'(c) = -w(c)$ almost everywhere. Coincident payoff segments have the same slope, so arbitrary assignments of indifferent types preserve this identity. Since $v(\bar{c}) = x$, integration under the uniform prior gives

$$\bar{c}q = v(0) - x, \quad U = x + \frac{1}{\bar{c}} \int_0^{\bar{c}} c w(c)\,\mathrm{d}c = x + K - c_2 q$$

Under protection, the zero-cost type can guarantee resistance and war by demanding more than h, so $v(0) \geq p$. Under exposure, it can imitate the surely accepted low demand and then attack, so $v(0) \geq p + \gamma x$. Thus, writing $z = \bar{c}q$,

$$z \geq t_j(x)$$

At a fixed war probability, conflict costs are minimized by concentrating war among the lowest-cost types. Specifically,

$$\bar{c}[K - I(z)] = \int_0^{\bar{c}} (c - z)[w(c) - \mathbb{1}[c < z]]\,\mathrm{d}c \geq 0$$

Both factors have the same sign below and above z. Equality requires $w(c) = \mathbb{1}[c < z]$ almost everywhere. Every demand above x has an acceptance probability $\lambda_d \in (0, 1)$ and a positive mass of resolved settlers: a pool containing only attackers and weak types would be resisted. These settlers fight with probability $1 - \lambda_d \in (0, 1)$, so equality is impossible and $K > I(z)$.

The target can resist every demand, giving

$$U_2 \geq 1 - hF(r)$$

Every coercer can imitate the low demand and obtain at least x. Together with $U + U_2 = 1 - K$, this implies $x \leq h\frac{r}{\bar{c}} < p$, and hence $t_j(x) > 0$. Combining the envelope identity, the target's option to resist, and the strict conflict-cost bound yields

$$hr \geq \bar{c}(U + K) = \bar{c}x + 2\bar{c}K - c_2 z > \bar{c}x + z^2 + c_2 z \geq \bar{c}x + t_j(x)^2 + c_2 t_j(x)$$

Therefore $\Psi_j(x) > 0$. Assumption 1 makes Ψ_j strictly decreasing on the nonnegative axis, so $x < x_j$. It follows that

$$q \geq \frac{t_j(x)}{\bar{c}} > q_j^*, \quad K > I(\bar{c}q) > K_j^*, \quad U \leq hF(r) - K < U_j^*, \quad V = 1 - K < V_j^*$$

War occurs only among resolved types, and high-demand settlers avoid it with positive probability. Hence $q < q_R$ and $K < K_R$. Every type obtains at least its resistance payoff, and the positive mass of weak types strictly gains, so $U > U_R$. Total welfare also strictly exceeds V_R.

If $\beta < 1$, the target is indifferent after every accepted demand. Dividing all acceptance probabilities by β preserves target best responses and multiplies all sender gains relative to resistance by the same positive factor. The resulting profile satisfies Lemma 1. For each outcome $Y \in \{q, K, U, V\}$, the original profile satisfies

$$Y(\beta) = \beta Y(1) + (1-\beta)Y_R$$

Its outcome lies strictly between the pure-low-response outcome and universal resistance. This proves all strict bounds, including mixed low-demand responses.

To establish sharpness and the entire marginal ranges, construct two-demand equilibria with low-demand indifference approaching the favorable pooling endpoint. Under protection, choose $k > t_P$ approaching t_P, and set

$$x = \frac{h(r-k)}{\bar{c}-k}, \quad \rho = \frac{x-p+k}{k+c_2}, \quad y = h$$

At the limit, $x = x_P$ and $\rho = 0$. Multiplying the numerator of ρ by $\bar{c} - k$ gives

$$H(k) = hr - p\bar{c} + (\bar{c} - c_2)k - k^2$$

Here $H(t_P) = 0$ and $H'(t_P) = \bar{c} - c_2 - 2t_P > 0$. Thus for k sufficiently close to t_P, the protected characterization holds with $0 < \rho < 1$. These equilibria approach the maximal pool, while their positive high-demand acceptance excludes the limit from the two-demand class.

Under exposure, choose $x < x_E$ approaching x_E, put $t = t_E(x)$ and $s = \Psi_E(x) > 0$, and set

$$\delta = \frac{-s + \sqrt{s^2 + 4\gamma x s(h-x)}}{2\gamma x}, \quad b = t - \delta, \quad \rho = \frac{\gamma x}{b + c_2}, \quad k = t + \frac{\rho\delta}{1-\rho}, \quad y = h$$

The root satisfies $\gamma x \delta^2 = s(h - x - \delta)$ and tends to zero. Since $t_E > \max\{0, t_h\}$ and $t_E < r$, all strict regular incentive restrictions hold sufficiently close to the limit. The target's low-demand surplus is

$$h(r-k) - x(\bar{c} - k + \gamma b) = \Psi_E(x) - (t-b)(k-t) = s - \frac{\rho\delta^2}{1-\rho} = 0$$

The last equality uses $\frac{\rho}{1-\rho} = \gamma \frac{x}{h-x-\delta}$. Thus these are regular equilibria with low-demand indifference. At $\beta = 1$,

$$\bar{c}q = t, \quad K = I(t) + \frac{\rho\delta^2}{2\bar{c}(1-\rho)} \to I(t_E)$$

Target indifference gives $U_2 = 1 - hF(r)$, so U and V also approach their pooling endpoints. The high-demand group has positive mass along the sequence but vanishes in the limit. Every finite member satisfies the Intuitive Criterion by Lemma 1.

At each sequence member, any $\beta \in (0, 1]$ is admissible. Varying β fills the interval between its pure-low-response outcome and universal resistance for each scalar outcome. Taking sequence members sufficiently close to the pooling limit reaches every interior value in each stated interval.

Finally, pure pooling war and conflict costs decrease as the accepted demand increases, whereas $U(x) = x + \frac{t_j(x)^2}{2\bar{c}}$ has derivative $1 - \frac{(1-g_j)t_j(x)}{\bar{c}} > 0$. Mixed pooling responses combine these outcomes with universal resistance. Thus no characterized pooling equilibrium exceeds the favorable endpoint or the opposite resistance endpoint. Maximal pooling with certain compliance and all-resisted equilibria attain them, respectively. This closes the four marginal ranges and proves the common extremal comparison. □

F.3. Bounds on Peaceful Concessions

Under the same conditions, protected maximal pooling with certain acceptance attains the largest expected peaceful transfer and coercive ability:

$$T_P^* = x_P[1 - F(t_P)] = h[F(r) - F(t_P)], \quad \Omega_P^* = x_P.$$

Every exposed equilibrium satisfies the first bound below; the second requires positive compliance:

$$T_E \leq h[F(r) - F(t_E)] < T_P^*,$$

$$\Omega_E \leq \frac{h[F(r) - F(t_E)]}{1 - F(t_E)} < \Omega_P^*.$$

The protected bounds are attained. The exposed bounds need not be sharp and do not imply that exposed maximal pooling maximizes either concession measure.

Proof. We use target participation and the preceding war bounds under Assumptions 1–2. Define $M_j = E[D\mathbb{1}[C \cap W]] \geq 0$, the expected concession followed by attack. It is zero under protection. The target can resist every demand, so

$$U_{2,j} = 1 - T - hq - g_j M_j \geq 1 - hF(r), \quad T \leq h[F(r) - q] \leq h\big[F(r) - F\big(t_j\big)\big].$$

Protected maximal pooling attains equality. Since $t_E > t_P$, the exposed upper bound is strictly below T_P^*. All-resisted equilibria have $T = 0$ and satisfy the bound as well.

For coercive ability, take an equilibrium with positive compliance and let β be the largest acceptance probability among its on-path accepted demands. If $\beta < 1$, the target is indifferent after every accepted demand. Dividing those acceptance probabilities by β preserves target optimality and multiplies each coercer's gains over resistance by the same factor. Demand choices remain optimal, and Lemma 1 supplies the off-path responses that complete the resulting equilibrium. Both T and P_C are divided by β, leaving Ω unchanged.

In this normalized equilibrium, weak types choose the smallest demand receiving compliance, which is now accepted with certainty. Backing down therefore does not occur. Settlement requires compliance, giving $P_C \geq 1 - q > 0$. Applying the war bound to the normalized equilibrium,

$$\Omega \leq \frac{T}{1-q} \leq \frac{h[F(r) - q]}{1-q} \leq \frac{h\big[F(r) - F\big(t_j\big)\big]}{1 - F\big(t_j\big)}.$$

The function $\frac{h[F(r)-q]}{1-q}$ has derivative $\frac{h[F(r)-1]}{(1-q)^2} < 0$. Protected maximal pooling attains its bound x_P, while $t_E > t_P$ makes the exposed bound strictly smaller. This proves both concession comparisons without matching bargaining terms across games.

Protected pooling at x_P with varying acceptance supplies every $T \in (0, T_P^*]$; all-resisted play supplies zero. Protected pooling at demands $x \in (0, x_P]$ supplies every $\Omega \in (0, x_P]$. Thus every exposed value of either concession measure is also attainable under protection, while a nonempty interval below each protected maximum is unavailable under exposure. This inclusion concerns each measure separately. $\square$

F.4. Pairwise Reversals with Certain Low-Demand Acceptance

The nested marginal ranges do not rank arbitrary equilibrium pairs. At every admissible parameter vector, one game's equilibria can approach its favorable endpoint while the other's approach universal resistance. This gives both directions of comparison for q, K, U, V.

Reversals also occur with $\beta = 1$ and strict low-demand acceptance incentives. Set $p = 0.60$, $c_2 = 0.02$, $a = 0.30$, $\gamma = 0.30$, and $\bar{c} = 1.50$, with $y = h = 0.62$ in all four profiles. Consider protected profiles with $(x, k, \rho) = \left(0.01, 0.60, \frac{1}{62}\right)$ and $\left(0.30, 0.32, \frac{1}{17}\right)$. Their unnormalized low-demand surpluses are 0.177 and 0.0056. The first exposed profile has $x = \rho = 0.20$, $b = 0.28$, and $k = 0.505$, with low-demand surplus $0.0291 > 0$. The second has $x = 0.01$, $\rho = \frac{3}{220}$, $b = 0.20$, and $k = \frac{6493}{10850}$, also with strictly positive low-demand surplus. All four satisfy their characterization conditions.

Comparing the interior exposed profile with the first protected profile gives

$$\Delta P_C = \frac{212}{775} > 0, \quad \Delta T = \frac{2119}{15000} > 0, \quad \Delta q = -\frac{202}{2325} < 0$$

$$\Delta K = -\frac{43309}{930000} < 0, \quad \Delta U = \frac{135007}{930000} > 0, \quad \Delta V = \frac{43309}{930000} > 0$$

Comparing the second exposed profile with the second protected profile reverses all six signs, and the sign of $\Delta\Omega$ reverses as well. Strict inequalities and continuity preserve these comparisons in neighborhoods of the primitives and the independently chosen menus. Thus the reversals do not require low-demand compliance to approach zero.

G. Zero Exposure with an Attack Option

Set $\gamma = 0$ while retaining the attack option, and allow the general prior specified in SI E. We compare the outcomes at each game's largest pooling demand when the target accepts with the same probability.

Lemma 5 (Zero exposure preserves maximal-pooling material outcomes): Set $\gamma = 0$ while retaining the exposed game's attack option. Let $t(x) = \max\{p - x, 0\}$ and

$$\Psi_F(x) = h[F(r) - F(t(x))] - x[1 - F(t(x))]$$

Both games have the same maximal pooling demand

$$x^* = \max\{x \in [0,1] : \Psi_F(x) \geq 0\} \in (0,h), \quad t^* = t(x^*)$$

with $\Psi_F(x^*) = 0$. Every common compliance rate $\lambda \in (0,1]$ is supportable at this demand. Terminal settlement, war, and backing-down probabilities and payoffs agree type by type almost surely. Consequently $(T,q,K,U,U_2,V)_E = (T,q,K,U,U_2,V)_P$.

In the exposed game, any prior mass $m \in [0, F(t^*)]$ of attacking types may join the accepted pool. Thus

$$P_{C,P} = \lambda[1 - F(t^*)], \quad P_{C,E} = \lambda[1 - F(t^*) + m]$$

$$\Omega_P = x^*, \quad \Omega_E = x^* \frac{1 - F(t^*)}{1 - F(t^*) + m}$$

A strict gap in Ω occurs exactly when $m > 0$. Regular feigned restraint with strict demand preferences disappears at $\gamma = 0$.

Proof. Under protection, types above $t(x)$ strictly prefer an accepted positive demand to resistance, and types below it prefer resistance. Under exposure at $\gamma = 0$, the former still join the accepted pool. The latter earn $p - c$ from either acceptance followed by attack or resistance followed by war. They are indifferent about joining. Full support and atomlessness permit any accepted-attacker mass between zero and $F(t(x))$.

An accepted attacker gives the target identical compliance and resistance payoffs. The remaining accepted types generate total acceptance surplus $\Psi_F(x)$ in either game. The different posterior denominators do not change its sign. Thus $\Psi_F(x) \geq 0$ characterizes demands feasible at some positive compliance rate. Strict inequality requires $\lambda = 1$; equality permits every positive $\lambda \leq 1$. Assign excluded attackers to demand 1, which is surely resisted because $h < 1$. On-path incentives hold as stated. The completion in Lemma 1 extends to zero exposure by retaining attacking types that are indifferent in its survivor set: for $d < w$, attacking at an unused demand yields $R(c) \leq v(c)$, so no deviation is profitable even when dominance is weak. The remaining off-path cases are unchanged.

Continuity gives $\Psi_F(0) = h[F(r) - F(p)] > 0$ and $\Psi_F(h) = h[F(r) - 1] < 0$. No $x \geq h$ is feasible. The largest feasible demand therefore exists in $(0, h)$ and binds the target's constraint. This proves common maximality and independence from λ without requiring a unique positive root.

At x^*, types below t^* fight surely in both games. Above t^*, the common acceptance rate gives the same settlement, war, and backing-down lotteries. Since the concession is not retained separately after war, changing the history from resistance and war to compliance and attack changes no material payoff at $\gamma = 0$. This proves typewise equivalence and the aggregate equality. Counting accepted attackers adds λm to compliance without changing the peaceful transfer $\lambda x^*[1 - F(t^*)]$, which gives the formulas for Ω.

Finally, in regular feigned restraint the high-demand pool consists of resolved settlers, so target mixing requires $y = h$. A positive-cost low-demand attacker gains zero from compliance with x but gains $c + c_2 > 0$ from compliance with h. Positive high-demand compliance therefore rules out its strict preference for the low demand. The cutoff equation $\rho(b + c_2) = \gamma x = 0$ cannot hold with $\rho, b > 0$. $\square$

This result concerns maximal pooling, not the full equilibrium correspondence. Indifferent attackers can still be assigned to multiple demands at zero exposure.