Learning Before Crisis Bargaining*

Michael Gibilisco† Alexander Hirsch‡ Sota Minowa§

April 25, 2026

Abstract

States often make costly investments to privately learn about their military capabilities via secret military exercises, weapons tests, or proxy conflicts. We study how such learning creates "endogenous" asymmetric information that warps crisis bargaining. We consider the canonical crisis-bargaining setup, but allow the receiver to pay a cost to privately learn his military strength before bargaining. We show that the sender's attempts to extract surplus actually incentivizes the receiver to privately learn about his military strength, sometimes leading to bargaining breakdown. An informed receiver may elicit both more and less generous offers than an uninformed one, depending on the ratio of ex-ante uncertainty about his strength to the aggregate costs of war. When the receiver's learning costs are moderate, the unique equilibrium exhibits uncertainty about whether he has privately learned his strength; the proposer might therefore "accidentally" trigger war by making an unacceptable offer to an uninformed receiver. Finally, moderate learning costs generate the highest risk of conflict, despite exhibiting less learning and asymmetric information than smaller costs. Broadly, states' incentives to create asymmetric information via private learning significantly complicates the relationship between conflict fundamentals, asymmetric information, and war.

*This project has benefited from audience feedback at the Caltech Applied Political Theory Lab, the 2026 Park City Political Economy Conference, and seminar audiences at Vanderbilt University.

†Division of Humanities & Social Sciences, Caltech. Email: michael.gibilisco@caltech.edu.
‡Division of Humanities & Social Sciences, Caltech. Email: avhirsch@caltech.edu.
§Division of Humanities & Social Sciences, Caltech. Email: sminowa@caltech.edu.

1 Introduction

Canonical theories in International Relations emphasize information asymmetries as a cause of war. However, they typically treat such asymmetries as exogenous. In reality, information asymmetries are often the result of deliberate decisions by states themselves, since states can make costly investments to learn their military capabilities, conflict readiness, and domestic political costs of war. For example, Russia and NATO routinely conduct military exercises; one reason for Russia's Zapad 1999 exercises was to assess how aggressive Russian foreign policy should be (Luzin 2017). Likewise, American presidential administrations have privately tried to gauge the domestic political costs of war before entering foreign conflicts such as the Gulf War in 1991 and the invasion of Iraq in 2003.[1]

The fact that asymmetric information can be the result of deliberate decisions by state actors implies that its relationship with conflict, and even the relationship between seemingly well-understood fundamentals (like the cost of war) and conflict, is far from straightforward. For example, suppose that asymmetric information leads to more war ceteris paribus, but is also likelier to be "chosen" when the fundamentals disfavor war. If wars become likelier overall due to the asymmetric information, then the empirical relationship between the fundamentals and war will be reversed. Conversely, if wars become less likely overall due to the fundamentals, then the empirical relationship between asymmetric information and war will be reversed. Despite this complexity, previous crisis bargaining models have only incorporated endogenous asymmetric information *as a byproduct* of other strategic decisions that affect war fundamentals, such as private arming (e.g. Debs and Monteiro (2014); Meirowitz and Sartori (2008)). Thus, the literature lacks a theoretical account of how different aspects of *learning itself* — its costs, the degree of ex ante uncertainty, and whether states learn about relative capabilities or private costs—can warp crisis bargaining.

In this paper we develop such an account. We show that allowing states to privately learn about their strength fundamentally alters crisis bargaining, and changes seemingly straightforward empirical relationships. Most strikingly, we demonstrate that the relationship between asymmetric information and war is non-monotonic: when private learning about capabilities entails moderate costs the likelihood of war is highest, *even though* the equilibrium probability of learning (and hence of asymmetric information) is only moderate. We also find that the relationship between ex ante uncertainty and war can be non-monotonic, and specifically establish conditions where more uncertainty leads to more peace. Both phenomena emerge because states strategically adjust their offers in anticipation of their adversaries' learning, which feeds back into states' incentives to learn, the *ceteris-paribus* effect of learning on war, and the likelihood of bargaining breakdown.

[1] Of course, these expectations may be incorrect (Kertzer et al. 2025).

The Model We study the canonical crisis-bargaining environment. In the model, a proposer makes a take-it-or-leave-it offer to a receiver, who either accepts the offer or rejects it by going to war. The receiver's payoff from war depends on an initially unknown level of strength. At the beginning of the interaction, the receiver may pay a cost to learn this strength. However, his learning decision is private: when the proposer makes an offer, she does not know whether the receiver has learned his strength, or what that strength is if he has. We also allow the effect of the receiver's strength on the players' relative capabilities (i.e., the likelihood of victory) vs. the receiver's private war costs to vary, in order to capture the extent to the players' conflict payoffs are "common-valued" versus "private-valued" (Fey and Ramsay (2011)). We term this *interdependence*, as it captures the extent to which the receiver's war payoff is correlated with the proposer's war payoff.

We first characterize offer behavior when the proposer believes that the receiver does not know his strength. In this case, she naturally makes an offer that leaves the receiver indifferent between accepting and fighting, thereby extracting all of the bargaining surplus (as in standard ultimatum bargaining). Importantly, however, we also show that by doing so, the proposer inadvertently *maximizes the receiver's incentive to learn his strength.* This is because a receiver who is just barely willing to accept an offer is also the likeliest to have his decision affected by additional information, and the receiver's willingness to acquire information depends on the likelihood it will be useful. Thus, there is a general sense in which "standard" ultimatum bargaining behavior gives rise to asymmetric information.

We next characterize offer behavior when the proposer believes that the receiver knows his strength. We show that the optimal offer may be both more "generous" *or* more "stingy" than the optimal offer for an uninformed receiver, depending on the fundamentals. When uncertainty is low and/or war costs are high, the proposer is more generous to an informed receiver; either because generosity more predictably reduces the chance of war, and/or because war is better to avoid. Conversely, when uncertainty is high and/or the costs of war are low, the proposer is stingier toward an informed receiver; because generosity does not predictably reduce the chance of war, or because war is not as important to avoid. While these relationships are straightforward, they have an important but subtle implication; when the fundamentals lead the proposer to make a stingier offer to an informed receiver, the receiver becoming informed will *reduce* the chance of war, since he will surely reject the offer if uninformed. Thus, the standard relationship between asymmetric information and war can be reversed in the presence of *higher order uncertainty* about whether the receiver is informed, since the proposer may "accidentally" make the optimal offer for an informed receiver to an uninformed one.

We show that such higher order uncertainty is a necessary part of equilibrium when the receiver's costs of learning are *moderate.* The reason is that any shift by the proposer away from the optimal uninformed offer reduces the receiver's willingness to become informed as he becomes more confident that his uninformed decision is correct – that is, more confident that he should accept an increasingly generous offer outright or reject an increasingly stingy one. The receiver must therefore *randomize* over becoming informed when his learning costs are moderate, since the optimal uninformed offer will incentivize him to learn while the optimal informed offer will not.

After analyzing the relationship between the proposer's anticipated offer and the receiver's learning incentives we complete the equilibrium characterization, which describes: (a) a threshold for learning costs to be sufficiently high that the receiver will remain uninformed even anticipating the uninformed offer, so that he is uninformed in equilibrium (b) a threshold for learning costs to be sufficiently low that the receiver will become informed even anticipating the informed offer, so that he becomes informed in equilibrium, and (c) the form of equilibrium when learning costs are moderate, so that the receiver randomizes and there is higher order uncertainty in equilibrium. With higher order uncertainty, the proposer's offer skews away from the optimal uninformed offer toward the optimal informed one, albeit in slightly different ways depending on the fundamentals. Specifically, if uncertainty is low and/or war costs are high she will surely make an offer between the two, whereas if uncertainty is high and/or war costs are low she will randomize between the two.

After characterizing equilibria, we use our theory to explore three foundational questions in the rationalist explanations for war literature. We first establish that the probability of learning, and hence the probability that there is asymmetric information, is decreasing in the cost of learning. Thus, learning costs proxy for the likelihood of endogenous asymmetric information in our model. We next examine the equilibrium relationship between learning costs and the probability of war, and find a non-monotonic relationship; moderate costs yield the highest likelihood of war. When learning costs are large, learning never occurs, information is symmetric, and peace prevails. When learning costs are small, learning surely occurs, and war breaks out if the receiver's actual war payoff exceeds the proposer's offer. When learning costs are moderate, however, the receiver must randomize, and higher learning costs (corresponding to less asymmetric information) can lead to *more* conflict; although in two very differing ways depending on the fundamentals.

The first way occurs if ex-ante uncertainty is small relative to the cost of war, so that the proposer will make a more generous offer to an informed receiver. In this case learning is necessary for war, since an uninformed receiver will accept the receiver's equilibrium offer outright. In equilibrium, higher learning costs simultaneously reduce learning – which decreases the chance of war ceteris paribus – and make the receiver's offer less generous – which increases the chance of war ceteris paribus. The latter occurs because the proposer's equilibrium offer must more strongly incentivize learning when it becomes costlier to maintain higher order uncertainty about learning, and therefore more closely resemble the optimal uninformed offer, i.e., become less generous. The net effect of these competing forces is to increase the chance of war, reversing the standard relationship between asymmetric information and war.

The second way occurs when ex-ante uncertainty is large relative to the cost of war, so that the proposer will make a stingier offer to an informed receiver. When this is the case, the proposer will sometimes make the optimal offer for an informed receiver in equilibrium, and thus "accidentally" trigger war if the receiver remains uninformed. Thus, asymmetric information will *decrease* the chance of war ceteris paribus, since it allows the receiver to learn when to accept the stingier offer. As a result, an equilibrium with moderate learning costs and partial learning can exhibit a higher chance of war than an equilibrium with low learning costs and full learning, again reversing the standard relationship between asymmetric information and war.

We next examine the relationship between interdependence and asymmetric information. We find that as interdependence increases – so that the receiver's private information is less about his own costs and more about relative capabilities – equilibrium learning (and hence asymmetric information) decreases. The reason is that greater interdependence increases the *wedge* between the optimal offer for an uninformed receiver and the optimal offer for an informed one, which in turn drives down the value of learning when the proposer believes the receiver to be informed. The wedge increases because the proposer targets the *marginal* type of the receiver when choosing her offer, so it depends on how strongly she wishes to avoid war against this marginal type. When uncertainty is low the optimal offer is generous, so the marginal type of receiver is stronger than the average type – greater interdependence therefore makes the proposer wish to avert war more, pushing the optimal offer up. When uncertainty is large the optimal offer is stingy, so the marginal type of receiver is weaker than the average type – greater interdependence therefore makes the proposer wish to avert war less, pushing the optimal offer down.

We last examine the relationship between ex-ante uncertainty about the receiver's strength and war. For the most part, we find that greater uncertainty leads to more war, consistent with the conventional wisdom. However, we also establish necessary and sufficient conditions for this relation to be reversed. Specifically, for more uncertainty to lead to more peace, interdependence must be low (i.e., the receiver's uncertainty must be about private values to a sufficient degree), and uncertainty must also be sufficiently high. Under these circumstances, learning can lead to peace in the manner described above: because the receiver sometimes learns that he should accept a stingy offer he would have otherwise rejected. As uncertainty increases, the receiver generally learns more often. When learning costs are not too large, more uncertainty can therefore lead to a greater chance of peace.

Related literature Previous work endogenizing asymmetric information in crisis bargaining largely focuses on private military investments instead of learning. Meirowitz and Sartori (2008) show that when states privately arm before bargaining, they deliberately create ambiguity about their military strength in equilibrium. In a dynamic environment, Debs and Monteiro (2014) demonstrate how unobserved investment decisions can create commitment problems that lead to bargaining failures. We view our paper as complementary to these works, but focus on costly learning rather than capability accumulation. A key difference is that, unlike investment-based models, learning can reveal that a state is weaker than it previously believed. This gives rise to an important feature of our model – that asymmetric information can reduce conflict ceteris paribus because a state, after learning it is weaker than expected, may accept a stingy offer it would have otherwise rejected.

There is also a large empirical literature that attempts to test hypotheses derived from crisis-bargaining models. A key hypothesis of these works is that uncertainty and asymmetric information are positively correlated with war (Kaplow and Gartzke 2015; Reed 2003; Reiter 2003; Smith and Spaniel 2019; Spaniel and Smith 2015; Thyne 2012).[2] Our analysis illustrates that such correlations are difficult to interpret as evidence either for or against bargaining models when asymmetric information is endogenous. Even when learning increases the likelihood of asymmetric information, it can either increase or decrease the probability of war depending on how learning alters bargaining behavior.

Relatedly, Tchaouchev (2025) studies the diffusion of information about gunpowder artillery during the fifteenth century; he shows that early uncertainty about the effectiveness of cannons fueled conflict, but subsequent battlefield demonstrations reduced conflict by aligning the beliefs of states. If the introduction of gunpowder artillery reflected heightened uncertainty around relative military capabilities, then this result can be interpreted as an empirical test of our model; when uncertainty is over relative power more uncertainty leads to a higher equilibrium likelihood of war.

Finally, we contribute to a broader theoretical literature in economics that studies bargaining when parties can learn about the "inside option" of agreement rather than the "outside option" of bargaining breakdown (as in our model). Much of this literature assumes that the proposer has commitment power: in Crémer and Khalil (1992) learning occurs after the proposer makes an offer, while in Crémer, Khalil and Rochet (1998) and Kessler (1998) the proposer commits to an offer as a function of a message from the receiver. In contrast, our proposer cannot commit, so our bargaining setup is more similar to ultimatum bargaining in Roesler and Szentes (2017), Ravid (2020), and Ravid, Roesler and Szentes (2022), with some notable differences. All three works study a continuous learning technology by the receiver, which precludes equilibrium uncertainty about *whether* the receiver is informed, as well as a full characterization of comparative statics in learning costs (in Ravid, Roesler and Szentes (2022)). In Ravid (2020) the proposer also begins privately informed about the receiver, so there is a signalling component. Finally, all three consider a receiver who can become informed about a private value of a good or cost of production,

[2] Some argue that it should be negatively correlated (e.g., Uzonyi and Wells 2016).

while we allow uncertainty to concern both private-value war costs and common-value war outcomes.[3] Our analysis reveals that the distinction is crucial. If uncertainty concerns common relative war capabilities rather than private war costs then more uncertainty leads to more war; the relationship can only reverse when it sufficiently concerns private war costs.

2 Model

Two states engage in crisis bargaining. One state is the proposer, labeled P, and the other is the receiver, labeled R. The policy space is $[0, 1]$, and a policy $x \in [0, 1]$ denotes the amount of territory or good given to state R. States have standard linear preferences over the allocation of territory: $\pi_P(x) = 1 - x$ and $\pi_R(x) = x$.

At the beginning of the interaction, there is an unknown state of world $s \in \left[-\frac{\delta}{2}, \frac{\delta}{2}\right]$ where $\delta > 0$. We interpret s as state R's strength in a war against state P. Specifically, we treat war as a costly lottery in which R wins with probability $\bar{p} + \lambda s$, and P wins with complementary probability $1 - \bar{p} - \lambda s$, where $\lambda \in [0, 1]$. War also entails some costs, where P's costs of war is $\bar{c}_P$ and R's cost of war is $\bar{c}_R - (1 - \lambda)s$. Thus, λ denotes the degree to which the R's unknown strength, s, affects his probability of victory versus his war cost. When $\lambda = 1$, R's strength s only affects his probability of victory and thus P's probability of victory, but when $\lambda = 0$, s only affects R's cost of war.

The strategic interaction proceeds as follows.

1. Nature draws s from $\left[-\frac{\delta}{2}, \frac{\delta}{2}\right]$ according to a cdf F with corresponding density f.
2. State R can privately learn s ($l = 1$) or not ($l = 0$).
3. State P makes a take it or leave it offer to R: $x \in [0, 1]$.
4. R sees the offer x and either accepts ($a = 1$) or rejects ($a = 0$).
    (a) If R accepts, then the game ends with x as the policy outcome.
    (b) If R rejects, then a war ensues and the winner chooses their ideal policy x to be implemented.

In step 2 state R's learning is *private*, in the sense that neither the decision to learn l nor the knowledge from learning s is observed by state P.

The proposer's expected payoffs conditional on (x, a, s) can be written as

$$u_P(x, a, s) = \begin{cases} 1 - x & \text{if } a = 1 \\ 1 - \bar{p} - \lambda s - \bar{c}_P & \text{if } a = 0 \end{cases}. \tag{1}$$

[3] Chatterjee, Dong and Hoshino (2025) consider a model where the proposer knows the receiver's value of the good, so their offer has a signaling component. They study equilibrium as the receiver's cost of learning goes to zero, and provide an example showing that their main results do not hold in common-value settings.

In the case of war, i.e., $a = 0$, two outcomes can arises. With probability $\bar{p} + \lambda s$, R wins and chooses policy $x = 1$, in which case P's payoff is $\pi_P(1) = 0$ minus her cost of war, $\bar{c}_P$. With complementary probability, $1 - \bar{p} - \lambda s$, P wins and chooses policy $x = 0$, in which case P's payoff is $\pi_P(1) = 1$ minus her cost of war, $\bar{c}_P$.

Similarly, the receiver's expected payoffs conditional on (x, a, l, s) are

$$u_R(x, a, l, s) = \begin{cases} x - lk_R & \text{if } a = 1 \\ \bar{p} - \bar{c}_R + s - lk_R & \text{if } a = 0 \end{cases}. \tag{2}$$

In Equation 2, $k_R > 0$ denotes the receiver's cost of learning, i.e., the effort needed to uncover his strength s.

Comparing Equations 1 and 2 illustrates that s is a shock to the receiver's war payoff, and λ governs the *interdependence* between the receiver and the proposer's war payoffs. Specifically, $\lambda > 0$ implies that the payoffs are negatively correlated via state R's probability of victory depending on s. When interdependence is minimal, i.e., $\lambda = 0$, we say the game has *private* values, because state R's strength only affects his own cost of war. When interdependence is maximal, i.e., $\lambda = 1$, we say the game has *common* values, because both states' war payoffs depend equally on s. When $\lambda \in (0, 1)$ the intermediate case holds, and the uncertainty has both common and private value components.

We make two assumptions to keep the analysis tractable and focus on the most interesting cases; they are maintained throughout the subsequent analysis.

Assumption 1. *F is the uniform distribution over $\left[-\frac{\delta}{2}, \frac{\delta}{2}\right]$.*

In light of Assumption 1, we interpret δ as the degree of *ex-ante symmetric uncertainty about the receiver's strength*. Because the probability of state R winning a war is $\bar{p} + \lambda s$ and must be between 0 and 1, we are also implicitly assuming that $\frac{\delta}{2} \leq \frac{1}{\lambda} \min\{\bar{p}, 1 - \bar{p}\}$. Further, because $\bar{p} \in (0, 1)$, a necessary condition for this to hold is $\delta < \frac{1}{\lambda}$.

Assumption 2. *State R's war payoffs are interior: $\frac{\delta}{2} < \bar{p} - \bar{c}_R < 1 - \frac{\delta}{2}$.*

Assumption 2 implies that state R would surely accept if he receives a sufficiently generous offer and surely reject if he receives a sufficiently stingy offer, regardless of whether he is informed.[4] We characterize subgame perfect equilibria under Assumptions 1 and 2.

3 Equilibrium Characterization

It is first useful to explicitly derive the distributions governing the states' war payoffs, w_i, $i = P, R$. To do so it is helpful to define $\underline{w}_R = \bar{p} - \bar{c}_R - \frac{\delta}{2}$ and $\overline{w}_R = \bar{p} - \bar{c}_R + \frac{\delta}{2}$ as the lower

[4] Along with Assumption 1, this implies an upper bound on δ of $\bar{\delta}_\lambda = 2 \min\left\{x_0^*, \frac{1}{\lambda}\bar{p}, 1 - x_0^*, \frac{1}{\lambda}(1 - \bar{p})\right\}$.

and upper bounds on state R's war payoffs, so that $\mathbb{E}[w_R] = \bar{p} - \bar{c}_R$.

Remark 1. *State R's war payoff is a uniformly distributed random variable over the interval $[\underline{w}_R, \bar{w}_R]$ with cdf*

$$G(w_R) = \frac{w_R - \underline{w}_R}{\delta}$$

State P's war payoff is

$$w_P(w_R) = \lambda(1 - w_R) + (1 - \lambda)(1 - \mathbb{E}[w_R]) - (\bar{c}_R + \bar{c}_P)$$

which is also a uniformly distributed random variable.

Remark 1 illustrates why the parameter λ captures the interdependence of realized war payoffs. When $\lambda = 0$, state P's war payoff does not depend on the specific realization of state R's war payoff w_R. As such, state R's private information pertains to private values. As λ increases, the specific realization of state R's war payoff has a larger effect on state P's war payoff, so state R's private information increasingly pertains to common values, i.e., to both states' payoffs.

A strategy for state P is a function $\sigma_P : [\underline{w}_R, \overline{w}_R] \to [0, 1]$ where $\sigma_P(x)$ denotes the probability that state P proposes x. Although state P must sometimes randomize over her proposals in equilibrium, we later show that they propose a finite number policies in any such equilibrium. As such, we sidestep measurability issues when defining her proposal strategy. Further, it is without loss of generality to restrict state P proposals to the interval $[\underline{w}_R, \overline{w}_R]$, since proposing $x > \overline{w}_R$ is strictly dominated by proposing $x = \overline{w}_R$ (as both will always be accepted), whereas proposing $x \leq \underline{w}_R$ is equivalent to proposing $x = \underline{w}_R$ (as both will always be rejected).

A strategy for state R determines both the probability of learning, and which policies will be accepted after each learning decision $l \in \{0, 1\}$. The former is denoted $\sigma_R \in [0, 1]$. For the latter, note that state R's optimal acceptance decisions are straightforward; if state R learns then he accepts if and only if $x \geq w_R$, whereas if he does not learn he accepts if and only if $x \geq \mathbb{E}[w_R]$.[5] In the subsequent analysis and exposition we characterize equilibrium profiles of the form $\sigma = (\sigma_P, \sigma_R)$, and assume throughout that state R is playing an optimal accept/reject decision.

3.1 Pure strategy equilibria

In a pure strategy equilibrium, state R surely learns his strength or not. To characterize such equilibria we first analyze state P's proposals when they are best responding to state

[5] More formally, subgame perfect equilibrium pins down acceptance decisions when the inequalities are strict. When indifferent we assume state R always accepts, since any equilibrium where they do not do so remains an equilibrium when they do.

R's learning decision, and then proceed to analyze R's learning decision.

Proposer's offers. Let $U_P(x|l)$ denote state P's expected utility from offering $x \in [\underline{w}_R, \overline{w}_R]$ given state R's learning decision $l \in \{0, 1\}$. When state R does not learn his own strength ($l = 0$) he will accept any offer weakly above $\mathbb{E}[w_R]$, so

$$U_P(x|l=0) = \begin{cases} 1 - x & \text{if } x \geq \bar{p} - \bar{c}_R \\ 1 - \bar{p} - \bar{c}_P & \text{if } x < \bar{p} - \bar{c}_R \end{cases}. \tag{3}$$

In the standard fashion state P's optimal offer is $x_0^* = \mathbb{E}[w_R] = \bar{p} - \bar{c}_R$; this makes state R indifferent between accepting and war, thereby allowing state P to capture the full bargaining surplus.

When state P expects state R to privately learn his own strength ($l = 1$), state P anticipates that R will accept any offer $x \in [\underline{w}_R, \overline{w}_R]$ that exceeds his actual war payoff ($x \geq x_R$). Thus, state P's expected utility from proposing x is

$$U_P(x|l=1) = \int_{\underline{w}_R}^{x} (1-x) g(w_R) dw_R + \int_{x}^{\overline{w}_R} w_P(w_R) g(w_R) dw_R \tag{4}$$

where $g(w_R) = \frac{1}{\delta}$ is the pdf of state R's war payoff from Remark 1. Differentiating with respect to the offer x yields

$$\frac{\partial U_P}{\partial x}(x|l=1) = ((1-x) - w_P(x)) \cdot g(x) - G(x).$$

Intuitively, the benefit to the proposer of marginally increasing her offer x is the net benefit $(1-x) - w_P(x)$ of avoiding war against a receiver $w_R = x$ just indifferent between accepting and rejecting her offer, *times* the likelihood $g(x)$ that the receiver is said type. The marginal cost is the increased expense when the receiver is a type $w_R \leq x$ who would have accepted the offer regardless (which is the case with probability $G(x)$).

The derivative satisfies a number of intuitive properties. First, it is easily verified that $\frac{\partial U_P}{\partial x}(\underline{w}_R|l=1) > 0$; that is, it is never optimal to make an offer $x \leq \underline{w}_R$ that will be rejected for sure, since the marginal cost of increasing it is 0. Second, increasing the offer yields greater benefits when there is a higher likelihood $g(x)$ that the receiver is exactly the indifferent type x who will be impacted by the increase. Finally, we show in the Appendix that U_P is strictly concave over $(\underline{w}_R, \overline{w}_R)$, so that the unique optimal offer can be characterized using a first-order approach (see Appendix A.2).

Applying Assumption 1, the functional form of the cdf from Remark 1, and bounding within $[\underline{w}_R, \overline{w}_R]$ yields that the optimal offer x_1^* is equal to

$$x_1^* = \min\left\{x_0^* + \left(\frac{1}{2-\lambda}\right)\left(\bar{c}_P + \bar{c}_R - \frac{\delta}{2}\right), \overline{w}_R\right\} \tag{5}$$

It is straightforward to see that the relative magnitude of ex-ante uncertainty δ about R's strength to the aggregative cost of war $\bar{c}_P + \bar{c}_R$ determines whether P will make a higher offer when she believes R to be privately informed ($x_1^* > x_0^*$) vs when she believes R to be uninformed R ($x_1^* < x_0^*$), as follows.

Remark 2. *The learning offer, x_1^*, is smaller than the no-learning offer, x_0^*, if and only if uncertainty is high relative to the aggregate cost of war ($\delta > 2(\bar{c}_P + \bar{c}_R)$).*

When uncertainty is low relative to the cost of war (i.e., $\delta < 2(\bar{c}_P + \bar{c}_R)$), P is relatively confident about how an informed receiver will respond to her offer. As such, P can meaningfully increase the probability of acceptance by marginally increasing x, thereby incentivizing her to make a more generous offer. When uncertainty is high relative to the cost of war (i.e., $\delta > 2(\bar{c}_P + \bar{c}_R)$), P is relatively uncertain about how an informed receiver will respond to her offer. As such, P does not expect a marginal increases in her offer to meaningfully increase the probability of acceptance, disincentivizing her from making a generous offer.

Receiver's learning. State R's willingness to learn more about his own strength depends on the offer x that he anticipates receiving. Let $U_R(l|x)$ denote P's expected payoff from decision $l \in \{0, 1\}$ when anticipating offer x. If state R learns his strength s, then he anticipates that he will accept the offer x if $x \geq w_R = \bar{p} - \bar{c}_R + s$ and go to war otherwise. Thus, state R's expected utility from privately learning his strength is

$$U_R(1|x) = \Pr(w_R \leq x)x + \Pr(w_R > x)\mathbb{E}[w_R|w_R > x] - k_R. \tag{6}$$

If state R does not learn, then he anticipates either accepting x or rejecting and getting their uninformed war payoff of $\mathbb{E}[w_R] = x_0^*$, depending on which is greater. That is,

$$U_R(0|x) = \begin{cases} x_0^* & \text{if } x < x_0^* \\ x & \text{if } x \geq x_0^* \end{cases}. \tag{7}$$

To understand when state R sometimes chooses to learn his own strength, we can subtract the latter from the former:

$$U_R(1|x) - U_R(0|x) = \begin{cases} \Pr(w_R \leq x)\mathbb{E}[x - w_R|w_R \leq x] - k_R & \text{if } x < x_0^* \\ \Pr(w_R > x)\mathbb{E}[w_R - x|w_R > x] - k_R & \text{if } x \geq x_0^* \end{cases}.$$

Recalling that k_R is the cost of learning, what remains is the net benefit of learning, or the *value of information*, to state R.

Specifically, when $x \geq x_0^*$ state R anticipates that he will accept the offer and avoid a fight if he remains uniformed. As such, the benefit of learning derives from the possibility that he will discover *positive* information about his strength indicating that he is stronger than expected and should instead fight. Formally,

$$\phi^+(x) \equiv \Pr(w_R > x)\mathbb{E}[w_R - x | w_R > x].$$

represents the value of positive information about R's strength. Alternatively, when $x < x_0^*$ state R anticipates that he will reject the offer and fight if he remains uninformed. Thus, the benefit of learning derives from the possibility of discovering *negative* information about his strength indicating that he is weaker than expected and should instead avoid a fight by accepting. Formally,

$$\phi^-(x) \equiv \Pr(w_R \leq x)\mathbb{E}[x - w_R | w_R \leq x].$$

represents the value of negative information about R's strength.

Notice that ϕ^- is strictly increasing in the anticipated proposal x; the intuition is two-fold. First, as the expected offer x increases, there is a higher probability that accepting it is the optimal decision (in the ex post sense) for state R instead of war. Second, conditional on acceptance being optimal, the net benefit of doing so is larger (since the offer is stronger). For similar reasons ϕ^+ is strictly decreasing in x, as it becomes increasingly likely that learning his strength will not affect his decision to reject. Finally, the two functions intersect at x_0^* because $\phi^-(x) - \phi^+(x) = \Pr(w_R \leq x)\mathbb{E}[x - w_R | w_R \leq x] + \Pr(w_R > x)\mathbb{E}[x - w_R | w_R > x] = x - \mathbb{E}[w_R] = x - x_0^*$; thus, the value of learning can be written as

$$\phi(x) = \min\{\phi^-(x), \phi^+(x)\}.$$

This achieves a maximum at x_0^*, which is the offer that leaves an uniformed state R indifferent between accepting and going to war. The following result summarizes this discussion.

Lemma 1. *For every proposal $x \in [\underline{w}_R, \overline{w}_R]$, the value of learning ϕ satisfies*

$$\phi(x) = \min\left\{\frac{(x - \underline{w}_R)^2}{2\delta}, \frac{(\overline{w}_R - x)^2}{2\delta}\right\},$$

where $\phi(x)$ is strictly increasing (decreasing) in x when $x < (>) x_0^*$.

The lemma illustrates our first key insight; that the proposer's optimal offer to an uninformed receiver necessarily *maximizes his willingness to become informed.* The reason is that the receiver's incentive to learn is driven by the likelihood that he expects additional information to be decision-relevant, which is most likely precisely when he anticipates an offer that will make him indifferent over acceptance. Thus, there is a sense in which optimal behavior in ultimatum bargaining naturally gives rise to asymmetric information.

Finally, combining this insight with our analysis of state P's proposals allows us to characterize pure-strategy equilibria.

Proposition 1. *Assume $x_1^* \neq x_0^*$. Then $\phi(x_1^*) < \phi(x_0^*)$, and a pure-strategy equilibrium exists if and only if $x \notin (\phi(x_1^*), \phi(x_0^*))$, which takes the following form.*

1. *An equilibrium in which state R never learns exists if and only if $k_R \geq \phi(x_0^*)$. In this equilibrium, state P proposes x_0^*.*

2. *An equilibrium in which state R surely learns exists if and only if $k_R \leq \phi(x_1^*)$. In this equilibrium, state P proposes x_1^*.*

Alternatively, if $x \notin (\phi(x_1^*), \phi(x_0^*))$ then in equilibrium state R must randomize over becoming informed.

3.2 Mixed strategy equilibria

The previous results demonstrate that for a range of intermediate learning costs ($k_R \in (\phi(x_1^*), \phi(x_0^*))$) the receiver must necessarily *randomize* over becoming informed; the reason is that the optimal uninformed offer is sufficiently aggressive to incentivize learning, while the optimal informed offer is either too generous or too stingy. Thus, uncertainty about *whether* the receiver is better informed arises naturally in our model. Equilibrium for intermediate learning costs further requires that state R anticipate offers from state P that make him *just indifferent* over privately learning his strength, so that he will randomize. This in turn implies that P's equilibrium offer behavior will depend both on ex-ante uncertainty δ and on aggregate war costs $\bar{c}_P + \bar{c}_R$.

To see why, suppose that state P believes state R to have privately learned his strength with probability $\sigma_R \in (0, 1)$. Then his expected payoff from proposing x is

$$U_P(x|\sigma_R) = \sigma_R U_P(x|l = 1) + (1 - \sigma_R) U_P(x|l = 0).$$

That is, it is a strict convex combination of $U_P(x|l = 1)$ and $U_P(x|l = 0)$, weighted by σ_R. Both functions are single peaked with maxima at x_1^* and x_0^* respectively; the former is continuous while the latter exhibits a discontinuous jump at x_0^*. However, the *order* of their maxima, and thus the exact shape of their convex combination $U_P(x|\sigma_R)$, will depend on the relationship between uncertainty and war costs, i.e., $\delta > (<)(=)\, 2(\bar{c}_P + \bar{c}_R)$.

Figure 1 illustrates. $U_P(x|l = 1)$ and $U_P(x|l = 0)$ are single-peaked over the range $[\underline{w}_R, \overline{w}_R]$. As such, any maximizer of their convex combination must fall within the closed

Figure 1: State P's expected payoffs.

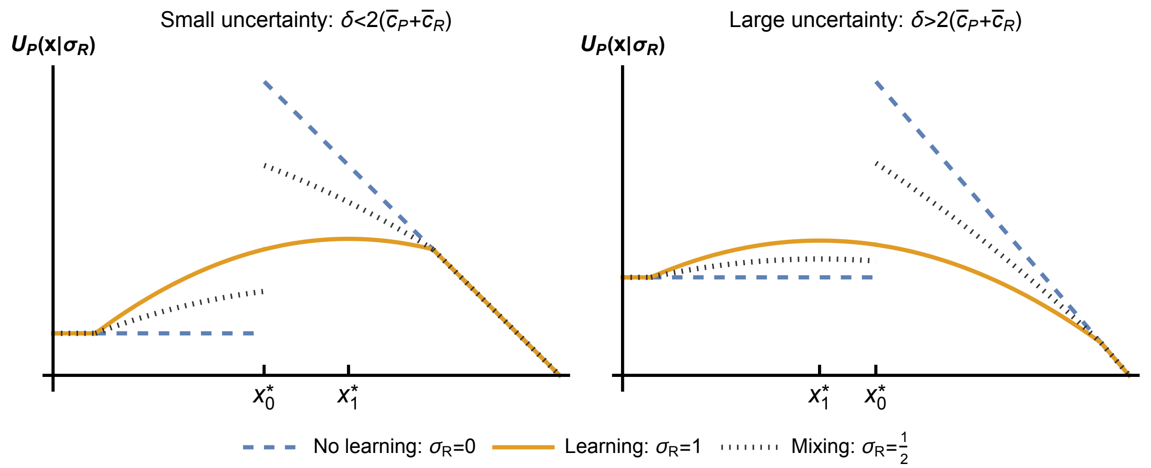

Notes: Figure generated assuming $\bar{p} = 0.5$ and $\lambda = 1$. In the left panel, $\bar{c}_P = \bar{c}_R = \frac{3}{16}$ and $\delta = \frac{1}{2}$. In the right panel, $\bar{c}_P = \bar{c}_R = \frac{1}{8}$ and $\delta = \frac{2}{3}$.

interval between their peaks x_0^* and x_1^*. Moreover, whether $U_P(x|\sigma_R)$ is *also* single-peaked over this interval (and therefore has a unique maximum) will depend ρ, because $U_P(x|l = 0)$ is discontinuous at its peak x_0^* and constant below the peak, which introduces a convexity into $U_P(x|\sigma_R)$. Specifically, when uncertainty is low relative to costs (i.e., $\delta < 2(\bar{c}_P + \bar{c}_R)$ so that $x_0^* < x_1^*$), then $U_P(x|\sigma_R)$ is strictly increasing over $x \leq x_0^*$ and a convex combination of a strictly concave and linearly decreasing function over $x > x_1^*$. As such, it must have a unique maximum at some $x_m^* \in [x_0^*, x_1^*)$, as in the left panel of Figure 1. Alternatively, when uncertainty is more substantial (i.e., $\delta > 2(\bar{c}_P + \bar{c}_R)$ so that $x_0^* < x_1^*$), then $U_P(x|\sigma_R)$ has two local maxima at x_1^* and x_0^*, as in the right panel of Figure 1. Which local maximum is the global one will depend on state R's probability of learning σ_R, and in a mixed equilibrium both must be maxima so that state P is willing to randomize over x_0^* and x_1^*. The following summarizes the form of offers in any mixed strategy equilibrium.

Lemma 2. *If $x \in (\phi(x_1^*), \phi(x_0^*))$ so that state R must randomize over learning in equilibrium, then the following hold.*

1. *If $\delta < 2(\bar{c}_P + \bar{c}_R)$ then state P surely proposes some $x_m^* \in (x_0^*, x_1^*)$.*

2. *If $\delta > 2(\bar{c}_P + \bar{c}_R)$ then state P randomizes between proposing x_0^* and proposing x_1^*.*

Having established the form of offers in any mixed strategy equilibrium, we next precisely characterize mixed equilibria. We begin with the case where ratio between uncertainty and war costs is small, i.e., $\delta < 2(\bar{c}_P + \bar{c}_R)$, so the informed offer is more generous than the uninformed offer ($x_1^* > x_0^*$).

Proposition 2. *A mixed strategy equilibrium in which state P surely offers $x_m^* \in (x_0^*, x_1^*)$ and state R randomizes over learning his strength exists if and only if $k_R \in (\phi(x_1^*), \phi(x_0^*))$ and $\delta < 2(\bar{c}_P + \bar{c}_R)$. In this equilibrium,*

$$\phi(x_m^*) = k_R \qquad \textit{and} \qquad \sigma_R^* = \frac{\delta}{\delta + (2-\lambda)(x_1^* - x_m^*)}.$$

In equilibrium, state P makes a proposal $x_m^* \in (x_0^*, x_1^*)$ that is strictly in between the optimal uninformed and informed proposals, so that state R is exactly indifferent between learning and not given her moderate learning cost $k_R \in (\phi(x_1^*), \phi(x_0^*))$. State R's probability of learning must in turn incentivize state P to make precisely this offer, implying that it is directly affected by uncertainty δ, strength interdependence λ, and the distance $x_1^* - x_m^*$ between the learning offer and the equilibrium offer. Also worth noting is that state R's equilibrium learning probability depends *indirectly* on her own cost of learning k_R via its equilibrium relationship with x_m^*. Specifically, as the cost of learning k_R increases, state P must be expected to make a stingier equilibrium offer x_m^* that more strongly incentivizes learning, so that state R is still willing to learn. However, in order to incentivize P to make a stingier offer, state R must in turn learn his own strength less often.

We next consider the case where ratio between uncertainty and war costs is large, i.e., $\delta < 2(\bar{c}_P + \bar{c}_R)$, so the informed offer is stingier than the uninformed offer ($x_1^* < x_0^*$).

Proposition 3. *A mixed strategy equilibrium in which state P randomizes over proposing x_0^* and proposing x_1^*, and state R randomizes over learning his strength, exists if and only if $k_R \in (\phi(x_1^*), \phi(x_0^*))$ and $\delta > 2(\bar{c}_P + \bar{c}_R)$. In this equilibrium,*

$$\sigma_P^*(x_0^*) = \frac{\phi(x_1^*) - k_R}{\phi(x_1^*) - \phi(x_0^*)} \qquad \textit{and} \qquad \sigma_R^* = \frac{2\delta(\bar{c}_P + \bar{c}_R)}{(2-\lambda)(x_0^* - x_1^*)^2 + 2\delta(\bar{c}_P + \bar{c}_R)}.$$

As in the previous case, state P's anticipated offers must make state R willing to randomize over learning his strength. However, as described in Lemma 2, it is no longer possible to incentivize state P to make an *interior* offer with this property. Instead, state P must be incentivized to *randomize* over the more-generous no learning offer x_0^* and stingier learning offer $x_1^* < x_0^*$ in a manner that accomplishes the same end. The learning probability σ_R^* that does so again depends directly uncertainty δ, interdependence λ, and offer distance $x_0^* - x_1^*$, but also the aggregate cost of war $\bar{c}_P + \bar{c}_R$. Finally, P's offer probabilities must incentivize state R to randomize over learning; thus, as learning becomes costlier P must become more likely to make the no learning offer x_0^* that better incentivizes learning relative to the learning offer $x_1^* < x_0^*$. We note that in contrast to the previous case, state R's learning probability no longer depends on her own cost of learning k_R, since the learning probability required required to incentivize P to randomize does not depends on k_R.

Figure 2 illustrates the form of equilibrium for different values of state R's learning cost

Figure 2: Summary of equilibrium characterization.

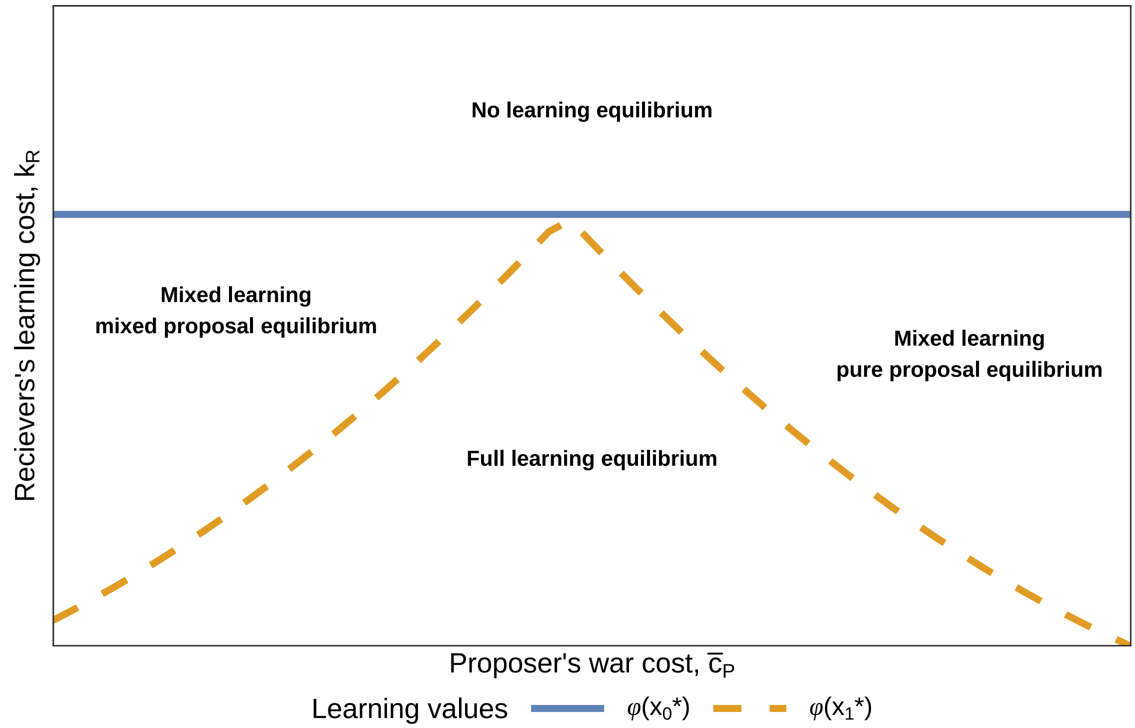

Notes: Figure generated assuming $\bar{p} = 0.45$, $\lambda = 0.5$, $\delta = 0.6$, and $\bar{c}_R = 1/6$. $\phi(x_1^*) = \phi(x_0^*)$ if and only if $\bar{c}_P = \frac{\delta}{2} - \bar{c}_R$.

k_R on the vertical axis and state P's cost of war on the horizontal axis. First note that the threshold $\phi(x_0^*)$ separating no-learning from mixed learning equilibria is constant as a function of the proposer's war cost; the reason is this cost does not affect state P's optimal no learning offer x_0^*, as she has all the bargaining power. In contrast, the threshold $\phi(x_1^*)$ separating mixed learning from full learning equilibria is first increasing and then decreasing in $\bar{c}_P$. As P's cost of war increases, she makes an increasingly generous offer (i.e., a higher x_1^*) to an informed R. Recall that if $\bar{c}_P < \frac{\delta}{2} - \bar{c}_R$ then uncertainty is large relative to the aggregate cost of war, so that $x_1^* < x_0^*$. From Lemma 1, as $\bar{c}_P$ increases x_1^* approaches x_0^* from below, and consequently the value of learning increases. Conversely, if $\bar{c}_P > \frac{\delta}{2} - \bar{c}_R$ then uncertainty is small relative to the aggregate cost of war ($\bar{c}_P > \frac{\delta}{2} - \bar{c}_R$), so that $x_1^* < x_0^*$. Thus, the value of learning decreases in $\bar{c}_P$ as x_1^* becomes increasingly distant from x_0^*.

To conclude this section, we note that outside of a measure-zero subset of learning costs there is a unique equilibrium. In our subsequent comparative statics analysis, we select the unique pure strategy equilibrium when multiple equilibria exist.

Remark 3. *If $k_R \notin \{\phi(x_0^*), \phi(x_1^*)\}$, then there is a unique equilibrium characterized by Propositions 1–3.*

4 Comparative Statics

4.1 Learning and war

What is the relationship between learning costs and war? It is well known that bargaining failure can result from asymmetric information (Fearon 1995; Powell 2002; Ramsay 2017). Extrapolating to our setting, endogenous learning creates asymmetric information, so it is reasonable to suppose that smaller learning costs will increase the likelihood of war. Our results reflect this to some degree; when $k_R < \phi(x_1^*)$ learning surely occurs, asymmetric uncertainty exists, and war occurs with positive probability. When $k_R > \phi(x_0^*)$ learning never occurs, so symmetric uncertainty exists, and war never occurs. In between these extremes, however, the relationship between learning costs and war is more complex. Moreover, the nature of this complexity differs substantially between the low uncertainty and high uncertainty cases.

Proposition 4. *The probability of learning, and hence of asymmetric information, is decreasing in the cost of learning. If $\delta \neq 2(\bar{c}_P + \bar{c}_R)$, then the probability of war is non-monotonic in the cost of learning. Specifically, the following hold.*

1. *The probability of war is positive and constant when $k_R < \phi(x_1^*)$, and it is zero and constant when $k_R > \phi(x_0^*)$.*
2. *If uncertainty is small ($\delta < 2(\bar{c}_P + \bar{c}_R)$) and $k_R \in [\phi(x_1^*), \phi(x_0^*))$, then it is continuous and strictly increasing.*
3. *If uncertainty is large ($\delta > 2(\bar{c}_P + \bar{c}_R)$), it is continuous and strictly decreasing when $k_R \in (\phi(x_1^*), \phi(x_0^*))$, but with a discontinuous increase at $k_R = \phi(x_1^*)$.*

Figure 3 illustrates Proposition 4. In the left panel, uncertainty is small. Somewhat surprisingly, in this case the probability of war is *increasing* in the cost of learning k_R when it is intermediate ($k_R \in [\phi(x_1^*), \phi(x_0^*))$). To see why, first recall that small uncertainty means that P's offer to an informed R is more generous than to an uninformed one ($x_1^* > x_0^*$), which further implies that her equilibrium offer $x_m^* > x_0^*$ averts war if R does not learn (see Proposition 2). Thus, in this case war only occurs if *both* R learns his own strength *and* it exceeds the offer ($w_R > x_m^*$), so that the probability of war is

$$\sigma_R^* \Pr(w_R > x_m^*), \tag{8}$$

As described in Proposition 2, the probability R learns is *decreasing* in his cost of learning k_R, which reduces wars ceteris paribus. However, P's equilibrium offer x_m^* also must become stingier as the learning cost k_R increases to better incentivize learning, which increases wars

Figure 3: Learning costs and the probability of war.

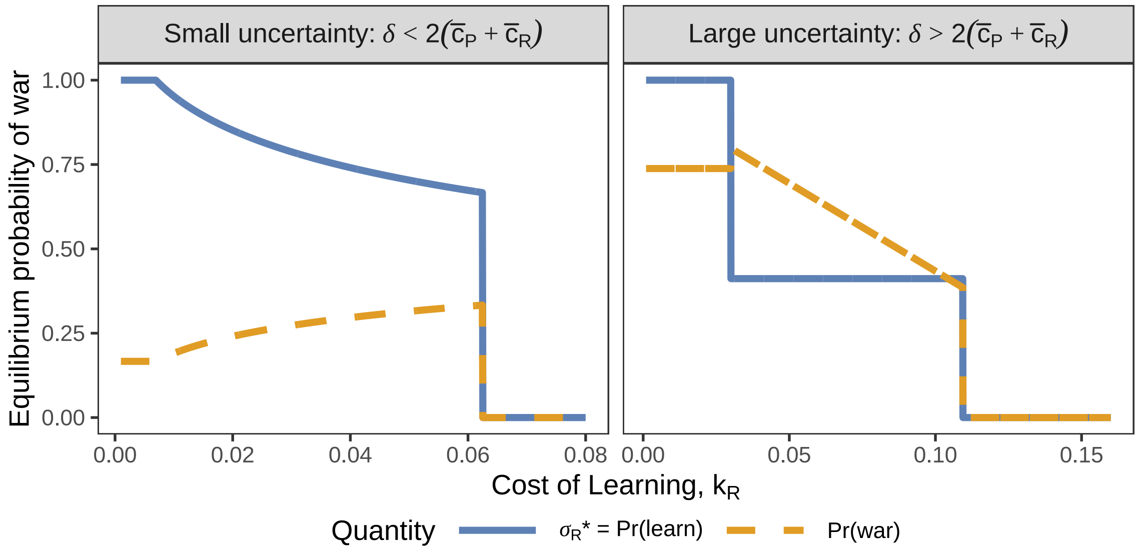

Notes: Figure generated assuming $\bar{p} = 0.5$ and $\lambda = 0.5$. In the left panel, $\bar{c}_i = 4/16$ and $\delta = 0.5$. In the right panel, $\bar{c}_i = 1/16$ and $\delta = 7/8$.

ceteris paribus. Thus, there are countervailing effects. In the proof of Proposition 4, we show that the effect of stingier offers always dominates the effect of reduced learning.

In the right panel, uncertainty is large. When the learning cost is intermediate its effect is straightforward—higher learning costs decrease war—but for somewhat counterintuitive reasons. In this case, P randomizes over the offer x_0^* that averts war if R is uninformed and a stingier offer x_1^* that will trigger war if R is uninformed. As the cost of learning k_R increases, P must become more likely to make the more generous offer x_0^* to in order to incentivize R to sometimes learn rather than reject outright; however, the probability R learns is unchanged. Thus, higher learning costs decrease wars via more generous offers.

The reason for the discontinuous increase at $k_R = \phi(x_1^*)$ is more nuanced. To understand this effect it is helpful to instead think about how strategies change as learning costs *decrease* continuously toward $\phi(x_1^*)$. As this occurs, P becomes increasingly *and continuously* more likely to make a "stingy" offer $x_1^* < x_0^*$ targetted at an informed R. While P's offer behavior cannot account for the discontinuity, the probability R learns also *jumps discontinuously* from σ_R^* to 1. Overall, the discontinuous change in the probability of war at the point of transition to the full learning equilibrium is $(1 - \sigma_R^*) \cdot \Pr(w_R \leq x_1^*)$, i.e., the increase in the probability R learns, times the probability an informed R will accept the stingier offer x_1^*.

Somewhat surprisingly, this expression is *strictly positive*, meaning that more learning by R actually *reduces* wars, contra the conventional wisdom. Why is this? The key insight is that in our model, learning and asymmetric information do not always increase the chance of war *ceteris paribus*. Specifically, if P makes a stingy offer $x_1^* < x_0^*$ that R is inclined to *reject* absent learning, then learning actually *decreases* the chance of war, by creating the possibility that R learns his war actual payoff w_R is below even the stingy offer ($w_R \leq x_1^*$). This situation is exactly what occurs when uncertainty is large and learning costs are intermediate, so that P sometimes "accidentally" makes the optimal offer x_1^* for an informed R to an uninformed R, who rejects it outright. Lastly, we note that this possibility requires that state P be *uncertain* about whether state R is better informed, which is a key novelty of our model; otherwise, she would never rationally make an offer that would trigger war to an uninformed receiver.

4.2 Interdependence, learning, and war

How do different types of uncertainty shape the receiver's incentive to learn? Fey and Ramsay (2011) emphasize the importance of common versus private value uncertainty in shaping the relationship between asymmetric information and war; peace is more difficult to sustain when there is uncertainty about common value information like the balance of power. In our model, interdependence λ parameterizes the extent to which the initially-unknown information about R's strength is about common versus private values. We next examine how interdependence affects state R's incentive to learn.

We first recall how interdependence λ affects the learning offer x_1^*, as characterized in Equation 5 (the no learning offer x_0^* is unaffected by interdependence).

Remark 4. *The gap $|x_1^* - x_0^*|$ between the learning and no-learning offer increases as interdependence grows, i.e., λ increases.*

Recall that the learning offer x_1^* balances the marginal benefit of a more generous offer (which will persuade the marginal type of receiver to accept rather than reject and fight) against the marginal cost (paying more to the receiver when he is a lower type who would have accepted regardless). When δ is small the learning offer is more generous than the no learning offer ($x_1^* > x_0^*$), so the marginal type of receiver ($w_R = x_1^*$) is *stronger* than the average type ($w_R = E[w_R] = x_0^*$). As such, with greater interdependence λ, state P's incentive to avoid war against the marginal type grows, pushing the optimal offer up and away from x_0^*. Conversely, when δ is large the learning offer is less generous than the no learning offer ($x_1^* < x_0^*$), so the marginal type of receiver is *weaker* than the average type. With greater interdependence, state P's incentive to avoid war against the marginal type therefore weakens, pushing the optimal offer down and again away from x_0^*.

The main implication of the preceding that greater interdependence discourages learning, and thus reduces the likelihood of asymmetric information. The reason is that greater interdependence makes state P's proposals more *sensitive* to her beliefs about the likelihood that state R is better informed, which in turn reduces state R's incentive to become informed for any non-zero belief that state P may hold (and optimally respond to). An immediate implication is that, because learning can either increase or decrease peace (Proposition 4), interdependence has countervailing effects on the likelihood of war.

Proposition 5. *As interdependence* λ *increases, learning and asymmetric information decrease. Specifically,*

- *The value of learning* $\phi(x_1^*)$ *when expecting the learning offer* x_1^* *is weakly decreasing in* λ*, and strictly decreasing if* $\delta > \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$*. The value of learning* $\phi(x_0^*)$ *when expecting the no learning offer* x_0^* *is unaffected by* λ*.*
- *The range of learning costs* k_R *for which equilibrium exhibits full vs. mixed learning is strictly decreasing in* λ*.*
- *Within the mixed learning region, the probability of learning is strictly decreasing in* λ*.*

Figure 4 illustrates; in each panel the horizontal axis is state R's cost of learning k_R and the vertical axis is the equilibrium probability σ_R^* state R learns. It is straightforward to see that the full learning region (where $\sigma_R^* = 1$) is strictly smaller for higher λ, and a learning cost k_R with partial learning exhibits even less learning for higher λ. To understand the latter, observe that state P's *actual* equilibrium offer must remain unchanged in the mixed learning region when interdependence grows, since P's offers must incentivize state R to randomize over learning. However, state P's *incentive* to make an offer diverging from x_0^* grows. To disincentivize such offers, state R's must therefore learn less often in equilibrium.

Figure 4: Equilibrium Learning Probability.

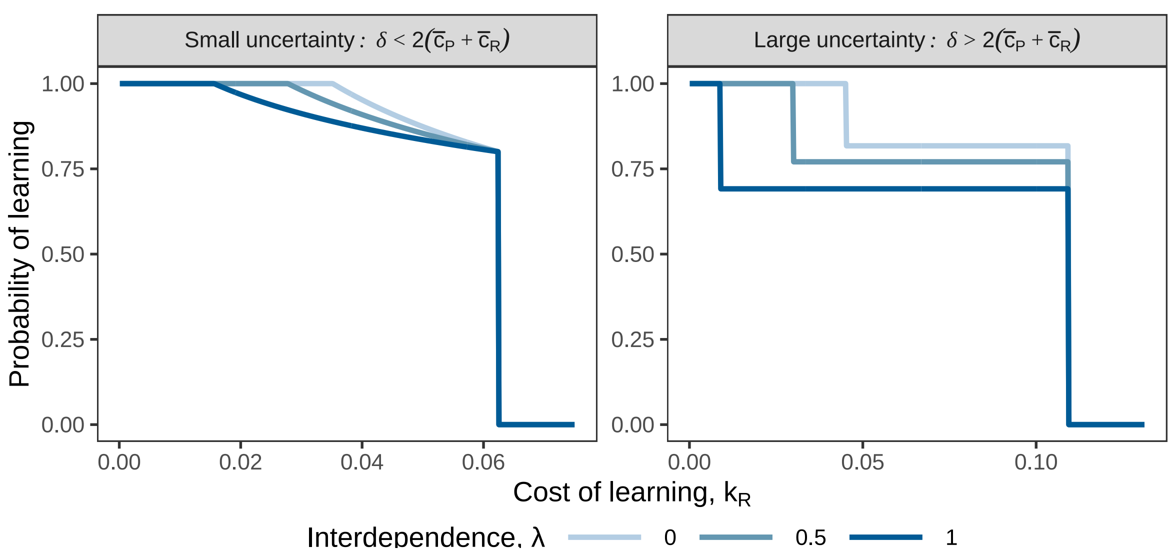

Notes: The figure are generated assuming that $\bar{p} = 0.5$. In the left panel, $\bar{c}_P = \bar{c}_R = \frac{3}{16}$ and $\delta = \frac{1}{2}$. In the right panel, $\bar{c}_P = \bar{c}_R = \frac{1}{16}$ and $\delta = \frac{7}{8}$.

4.3 Ex ante uncertainty and war

A common hypothesis in the empirical conflict literature is the principle of convergence: reductions in uncertainty about resolve or relative capabilities should reduce the onset of war (Kaplow and Gartzke 2015; Reiter 2003). For example, several scholars attempt to explain variation in the onset and duration of conflict by using leader tenure to proxy for uncertainty, where the key hypotheses monotonically relate uncertainty to conflict (Smith and Spaniel 2019; Spaniel and Smith 2015; Thyne 2012; Uzonyi and Wells 2016). As might be expected from the previous analysis, however, the relationship between uncertainty and war is more complex when the bargaining parties may *also* be uncertain about what each side knows about their own strength, as in our model. As before, it is helpful to decompose the discussion between the case with low uncertainty $\delta < 2(\bar{c}_P + \bar{c}_R)$ and high uncertainty, $\delta > 2(\bar{c}_P + \bar{c}_R)$. Moreover, we also consider two separate cases of pure private values $\lambda = 0$ and pure common values $\lambda = 1$.

Common to all these cases will be how uncertainty affects state R's value of learning given offer x_0^*, which does not depend on uncertainty δ or interdependence λ. The next remark illustrates a standard result in models of information acquisition.

Remark 5. *The threshold $\phi(x_0^*) = \frac{\delta}{8}$ is strictly increasing in uncertainty, δ.*

Recall that the offer x_0^* makes an uninformed state R indifferent between war and accepting ex-ante. As such, the value of information only depends on the amount of uncertainty δ via its *direct effect* on the receiver's incentive to learn.

Comparative statics on $\phi(x_1^*)$ with respected to uncertainty δ are more complicated, however, because uncertainty has both a direct effect on the value of learning, and an *indirect effect* via its influence on the learning offer x_1^*. Moreover, the strength of the indirect affect is mediated by interdependence λ, which determines how sensitive x_1^* is to greater uncertainty δ. This complexity has a crucial effect on equilibrium outcomes.

Low uncertainty ($\delta < 2(\bar{c}_P + \bar{c}_R)$) When ex-ante uncertainty about R's strength is low, the proposer will always makes a generous offer that an uninformed receiver will accept, so war will only occur if R learns that its war payoff is larger than the offer. Consequently, asymmetric information is a necessary condition for war. Expanding the expression in Equation 8 for the probability of war and explicitly denoting the dependence of equilibrium strategies on δ yields that:

$$\Pr(War) = \underbrace{\sigma_R^*(\delta)}_{\Pr(l=1)} \cdot \underbrace{\left(\frac{1}{2} - \frac{x_m^*(\delta) - x_0^*}{\delta}\right)}_{\Pr(w_R > x^*(\delta))}.$$

In the expression, greater uncertainty δ and lower offers $x_m^*(\delta)$ lead to more war *ceteris*

Figure 5: Uncertainty and the probability of war with private values ($\lambda = 0$).

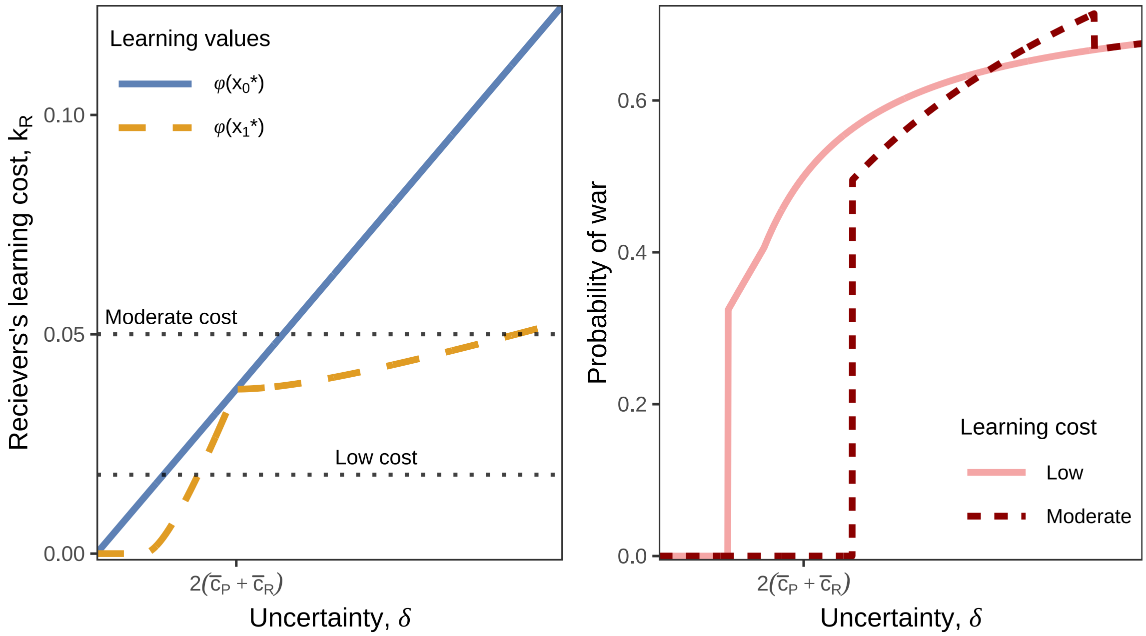

Notes: The left panel graphs the values of learning as a function of uncertainty, δ, where we assume $\bar{p} = 0.6$, $\bar{c}_R = \frac{1}{10}$, $\bar{c}_P = \frac{1}{20}$ and $\lambda = 0$. Fixing learning costs at the low value ($k_R = 0.02$) and the moderate value ($k_R = 0.05$), the right panel graphs the resulting probability of war as a function of δ.

paribus via a higher likelihood that R's actual war payoff w_R exceeds P's offer $x^*_m(\delta)$, while more learning also leads to more war *ceteris paribus* via a higher likelihood that R actually learns when her war payoff exceeds P's offer. Characterizing the equilibrium relationship between uncertainty and war, however, also requires accounting for how uncertainty affects the players' equilibrium strategies $\sigma^*_R(\delta)$ and $x^*_m(\delta)$. Despite this potential complexity, the overall equilibrium relationship is relatively simple.

Proposition 6. *Suppose uncertainty is low, i.e.,* $\delta < 2(\bar{c}_R + \bar{c}_R)$.

- *If* $\delta < \frac{2(\bar{c}_P + \bar{c}_R)}{3-\lambda}$ *then the threshold* $\phi(x^*_1)$ *that separates the full and mixed learning regions is* 0*. If* $\delta > \frac{2(\bar{c}_P + \bar{c}_R)}{3-\lambda}$ *then it is strictly positive and increasing in* δ*.*
- *If learning is not too costly* ($k_R < \frac{\bar{c}_P + \bar{c}_R}{4}$)*, then the probability of war is constant for* $\delta \in (0, 8k_R)$ *and strictly increasing for* $\delta > 8k_R$*. If learning is costly* ($k_R > \frac{\bar{c}_P + \bar{c}_R}{4}$)*, then the probability of war is zero and constant.*

Despite this relative simplicity, the relationship between uncertainty and war is sometimes flat and exhibits discontinuities. To see why, it is helpful to examine the left panel of Figure 5, which resembles Figure 2 but places ex-ante uncertainty δ on the horizontal axis, and demarcates the low ($\delta < 2(\bar{c}_P + \bar{c}_R)$) and high ($\delta > 2(\bar{c}_P + \bar{c}_R)$) uncertainty subregions.

As in Figure 2, the vertical axis is the cost of learning k_R, and the relationship between k_R and the receiver's values of learning $\phi(x_0^*)$ and $\phi(x_1^*)$ determines the form of the equilibrium.

First consider the subregion where $\delta < 2(\overline{c}_R + \overline{c}_R)$; here both learning thresholds $\phi(x_0^*)$ and $\phi(x_1^*)$ are increasing in δ, because the *indirect* effect of greater uncertainty is to decrease the learning offer x_1^* in the direction of the no learning offer x_0^*. This implies that increasing δ over the range $\delta < 2(\bar{c}_P + \bar{c}_R)$ either (1) has no effect on the form of equilibrium since it always exhibits no learning or war (if the cost of learning is sufficiently high), or (2) the equilibrium transitions from no learning, to partial learning, and finally to pure learning.

Next consider the subregion where $\delta > 2(\overline{c}_R + \overline{c}_R)$. Within the equilibrium exhibits no learning the relationship between uncertainty and war is necessarily flat, since wars never occur. When uncertainty increases sufficiently to transition to a partial learning equilibrium, the probability of war jumps abruptly because R may learn when her true war payoff w_R exceeds P's offer $x_m^*(\delta)$. Further increases in uncertainty then cause the probability of war to increase via a higher equilibrium likelihood $\sigma_R^*(\delta)$ that R learns, which outweighs the countervailing effect of P's increasingly generous offers $x_m^*(\delta)$ to disincentivize learning. Finally, uncertainty becomes sufficiently high to transition to the pure learning equilibrium; here further increases in uncertainty increase the chance of war by both increasing the chance that R is very strong and reducing P's equilibrium offer $x_1^*(\delta)$.

High uncertainty with private values ($\delta > 2(\bar{c}_P + \bar{c}_R)$ and $\lambda = 0$) When ex-ante uncertainty about R's strength is high, the proposer will propose either x_0^* or x_1^* (and may randomize between them). Notably, however, the latter offer is a "stingy" one that R will reject if he chooses to remain uninformed. High uncertainty thus creates an additional path to war overlooked in the previous literature; that P, entertaining the mere *possibility* that R is better-informed, will make a stingy offer $x_1^*(\delta) < x_0^*$ intended for an informed R that an uninformed R will reject. Consequently, the overall probability of war is

$$\Pr(War) = \sigma_R^*(\delta) \cdot \left(\frac{1}{2} + (1 - \sigma_P^*(\delta)) \cdot \frac{(x_0^* - x_1^*(\delta))}{\delta} \right) + (1 - \sigma_R^*(\delta)) \cdot (1 - \sigma_P^*(\delta)),$$

where $\sigma_P^*(\delta)$ denotes the probability that P proposes x_0^*. This additional path to war implies that asymmetric information may not always contribute to war. Rather, *if* P chooses to make the stingy offer $x_1^*(\delta)$ that is optimal for an informed receiver, then R learning his strength will actually *reduce* the chance, of war by allowing him to learn when his actual war payoff w_R is below even this stingy offer.

To present the next result, we introduce two additional pieces of notation. First, when $\lambda = 0$, Assumption 2 implies that $\delta < \bar{\delta}_0 = 2\min\{x_0^*, 1 - x_0^*\}$. Viewed as a function of δ, we may calculate an upper bound on state R's value of learning when he anticipates the learning offer x_1^*, which is $\bar{\phi}_1(\bar{\delta}_0) = \phi(x_1^*)|_{\delta=\bar{\delta}_0}$. In words, $\bar{\phi}_1(\bar{\delta}_0)$ denotes the threshold between the full and mixed learning regions when δ takes its maximum value and $\lambda = 0$.

Proposition 7. *Suppose ex ante uncertainty is high ($\delta > 2(\overline{c}_R + \overline{c}_R)$) and R's potential private information is private value ($\lambda = 0$).*

- *The threshold demarcating the always-learn and mixed-learning equilibrium, $\phi(x_1^*)$, is strictly increasing in uncertainty, δ.*
- *If the cost of learning is moderate, i.e., $\frac{\overline{c}_R+\overline{c}_R}{4} < k_R < \bar{\phi}_1(\bar{\delta}_0)$, then there exists $\delta < \delta'$ such that the probability of war is greater for δ than for δ'*

With high ex-ante uncertainty, increases in uncertainty can therefore reduce the chance of war through the mechanism previously described: it elicits more private learning from R, allowing him to sometimes avoid war despite a stingy offer $x_1^*(\delta)$. This is illustrated in Figure 5's right panel with moderate costs; as $\delta > 2(\bar{c}_P + \bar{c}_R)$ increases the equilibrium transitions from mixed learning with random proposals to full learning, decreasing the chance of war.

To better understand this novel effect, consider Figure 5's left panel. As previewed by Proposition 7, the key property to observe is that the value of learning when anticipating either equilibrium offer is *increasing* in ex-ante uncertainty δ. Thus, learning costs must be neither too small nor too large for equilibrium to transition from no learning, to partial learning, and then to pure learning. The first transition necessarily increases the chance of war, because absent learning the proposer P offers x_0^* and avoids war. The second transition, however, must decrease the chance of war, because at the point of transition P is certain to make the stingy offer $x_1^*(\delta)$ that triggers war absent learning, while R's probability of learning increases discontinuously to 1.

Finally, when the equilibrium exhibits partial learning, greater ex-ante uncertain increases war through two channels. First, as in the $\delta < 2(\bar{c}_P + \bar{c}_R)$ case, the proposer makes stingier offers, which increases the probability of war conditional on learning. Second, in contrast to the $\delta < 2(\bar{c}_P + \bar{c}_R)$ case, state R becomes less like likely to learn; and when state R does not learn, he surely rejects the stingier learning offer $x_1^* < x_0^*$.[6]

High uncertainty with common values ($\delta > 2(\bar{c}_P + \bar{c}_R)$ and $\lambda = 1$) Somewhat surprisingly, when uncertainty is over common values it is no longer the case that greater uncertainty can reduce war, and we recover the standard relationship.

Proposition 8. *Suppose ex-ante uncertainty is high ($\delta > 2(\overline{c}_R + \overline{c}_R)$) and R's potential private information is common value ($\lambda = 1$).*

- *The threshold demarcating the always-learn and mixed learning equilibrium, $\phi(x_1^*)$, is strictly decreasing in uncertainty, δ.*

[6] Although the effect on state P's probability of making the generous offer x_0^* is ambiguous, we show in the Appendix that these two effects dominate the effect of any potential increase in the likelihood P makes the generous offer (see Claim A.10).

Figure 6: Uncertainty and the probability of war with common values ($\lambda = 1$).

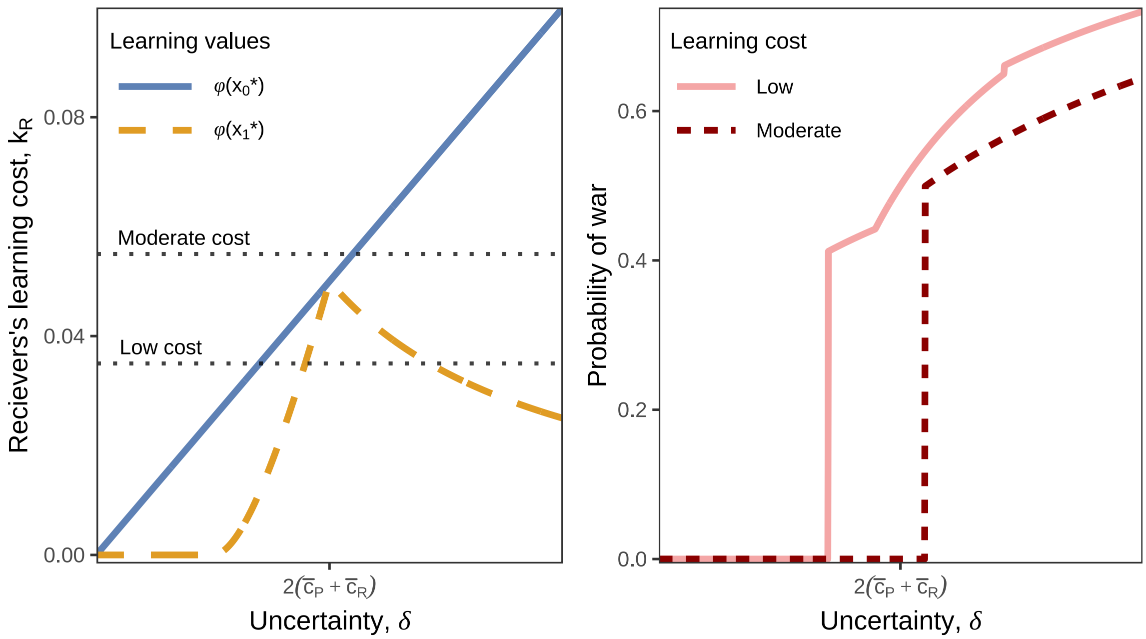

Notes: Figure generated assuming $\bar{p} = 0.6$, $\bar{c}_i = \frac{1}{10}$, $\lambda = 1$, and the learning costs are 0.035 and 0.055 for low and moderate costs, respectively.

- *The probability of war is weakly increasing in δ.*

The right panel of Figure 6 illustrates the relationship between uncertainty and war when uncertainty is high and $\lambda = 1$. As in the low uncertainty case it is uniformly increasing, but with flat regions and discontinuities. This occurs because (in contrast to the private values case) the value of learning $\phi(x_1^*)$ when anticipating the learning offer is *decreasing* in uncertainty δ, as illustrated in the left panel of Figure 6. Thus, as δ increases, the equilibrium must either transition from no learning to partial learning (if k_R is high), or from full learning to partial learning (if k_R is low). This precludes the possibility that greater uncertainty might increase wars, because either (a) higher δ increases learning but learning contributes to war, or (b) higher δ decreases learning but learning reduces wars.

Why is the receiver's value of learning $\phi(x_1^*(\delta))$ under high uncertainty increasing when interdependence is high ($\lambda = 1$)? Recall that interdependence affects the *sensitivity* of proposer's offers to her beliefs about the likelihood that R is informed. Thus, the *reduction* $x_0^* - x_1^*(\delta)$ in P's offer when she believes R to be informed is greater in the common values case than in the private values case, further implying that greater uncertainty δ reduces P's offer more in the common values case than in the private values case, thereby indirectly reducing the value of learning. This outweighs the direct effect of learning becoming more valuable in the face of greater uncertainty δ (holding the learning offer $x_1^*(\delta)$ fixed).

Overall, the model identifies necessary and sufficient conditions for receiver learning to overturn the literature's broadly held view of a positive relationship between uncertainty and war. First, Propositions 6 and 8 demonstrate that high ex ante uncertainty and some degree of private values are *necessary* conditions for a non-monotonic relationship between uncertainty and war. Put differently, if uncertainty is small *or* uncertainty is over common values, then greater uncertainty must increase war.[7] Second, Proposition 7 provides an additional necessary condition: that the cost of learning be moderate. When private values are combined with high uncertainty and moderate learning costs, a non-monotonic relationship between uncertainty and war emerges because greater uncertainty can incentivize the receiver to learn when to accept an apparently stingy offer.

5 Conclusion

This paper develops a crisis-bargaining model in which states can make costly private investments to learn about their own military strength before bargaining. Anticipating the possibility of learning, proposers adjust their offers in ways that feed back into receiver's incentives to acquire information, generating equilibria with uncertainty over whether learning has occurred. As a result, the relationship between learning, uncertainty, and war can be non-monotonic: asymmetric information can sometimes reduce conflict, moderate learning costs generate the highest risk of conflict, and greater ex ante uncertainty can sometimes promote peace rather than war. These findings complicate standard interpretations of uncertainty in crisis bargaining, and illustrate that the underlying cause of asymmetric information matters for its relationship to bargaining breakdown and war.

The model suggests a variety of fruitful avenues for future work. A key property of the model is that acquiring better information about the strategic situation does not always benefit the receiver, for two reasons. First, the proposer effectively "gives up" trying to buy peace with greater concessions when she believes a better informed receiver to be too unpredictable. Second, when the proposer is uncertain about whether the receiver is better informed, she may "accidentally" make an unacceptably stingy offer to an uninformed receiver who could have been predictably "bought off." Both paths to war illustrate the importance of considering additional ways for the bargaining parties to alleviate asymmetric information, including cheap talk communication, verifiable communication, and costly signalling. Some possibilities for communication are implausible; e.g., a receiver credibly communicating their ignorance when the fundamentals are such that an uninformed receiver will receive a more generous offer. Others seem more reasonable – e.g., a receiver credibly communicating *that* they have invested in becoming informed by publicizing intelligence

[7] For intermediate interdependence $\lambda \in (0, 1)$ that is neither fully common nor private value, the value of learning $\phi(x_1^*(\delta))$ first decreases and then increases in δ. Consequently, the form of equilibrium may transition from full learning to partial learning and then back again, thereby exhibiting both a discontinuous jump followed by a discontinuous drop for some values of k_R—see Appendix Figure B.1 for an example.

efforts or military exercises, while censoring exactly what they have learned.

Another interesting avenue of research is to consider a sequential game form where the proposer makes her offer *before* the receiver choose to acquire their information, which more closely resembles previous analyses of bargaining with endogenous learning in economics. In this alternative game form, the proposer can accurately forecast how her offer will affect the receiver's decision to learn, and will therefore sometimes try to "buy" the receiver's ignorance by making her offer enticing enough to accept outright rather than carefully consider. We view our model's simultaneous sequence as a more accurate proxy for a fluid interstate conflict situation, in which the exact actions the receiver can take to become informed, and the timing with which he take them (relative to the timing of bargaining), is nebulous. However, the sequential game form may be a better fit to some specific bargaining situations in which bargaining must occur at a specific place and time (as in a scheduled "summit"), and the parties cannot easily identify what sort of information would help them become better informed until offers are made.

Finally, extending our analysis to richer strategic environments is an important direction for future research. Two are particularly interesting. First, it is reasonable to consider the case of a proposer who can also learn about his strength. When this is the case, the proposer's offer also potentially signals whether and what she has learned, which significantly complicates the strategic interaction. Second, the information that states privately learn might accidentally or intentionally leak to their rivals. It would be interesting to examine the effects of such "information leakage," occuring either exogenously or due to the active efforts of one or both states in the interaction. Indeed, our analysis demonstrates that the receiver will sometimes benefit from such leaks, but it is unclear what the proposer's preferences are over leaks, and how leaks affect bargaining breakdown.

References

Chatterjee, Kalyan, Miaomiao Dong and Tetsuya Hoshino. 2025. "Bargaining and information acquisition." *American Economic Journal: Microeconomics* 17(3):75–102.

Crémer, Jacques and Fahad Khalil. 1992. "Gathering information before signing a contract." *American Economic Review* 82(3):566–578.

Crémer, Jacques, Fahad Khalil and Jean-Charles Rochet. 1998. "Strategic Information Gathering Before a Contract is Offered." *Journal of Economic Theory* 81(1):163–200.

Debs, Alexandre and Nuno P. Monteiro. 2014. "Known Unknowns: Power Shifts, Uncertainty, and War." *International Organization* 68(1):1–31.

Fearon, James. 1995. "Rationalist explanations for war." *International Organization* 49(3):379–414.

Fey, Mark and Kristopher Ramsay. 2011. "Uncertainty and incentives in crisis bargaining: Game-free analysis of international conflict." *American Journal of Political Science* 55(1):149–169.

Kaplow, Jeffrey and Erik Gartzke. 2015. "Knowing Unknowns: The Effect of Uncertainty in Interstate Conflict." Unpublished manuscript: `http://dl.jkaplow.net/uncertainty.pdf`.

Kertzer, Joshua D., Joshua Busby, Jonathan Monten, Jordan Tama and Craig Kafura. 2025. "Elite Misperceptions in Foreign Policy." Working paper: `https://jkertzer.sites.fas.harvard.edu/Research_files/Elite-Misperceptions-Web.pdf`.

Kessler, Anke S. 1998. "The Value of Ignorance." *RAND Journal of Economics* 29(2):339–354.

Luzin, Pavel. 2017. "The Political Logic of the Zapad Exercises." European Leadership Network: `https://europeanleadershipnetwork.org/commentary/the-political-logic-of-the-zapad-exercises/`.

Meirowitz, Adam and Anne Sartori. 2008. "Strategic uncertainty as a cause of war." *Quarterly Journal of Political Science* 3(4):327–352.

Powell, Robert. 2002. "Bargaining theory and international conflict." *Annual Review of Political Science* 5(1):1–30.

Ramsay, Kristopher W. 2017. "Information, uncertainty, and war." *Annual Review of Political Science* 20(1):505–527.

Ravid, Doron. 2020. "Ultimatum Bargaining with Rational Inattention." *American Economic Review* 110(9):2948–2963.

Ravid, Doron, Anne-Katrin Roesler and Balázs Szentes. 2022. "Learning Before Trading: on the Inefficiency of Ignoring Free Information." *Journal of Political Economy* 130(2):346–387.

Reed, William. 2003. "Information, Power, and War." *American Political Science Review* 97(4):633–641.

Reiter, Dan. 2003. "Exploring the bargaining model of war." *Perspectives on Politics* 1(1):27–43.

Roesler, Anne-Katrin and Balázs Szentes. 2017. "Buyer-optimal learning and monopoly pricing." *American Economic Review* 107(7):2072–2080.

Smith, Bradley and William Spaniel. 2019. "Militarized disputes, uncertainty, and leader tenure." *Journal of Conflict Resolution* 63(5):1222–1252.

Spaniel, William and Bradley Smith. 2015. "Sanctions, uncertainty, and leader tenure." *International Studies Quarterly* 59(4):735–749.

Tchaouchev, Denis. 2025. "Learning When to Fight: Technological Change and Conflict in the Gunpowder Revolution." Working paper: `https://dtchaouchev.github.io/sieges.pdf`.

Thyne, Clayton. 2012. "Information, commitment, and intra-war bargaining: The effect of governmental constraints on civil war duration." *International Studies Quarterly* 56(2):307–321.

Uzonyi, Gary and Matthew Wells. 2016. "Domestic institutions, leader tenure and the duration of civil war." *Conflict Management and Peace Science* 33(3):294–310.

APPENDIX (online only)

A Omitted proofs

A.1 Proof of Lemma 1

First, if $x \notin [\underline{w}_R, \overline{w}_R]$, then either $x < \underline{w}_R$, or $x > \overline{w}_R$. In the first case, $\Pr(w_R \leq x) = 0 = 1 - \Pr(w_R \geq x)$, so $\phi^+(x) > 0$ and $\phi^-(x) = 0$ by construction. Thus, $\phi(x) = 0$. In the second case, $\Pr(w_R \geq x) = 0 = 1 - \Pr(w_R \leq x)$, so $\phi^+(x) = 0$ and $\phi^-(x) > 0$, implying $\phi(x) = 0$.

Second, we prove that $\phi^-(x) = \frac{(x-\underline{w}_R)^2}{2\delta}$, the functional form of ϕ^+ can be derived in an identical way. By definition,

$$\begin{aligned}
\phi^-(x) &= \Pr(w_R \leq x)\mathbb{E}[x - w_R | w_R \leq x] \\
&= G(x)\left(x - \mathbb{E}[w_R | w_R \leq x]\right) \\
&= G(x)\left(x - \frac{\underline{w}_R + x}{2}\right) \\
&= \left(\frac{x - \underline{w}_R}{\delta}\right)\left(\frac{x - \underline{w}_R}{2}\right).
\end{aligned}$$

A.2 Proof of Proposition 1

We derive state P's optimal proposal when state R surely learns, x_1^*. The remainder of the proof follows straightforwardly from the arguments made in the main text.

Claim A.1. *If state R surely learns in equilibrium σ, then state P proposes x_1^* with probability one, where*

$$x_1^* = \min\left\{x_0^* + \left(\frac{1}{2-\lambda}\right)\left(\bar{c}_P + \bar{c}_R - \frac{\delta}{2}\right), \overline{w}_R\right\}.$$

Proof. Consider an equilibrium σ in which R surely learns. Recall that if R chooses $l = 1$, then P's expected payoffs are $U_P(x|l = 1)$ from Equation 4.

Taking the derivative with respect to x gives

$$\begin{aligned}
\frac{\partial U_P}{\partial x}(x|l = 1) &= (1 - x)g(x) - \int_{\underline{w}_R}^{x} g(w_R)d(w_R) - w_P(x)g(x) \\
&= (1 - x - w_P(x))g(x) - G(x).
\end{aligned}$$

We can use the definition of x_0^* to write $G(w_R) = \frac{w_R - x_0^* + \frac{\delta}{2}}{\delta}$ and $w_P(w_R) = \lambda(1 - w_R) + (1 - \lambda)(1 - x_0^*) - (\bar{c}_R + \bar{c}_P)$ from Remark 1. Substituting these functions into expressions above, gives

$$\frac{\partial U_P}{\partial x}(x|l=1) = (1-x-\lambda(1-x)-(1-\lambda)(1-x_0^*)+\bar{c}_R+\bar{c}_P)\frac{1}{\delta} - \frac{x-x_0^*+\frac{\delta}{2}}{\delta}$$
$$= \frac{(1-\lambda)(x_0^*-x)+\bar{c}_R+\bar{c}_P}{\delta} - \frac{x-x_0^*+\frac{\delta}{2}}{\delta}$$
$$= \frac{(2-\lambda)(x_0^*-x)+\bar{c}_R+\bar{c}_P-\frac{\delta}{2}}{\delta}.$$

It is clear that U_P is strictly concave in x, so the first-order condition will be necessary and sufficient to characterize P's optimal interior offer:

$$\frac{\partial U_P}{\partial x}(x|l=1) = 0 \iff (2-\lambda)(x_0^*-x)+\bar{c}_R+\bar{c}_P-\frac{\delta}{2} = 0$$
$$\iff x = x_0^* + \left(\frac{1}{2-\lambda}\right)\left(\bar{c}_P+\bar{c}_R-\frac{\delta}{2}\right).$$

Note that the optimal interior offer is larger than $\underline{w}_R = x_0 - \frac{\delta}{2}$. Thus, if the optimal interior offer is smaller than $\overline{w}_R$, then P proposes it. If the optimal interior offer is larger than $\overline{w}_R$, then P proposes $\overline{w}_R$. □

A.3 Proof of Proposition 2

To show that the conditions are sufficient, suppose $\delta < 2(\bar{c}_P + \bar{c}_R)$ and $k_R \in (\phi(x_1^*), \phi(x_0^*))$. Given a probability of learning σ_R, we can write state P's expected payoff from proposal x as

$$U_P(x|\sigma_R) = \sigma_R U_P(x|l=1) + (1-\sigma_R)U_P(x|l=0).$$

Recall, $U_P(x|l=1)$ and $U_P(x|l=0)$ are single-peaked in x with unique maximizers x_1^* and x_0^*, respectively. Thus, a maximizer of $U_P(x|\sigma_R)$ must be in the interval $[x_0^*, x_1^*]$. In addition, $U_P(x|l=1)$ is strictly concave over $[x_0^*, x_1^*]$ and $U_P(x|l=0)$ is strictly decreasing on the interval, and therefore concave. Thus, $U_P(x|\sigma_R)$ is strictly concave over $[x_0^*, x_1^*]$ in a mixed strategy equilibrium because $\sigma_R > 0$, which means $U_P(x|\sigma_R)$ has a unique maximizer in $[x_0^*, x_1^*]$. Call this x_m^*.

To characterize what state R's learning probability will be in a mixed strategy equilibrium, first note that if $x_m^* \in \{x_0^*, x_1^*\}$, then R's best response to is either learn or not with probability one because $k_R \in (\phi(x_1^*), \phi(x_0^*))$. Thus, in any mixed strategy equilibrium, $x_m^* \in (x_0^*, x_1^*)$, so P's first-order condition must hold. Differentiating U_P with respect to x

gives us

$$\begin{aligned}
\frac{\partial U_P}{\partial x}(x|\sigma_R) &= \sigma_R \frac{\partial U_P}{\partial x}(x|l=1) + (1-\sigma_R)\frac{\partial U_P}{\partial x}(x|l=0) \\
&= \sigma_R \left[(1-x-w_P(x))g(x) - G(x)\right] - (1-\sigma_R) \\
&= \sigma_R \left(\frac{(2-\lambda)(x_0^* - x) + \bar{c}_R + \bar{c}_P - \frac{\delta}{2}}{\delta} \right) - (1-\sigma_R) \\
&= \sigma_R \left(\frac{(2-\lambda)(x_1^* - x) + (2-\lambda)(x_0^* - x_1^*) + \bar{c}_R + \bar{c}_P - \frac{\delta}{2}}{\delta} \right) - (1-\sigma_R) \\
&= \sigma_R \left(\frac{2-\lambda}{\delta} \right)(x_1^* - x) - (1-\sigma_R)
\end{aligned}$$

The second and third inequalities are just invoking the properties of $\frac{\partial U_P}{\partial x}$ established in Proposition 1. We can now write P's first-order condition as

$$\begin{aligned}
\frac{\partial U_P}{\partial x}(x_m^*|\sigma_R) = 0 &\iff \sigma_R\,(2-\lambda)\,(x_1^* - x_m^*) = \delta(1-\sigma_R) \\
&\iff \sigma_R = \frac{\delta}{\delta + (2-\lambda)(x_1^* - x_m^*)}.
\end{aligned}$$

Notice that σ_R is strictly increasing in x_m^* and bounded above by 1. Moreover, as $x_m^* \to x_0^*$, $\sigma_R \to \frac{2\delta}{2(\bar{c}_P + \bar{c}_R) + \delta} \in (0,1)$. Thus, for all $x_m^* \in (x_0^*, x_1^*)$, there exists $\sigma_R^* \in (0,1)$ such that $\frac{\partial U_P}{\partial x}(x_m^*|\sigma_R^*) = 0$.

To characterize what proposal x_m^* will appear, it suffices to look at R's indifference condition:

$$\phi(x_m^*) = k_R.$$

Recall ϕ is continuous and decreasing on the interval $[x_0^*, x_1^*]$. Moreover, $\phi(x_1^*) < k_2 < \phi(x_0^*)$ by assumption, so there exists x_m^* such that $\phi(x_m^*) = k_R$.

To show that the conditions are necessary for a mixed equilibrium of this type: if $\delta \geq 2(\bar{c}_P + \bar{c}_R)$, then Lemma 2 implies state R cannot surely propose $x_m^* \in (x_0^*, x_1^*)$ in a mixed strategy equilibrium. If $k_R \geq \phi(x_0^*)$ and $\delta < 2(\bar{c}_P + \bar{c}_R)$, then R will never learn if P proposes any $x_m^* \in (x_0^*, x_1^*)$ because ϕ is uniquely maximized at x_0^*. If $k_R \leq \phi(x_1^*)$ and $\rho < 1$, then $x_1^* \in (\underline{w}_R, \overline{w}_R)$, or else $k_R \leq \phi(x_1^*) = 0$, which is a contradiction. Then state R will surely learn if P proposes any $x_m^* \in (x_0^*, x_1^*)$ because ϕ will be strictly decreasing on the interval $(x_0^*, x_1^*) \subset (x_0^*, \overline{w}_R)$.

A.4 Proof of Proposition 3

To show that the conditions are sufficient, suppose $\delta > 2(\bar{c}_P + \bar{c}_R)$ and $k_R \in (\phi(x_1^*), \phi(x_0^*))$. Given a probability of learning σ_R, we can write state P's expected payoff from proposal x

as

$$U_P(x|\sigma_R) = \sigma_R U_P(x|l=1) + (1-\sigma_R)U_P(x|l=0).$$

Recall, $U_P(x|l)$ is single-peaked in x with a unique maximizer at x_l^* for $l = 0, 1$. Thus, a maximizer of $U_P(x|\sigma_R)$ must be in the interval $[x_1^*, x_0^*]$. In addition $U_P(x|l=1)$ is strictly concave and decreasing on the interval $[x_1^*, x_0^*]$, whereas $U_P(x|l=0)$ is constant on $[x_1^*, x_0^*)$ and discontinuous at x_0^*. Thus, $\sigma_R > 0$ implies $U_P(x|\sigma_R) < U_P(x_1^*|\sigma_R)$ for all $x \in (x_1^*, x_0^*)$. As such, in any mixed strategy equilibrium P can only propose x_1^* or x_0^* with positive probability.

To characterize what state R's learning probability will be in a mixed strategy equilibrium, first note that if $\sigma_P(x_0^*) \in \{0, 1\}$, then R's best response to is either learn or not with probability one because $k_R \in (\phi(x_1^*), \phi(x_0^*))$. Thus, $\sigma_P(x_0^*) \in (0, 1)$, and P's indifference condition pins down R's probability of learning:

$$\underbrace{\sigma_R U_P(x_0^*|l=1) + (1-\sigma_R)U_P(x_0^*|l=0)}_{U_P(x_0^*|\sigma_R)} = \underbrace{\sigma_R U_P(x_1^*|l=1) + (1-\sigma_R)U_P(x_1^*|l=0)}_{U_P(x_1^*|\sigma_R)} \quad \text{(A.1)}$$

Recall that $x_1^* < x_0^*$, so R surely rejects x_1^* if they do not learn, which means $U_P(x_1^*|l = 0) = \mathbb{E}[w_P] = 1 - x_0^* - (\bar{c}_R + \bar{c}_P)$. In addition, R surely accepts x_0^* if they do not learn, so $U_P(x_0^*|l=0) = 1 - x_0^*$. Substituting these expressions in Equation A.1 and rearranging gives

$$(1-\sigma_R)(\bar{c}_R + \bar{c}_P) = \sigma_R \underbrace{\left(U_P(x_1^*|l=1) - U_P(x_0^*|l=1)\right)}_{\equiv D}. \quad \text{(A.2)}$$

We can rewrite the payoff difference D as

$$D = \int_{\underline{w}_R}^{x_1^*} [x_0^* - x_1^*]\, g(w_R)dw_R + \int_{x_1^*}^{x_0^*} [w_P(w_R) - (1-x_0^*)]\, g(w_R)dw_R$$

Recall $x_1^* < x_0^* < \overline{w}_R$, and we can show

$$x_1^* - \underline{w}_R = \frac{\delta}{2} - \frac{\delta}{2(2-\lambda)} + \frac{\bar{c}_R + \bar{c}_P}{2-\lambda} > 0.$$

So we can express D as

$$\begin{aligned}
D &= \frac{(x_0^*-x_1^*)(x_1^*-\underline{w}_R)}{\delta} - \frac{(1-x_0^*)(x_0^*-x_1^*)}{\delta} + \int_{x_1^*}^{x_0^*} w_P(w_R)\frac{1}{\delta}dw_R \\
&= \frac{(x_0^*-x_1^*)(x_1^*-\underline{w}_R-(1-x_0^*))}{\delta} + \int_{x_1^*}^{x_0^*} \left[\lambda(1-w_R)+(1-\lambda)(1-x_0^*)-(\bar{c}_R+\bar{c}_P)\right]\frac{1}{\delta}dw_R \\
&= \frac{(x_0^*-x_1^*)(x_1^*-\underline{w}_R-(1-x_0^*))}{\delta} + \frac{(x_0^*-x_1^*)}{\delta}\left(1-(\bar{c}_R+\bar{c}_P)-\frac{x_0^*(2-\lambda)+x_1^*\lambda}{2}\right) \\
&= \frac{(x_0^*-x_1^*)}{\delta}\left(x_1^*-\underline{w}_R-(\bar{c}_P+\bar{c}_R)+\frac{\lambda(x_0^*-x_1^*)}{2}\right).
\end{aligned}$$

Recall $\underline{w}_R = x_0^* - \frac{\delta}{2}$ and we can use the definition of x_1^* to get $\bar{c}_P+\bar{c}_R = \frac{\delta}{2}-(2-\lambda)(x_0^*-x_1^*)$. Substituting these expressions above, gives us

$$D = \frac{(x_0^*-x_1^*)}{\delta}\left(\frac{(2-\lambda)(x_0^*-x_1^*)}{2}\right) = \frac{(2-\lambda)(x_0^*-x_1^*)^2}{2\delta}.$$

Now substitute D into state P's indifference condition in Equation A.2 to get

$$(1-\sigma_R)(\bar{c}_R+\bar{c}_P) = \sigma_R\frac{(2-\lambda)(x_0^*-x_1^*)^2}{2\delta} \iff \sigma_R = \frac{2\delta(\bar{c}_P+\bar{c}_R)}{(2-\lambda)(x_0^*-x_1^*)^2+2\delta(\bar{c}_P+\bar{c}_R)}$$

It is straightforward to verify that $\sigma_R \in (0,1)$.

To pin down state P's probability of proposing x_0^*, we work with state R's indifference condition:

$$\sigma_P(x_0^*)U_R(0|x_0^*) + (1-\sigma_P(x_0^*))U_R(0|x_1^*) = \sigma_P(x_0^*)U_R(1|x_0^*) + (1-\sigma_P(x_0^*))U_R(1|x_1^*)$$

which is equivalent to

$$\sigma_P(x_0^*)\left(U_R(0|x_0^*)-U_R(1|x_0^*)\right) = (1-\sigma_P(x_0^*))\left(U_R(1|x_1^*)-U_R(0|x_1^*)\right).$$

Because $x_1^* < x_0^*$, by definition of ϕ, $U_R(1|x_1^*) - U_R(0|x_1^*) = \phi^-(x_1^*) - k_R$. In addition, $U_R(0|x_0^*) - U_R(1|x_0^*) = k_R - \phi^-(x_0^*)$ (because $\phi^-(x_0^*) = \phi^+(x_0^*) = \phi(x_0^*)$) Making these substitutions gives us

$$\sigma_P(x_0^*)\left(k_R-\phi^-(x_0^*)\right) = (1-\sigma_P(x_0^*))\left(\phi^-(x_1^*)-k_R\right).$$

Solving for $\sigma_P(x_0^*)$ then gives

$$\sigma_P(x_0^*) = \frac{\phi^-(x_1^*)-k_R}{\phi^-(x_1^*)-\phi^-(x_0^*)} = \frac{\phi(x_1^*)-k_R}{\phi(x_1^*)-\phi(x_0^*)}.$$

To show that the conditions are necessary for a mixed equilibrium of this type: if $\delta \leq 2(\bar{c}_P + \bar{c}_R)$, then Lemma 2 implies state R surely proposes some x between $x_0^* \leq x_1^*$. Now assume $\delta > 2(\bar{c}_P + \bar{c}_R)$. If state P is mixing, then they must be mixing between x_0^* and x_1^* by Lemma 2. When state P mixes between x_0^* and x_1^*, then R's expected value of learning is

$$\sigma_P(x_0^*) \underbrace{(U_R(1|x_0^*) - U_R(0|x_0^*))}_{=\phi(x_0^*)} + (1 - \sigma_P(x_0^*)) \underbrace{(U_R(1|x_1^*) - U_R(0|x_1^*))}_{=\phi(x_1^*)}.$$

If $k_R < \phi(x_1^*)$, then R will surely learn because its value of learning is at least as large as $\phi(x_1^*)$. If $k_R > \phi(x_0^*)$, then R will never learn because its value of learning is smaller than $\phi(x_0^*)$.

A.5 Proof of Proposition 4

Claim A.2. *The probability of learning decreases in the cost of learning.*

Proof. First, focus on on the probability of learning, σ_R^*. If $k_R \leq \phi(x_1^*)$, then only the pure-strategy equilibrium with learning exists and $\sigma_R^* = 1$, assuming we select the pure strategy equilibrium when $k_R = \phi(x_1^*)$. If $k_R \geq \phi(x_0^*)$, then only the pure-strategy equilibrium with no learning exists and $\sigma_R^* = 0$, assuming we select the pure strategy equilibrium when $k_R = \phi(x_0^*)$. If $k_R \in (\phi(x_1^*), \phi(x_0^*))$, then Propositions 2 and 3 characterize the mixing probabilities.

Consider the case when $\delta > 2(\bar{c}_P + \bar{c}_R)$. Then Proposition 3 implies that $\sigma_R^* \in (0, 1)$, and σ_R^* is independent of k_R, i.e., $\frac{\partial \sigma_R^*}{\partial k_R} = 0$. Thus, $\delta > 2(\bar{c}_P + \bar{c}_R)$ implies σ_R^* is decreasing in k_R.

Consider the case when $\delta < 2(\bar{c}_P + \bar{c}_R)$. Then Proposition 3 implies that $\sigma_R^* \in (0, 1)$. To see that σ_R^* is decreasing in k_R when $k_R \in (\phi(x_1^*), \phi(x_0^*))$, first note that the characterization of σ_R^* implies that σ_R^* is increasing as a function of x_m^*, i.e., $\frac{\partial \sigma_R^*}{\partial x_m^*} > 0$. Second, use the implicit function theorem and the characterization of x_m^* to get

$$\frac{\partial x_m^*}{\partial k_R} = -\frac{-1}{\frac{\partial \phi}{\partial x_m^*}},$$

which is less than zero because $\frac{\partial \phi}{\partial x_m^*} < 0$ when $x_m^* \in (x_0^*, x_1^*)$. Thus, if $k_R \in (\phi(x_1^*), \phi(x_0^*))$, then $\sigma_R^* \in (0, 1)$ and σ_R^* is strictly decreasing in k_R. We conclude that $\delta < 2(\bar{c}_P + \bar{c}_R)$ implies σ_R^* is decreasing in k_R. □

Claim A.3. *1. The probability of war is positive and constant when $k_R < \phi(x_1^*)$, and it is zero and constant when $k_R > \phi(x_0^*)$.*

2. *If uncertainty is small* ($\delta < 2(\bar{c}_P + \bar{c}_R)$) *and* $k_R \in [\phi(x_1^*), \phi(x_0^*))$, *then it is continuous and strictly increasing.*

3. *If uncertainty is large* ($\delta > 2(\bar{c}_P + \bar{c}_R)$), *it is continuous and strictly decreasing when* $k_R \in (\phi(x_1^*), \phi(x_0^*))$, *but with a discontinuous increase at* $k_R = \phi(x_1^*)$.

Proof. For (1), we consider the pure strategy equilibria. If $k_R \leq \phi(x_1^*)$, then $\phi(x_1^*) > 0$, which means $x_1^* \in (\underline{w}_R, \overline{w}_R)$. Furthermore, R learns with probability 1, and P always proposes x_1^*, which does not depend on k_R. War therefore occurs if $w_R > x_1^*$, and the probability of war is

$$\Pr(War \mid k_R \leq \phi(x_1^*)) = \Pr(w_R > x_1^*) = 1 - G(x_1^*) = 1 - \frac{2(\bar{c}_P + \bar{c}_R) - \delta(1-\lambda)}{2\delta(2-\lambda)},$$

which is greater than 0 because $x_1^* \in (\underline{w}_R, \overline{w}_R)$.

When the learning cost is sufficiently high ($k_R > \phi(x_0^*)$), R never learns at equilibrium and thus the two sides have symmetric information structure. P's proposal x_0^* is always accepted and thus the probability of war is $\Pr(War \mid k_R \geq \phi(x_0^*)) = 0$.

For (2), if $\phi(x_1^*) < k_R < \phi(x_0^*)$ and $\delta < 2(\bar{c}_P + \bar{c}_R)$, Proposition 2 implies that R learns with probability strictly between zero and one, and P proposes $x_m^* \in (x_0^*, x_1^*)$ for sure. Thus war occurs only when R learns and finds himself strong enough:

$$\begin{aligned}\Pr(War) &= \sigma_R^* \Pr(w_R > x_m^*) \\ &= \frac{\delta}{\delta + (2-\lambda)(x_1^* - x_m^*)}(1 - G(x_m^*)) \\ &= \frac{\delta}{\delta + (2-\lambda)(x_1^* - x_m^*)}\left(\frac{\sqrt{2}k_R}{\sqrt{\delta k_R}}\right) \\ &= \frac{2\sqrt{2\delta k_R}}{2(\bar{c}_P + \bar{c}_R) + 2(2-\lambda)\sqrt{2\delta k_R} - \delta(1-\lambda)}.\end{aligned}$$

Recall from the proof of Claim A.2, σ_R^* is strictly decreasing in k_R. Also from Claim A.2, x_m^* is decreasing in k_R, so $1 - G(x_m^*) = \Pr(war|l = 1)$ is increasing in k_R. To see what effect dominates, define $\omega = \sqrt{2\delta k_R}$. Straightforward differentiation reveals that the above expression in increasing in ω, which is strictly increasing in k_R.

To see that the war probability is continuous at the lower threshold of learning cost $\phi(x_1^*)$, note that σ_R^* approaches one and x_m^* approaches x_1^* as k_R tends to $\phi(x_1^*)$ from above. Both the probability of learning and the probability of war conditional on learning take finite values, so the limit of product is the product of the limits:

$$\begin{aligned}\lim_{k_R \downarrow \phi(x_1^*)} \Pr(War) &= \lim_{k_R \downarrow \phi(x_1^*)} \sigma_R^*(1 - G(x_m^*)) \\ &= \left(\lim_{k_R \downarrow \phi(x_1^*)} \sigma_R^*\right)\left(1 - \lim_{k_R \downarrow \phi(x_1^*)} G(x_m^*)\right) \\ &= 1 - G(x_1^*) \\ &= \Pr(w_R > x_1^*) = \Pr(War \mid k_R = \phi(x_1^*)).\end{aligned}$$

In the third inequality we are invoking that G is the uniform cdf so it is continuous.

For (3), if $\phi(x_1^*) < k_R < \phi(x_0^*)$ and $\delta > 2(\bar{c}_P + \bar{c}_R)$, Proposition 3 implies that P mixes between proposing x_0^* and x_1^* and R learns with probability strictly between zero and one. If R does not learn, then it surely rejects $x_1^* < x_0^* = \mathbb{E}[w_R]$. If state R learns, then it rejects if w_R is greater than the proposed offer. In this case, we have

$$\Pr(War) = (1 - \sigma_R^*)\sigma_P^*(x_1^*) + \sigma_R^*(\sigma_P^*(x_1^*) \Pr(w_R > x_1^*) + \sigma_P^*(x_0^*) \Pr(w_R > x_0^*)). \tag{A.3}$$

To see how changes in k_R affect the probability of war, note that $\frac{\partial x_l^*}{\partial k_R} = 0$, for $l \in \{0, 1\}$. Furthermore, Proposition 3 implies $\frac{\partial \sigma_R^*}{\partial k_R} = 0$ and

$$\frac{\partial \sigma_P(x_0^*)}{\partial k_R} = \frac{1}{\phi(x_0^*) - \phi(x_1^*)} = -\frac{\partial \sigma_P(x_1^*)}{\partial k_R} > 0.$$

We can then differentiate the probability of war in Equation A.3 with respect to learning cost:

$$\begin{aligned}\frac{\partial \Pr(War)}{\partial k_R} &= (1 - \sigma_R^*)\frac{\partial}{\partial k_R}\sigma_P^*(x_1^*) + \sigma_R^*\left(\Pr(w_R > x_1^*)\frac{\partial}{\partial k_R}\sigma_P^*(x_1^*) + \Pr(w_R > x_0^*)\frac{\partial}{\partial k_R}\sigma_P^*(x_0^*)\right) \\ &= (1 - \sigma_R^*)\frac{-1}{\phi(x_0^*) - \phi(x_1^*)} + \sigma_R^*\left(\Pr(w_R > x_1^*)\frac{-1}{\phi(x_0^*) - \phi(x_1^*)} + \Pr(w_R > x_0^*)\frac{1}{\phi(x_0^*) - \phi(x_1^*)}\right) \\ &= (1 - \sigma_R^*)\frac{-1}{\phi(x_0^*) - \phi(x_1^*)} + \sigma_R^* \Pr(x_1^* < w_R \leq x_0^*)\frac{-1}{\phi(x_0^*) - \phi(x_1^*)} < 0.\end{aligned}$$

Next we show that the war probability jumps up at the lower threshold of learning cost $\phi(x_1^*)$ and thus only piecewise decreasing in the cost of learning. To see this, first note that Propositon 3 implies that

$$\lim_{k_R \downarrow \phi(x_1^*)} \sigma_P^*(x_0^*) = 0 = 1 - \lim_{k_R \downarrow \phi(x_1^*)} \sigma_P^*(x_1^*).$$

Applying these limits to Equation A.3 gives us

$$\begin{aligned}
\lim_{k_R \downarrow \phi(x_1^*)} \Pr(War) &= (1-\sigma_R^*)\left(\lim_{k_R \downarrow \phi(x_1^*)} \sigma_P^*(x_1^*)\right) \\
&\quad + \sigma_R^*\left(\Pr(w_R > x_1^*) \lim_{k_R \downarrow \phi(x_1^*)} \sigma_P^*(x_1^*) + \Pr(w_R > x_0^*) \lim_{k_R \downarrow \phi(x_1^*)} \sigma_P^*(x_0^*)\right) \\
&= (1-\sigma_R^*) + \sigma_R^* \Pr(w_R > x_1^*) \\
&= \Pr(w_R > x_1^*) + (1-\sigma_R^*)\Pr(w_R \leq x_1^*) \\
&> \Pr(w_R > x_1^*) = \Pr(War \mid k_R < \phi(x_1^*))
\end{aligned}$$

Thus, we conclude that the probability of war discontinuously jumps at $k_R = \phi(x_1^*)$. □

A.6 Proof of Proposition 5

To see (1), first note that the value of learning ϕ does not depends on λ. Thus, we can write

$$\frac{d\phi(x_1^*)}{d\lambda} = \frac{\partial \phi(x_1^*)}{\partial x_1^*}\frac{\partial x_1^*}{\partial \lambda}.$$

From Lemma 1, we know $\frac{\partial \phi(x_1^*)}{\partial x_1^*} > 0$ if and only if $x_1^* < x_0^*$, which is the case when $\delta < 2(\bar{c}_P + \bar{c}_R)$. For $\frac{\partial x_1^*}{\partial \lambda}$, consider three cases. If $\delta < \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$, then $x_1^* = \overline{w}_R$ and $\frac{\partial x_1^*}{\partial \lambda} = 0 = \frac{d\phi(x_1^*)}{d\lambda}$. If $\delta \in \left(\frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}, 2(\bar{c}_P + \bar{c}_R)\right)$, then $x_1^* \in (x_0^*, \overline{w}_R)$, and $\frac{\partial x_1^*}{\partial \lambda} > 0$. Thus, $\frac{d\phi(x_1^*)}{d\lambda} < 0$. If $\delta > 2(\bar{c}_P + \bar{c}_R)$, then $x_1^* \in (\underline{w}_R, x_0^*)$, which means $\frac{\partial x_1^*}{\partial \lambda} < 0$. Thus, $\frac{d\phi(x_1^*)}{d\lambda} < 0$. In every case, $\phi(x_1^*)$ is at least weakly decreasing in λ, and it is strictly decreasing if $\delta > \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$.

From the equilibrium characterization, we know that the full learning equilibrium is attained when $k_R < \phi(x_1^*)$ and whichever type of mixed learning equilibrium is attained when $k_R \in (\phi(x_1^*), \phi(x_0^*))$. Thus (2) immediately follows from (1).

To see (3), first consider the low uncertainty case ($\delta < 2(\bar{c}_P + \bar{c}_R)$) and recall that the learning probability at the mixed strategy equilibrium is given by

$$\sigma_R^* = \frac{\delta}{\delta + (2-\lambda)(x_1^* - x_m^*)}.$$

Note that x_0^*, x_m^* and $(2-\lambda)(x_1^* - x_0^*)$ are constant in λ. First note that

$$(2-\lambda)(x_1 - x_0) = \bar{c}_P + \bar{c}_R - \frac{\delta}{2},$$

which does not depend on λ. Thus, we have

$$\frac{\partial}{\partial \lambda}(2-\lambda)(x_1^* - x_m^*) = \frac{\partial}{\partial \lambda}(2-\lambda)(x_0^* - x_m^*) = (x_0^* - x_m^*)\frac{\partial}{\partial \lambda}(2-\lambda) > 0.$$

Therefore, $\frac{\partial \sigma_R^*}{\partial \lambda} < 0$.

For high uncertainty case ($\delta > 2(\overline{c}_P + \overline{c}_R)$), recall that the learning probability at the mixed strategy equilibrium is given by

$$\sigma_R^* = \frac{2\delta(\overline{c}_P + \overline{c}_R)}{(2-\lambda)(x_0^* - x_1^*)^2 + 2\delta(\overline{c}_P + \overline{c}_R)}$$

Note that $2\delta(\overline{c}_P + \overline{c}_R)$ is constant in λ. For the remaining part, we have

$$(2-\lambda)(x_0^* - x_1^*)^2 = \frac{1}{2-\lambda}\left(\frac{\delta}{2} - \overline{c}_P - \overline{c}_R\right)^2,$$

and the derivative of the right-hand side with respect to λ is positive. Thus we have $\frac{\partial \sigma_R^*}{\partial \lambda} < 0$.

A.7 Proof of Proposition 6

The next claim establishes the comparative statics of how δ affects $\phi(x_1^*)$, assuming $\delta < 2(\bar{c}_P + \bar{c}_R)$.

Claim A.4. *If $\delta < \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$, then $\phi(x_1^*) = 0$ and $\phi(x_1^*)$ is constant in δ. If $\delta \in \left(\frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}, 2(\bar{c}_P + \bar{c}_R)\right)$, then $\phi(x_1^*) > 0$ and $\phi(x_1^*)$ is strictly increasing in δ.*

Proof. If $\delta < 2(\bar{c}_P + \bar{c}_R)$, $x_1^* > x_0^*$, so $\phi(x_1^*) = \phi^+(x_1^*)$. We want to compute the total derivative:

$$\frac{d\phi^+}{d\delta}(x_1^*) = \frac{\partial \phi^+}{\partial \delta}(x_1^*) + \frac{\partial \phi^+}{\partial x}(x_1^*)\frac{\partial x_1^*}{\partial \delta}. \tag{A.4}$$

First note that Lemma 1 implies that

$$\frac{\partial \phi^+}{\partial \delta}(x_1^*) = \begin{cases} \left(\frac{\overline{w}_R - x_1^*}{2\delta}\right)\left(1 - \frac{\overline{w}_R - x_1^*}{\delta}\right) & \text{if } x_1^* < \overline{w}_R \\ 0 & \text{if } x_1^* \geq \overline{w}_R \end{cases}$$

and

$$\frac{\partial \phi^+}{\partial x}(x_1^*) = \min\left\{-\frac{\overline{w}_R - x_1^*}{\delta}, 0\right\}.$$

Second, Equation 5 implies that $x_1^* = \overline{w}_R$ if and only if $\delta \leq \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$. Thus, $\delta < \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$ implies $x_1^* = \overline{w}_R$, in which case $\frac{d\phi^+(x_1^*)}{d\delta} = \frac{\partial \phi^+}{\partial x}(x_1^*) = 0$. Substituting these values in Equation A.4 shows that $\phi(x_1^*)$ is therefore constant.

Now assume $\delta \in \left(\frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}, 2(\bar{c}_P+\bar{c}_R)\right)$. Then $x_0^* < x_1^* < \overline{w}_R$, and

$$\frac{\partial x_1^*}{\partial \delta} = -\frac{1}{2(2-\lambda)}.$$

Using Equation A.4 and the partial derivatives above, we can write the total derivative as

$$\frac{d\phi^+}{dx_1^*}(x_1^*) = \left(\frac{\overline{w}_R - x_1^*}{2\delta}\right)\left[\left(1 - \frac{\overline{w}_R - x_1^*}{\delta}\right) + \frac{1}{2-\lambda}\right].$$

A sufficient condition for $\frac{d\phi^+}{dx_1^*}(x_1^*) > 0$ is

$$\begin{aligned} 1 > \frac{\overline{w}_R - x_1^*}{\delta} \iff 1 &> \frac{2(\bar{c}_P+\bar{c}_R) + \delta(1-\lambda)}{2\delta(2-\lambda)} \\ \iff \delta &> \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}, \end{aligned}$$

which holds by assumption. □

The next claim establishes the local comparative statics that, within the full-learning equilibrium, increases in δ increase the probability of war. Note that this claim does not impose the assumption that $\delta < 2(\bar{c}_P+\bar{c}_R)$.

Claim A.5. *In the always learning equilibrium, the probability of war is increasing in δ.*

Proof. First, assume $k_R < \phi(x_1^*)$, which is equivalent to the always learning equilibirum existing. Because $k_R > 0$, $x_1^* < \overline{w}_R$ (i.e., $\delta > \frac{2(\bar{c}_P+\bar{c}_R)}{3-\lambda}$) because if not $\phi(x_1^*) = 0 < k_R$. In the always learning equilibrium, the probability of war is

$$\Pr(war) = \Pr(w_R > x_1^*) = 1 - G(x_1^*).$$

We want to compute the total derivative:

$$\frac{d\Pr(war)}{d\delta} = -\frac{\partial G}{\partial \delta}(x_1^*) - g(x_1^*)\frac{\partial x_1^*}{\partial \delta},$$

where g is the pdf of cdf G. Because $x_1^* < \overline{w}_R$, we know the following:

$$g(x_1^*) = \frac{1}{\delta} \qquad \frac{\delta x_1^*}{\partial \delta} = -\frac{1}{2(2-\lambda)} \qquad \frac{\partial G}{\partial \delta}(x_1^*) = \frac{x_0^* - x_1^*}{\delta^2}.$$

Substituting these expressions into the total derivative above gives us

$$\begin{aligned}\frac{d\,\Pr(war)}{d\delta} &= -\frac{x_0^* - x_1^*}{\delta^2} + \frac{1}{\delta}\left(\frac{1}{2(2-\lambda)}\right)\\ &= \frac{\bar{c}_P + \bar{c}_R - \frac{\delta}{2}}{(2-\lambda)\delta^2} + \frac{1}{\delta}\left(\frac{1}{2(2-\lambda)}\right)\\ &= \frac{\bar{c}_P + \bar{c}_R}{\delta^2(2-\lambda)} > 0.\end{aligned}$$

□

The next claim establishes the local comparative statics that, assuming that $\delta < 2(\bar{c}_P + \bar{c}_R)$, the probability of war is increasing in δ within the mixed strategy equilibrium.

Claim A.6. *If $\delta < 2(\bar{c}_P + \bar{c}_R)$, then the probability of war is increasing in δ in the mixed strategy equilibrium, i.e., $\phi(x_1^*) < k_R < \phi(x_0^*)$.*

Proof. In the equilibrium, the probability of war is

$$\Pr(war) = \sigma_R^*(1 - G(x_m^*))$$

Recall that $x_m^* = \overline{w}_R - \sqrt{2\delta k_R}$, so $1 - G(x_m^*) = 1 - \frac{\delta - \sqrt{2\delta k_R}}{\delta} = \frac{\sqrt{2}k_R}{\sqrt{\delta k_R}}$. Likewise, $\sigma_R^* = \frac{\delta}{\delta + (2-\lambda)(x_1^* - x_m^*)}$. Making these substitutions gives us

$$\Pr(war) = \frac{2\sqrt{2\delta k_R}}{2(\bar{c}_P + \bar{c}_R) + 2(2-\lambda)\sqrt{2\delta k_R} - \delta(1-\lambda)}.$$

This expression is increasing in $w = \sqrt{2\delta k_R}$. Moreover, fixing w, the denominator is decreasing in δ and is positive when $\delta < 2(\bar{c}_P + \bar{c}_R)$. Thus, $\Pr(war)$ is increasing in δ. □

To conclude the proposition, note that $\phi(x_0^*)$ is strictly increasing as a function of δ and obtains a maximum value at

$$\phi(x_0^*)|_{\delta = 2(\bar{c}_P + \bar{c}_R)} = \frac{\bar{c}_P + \bar{c}_R}{4}.$$

If $k_R \geq \frac{\bar{c}_P + \bar{c}_R}{4}$, then $k_R \geq \phi(x_0^*)$ for all $\delta \in (0, 2(\bar{c}_P + \bar{c}_R))$, which means the no learning equilibrium prevails for all δ and the probability of war is zero and constant as a function of δ.

If $k_R < \frac{\bar{c}_P + \bar{c}_R}{4}$, note that $\phi(x_0^*)|_{\delta=0} = 0$. Thus, the intermediate value theorem implies that there exists a $\delta' \in (0, 2(\bar{c}_P + \bar{c}_R))$ such that $\phi(x_0^*)|_{\delta=\delta'} = k_R$ — i.e., $\delta' = 8k_R$. Moreover, δ' is unique because $\phi(x_0^*)$ is strictly increasing in δ. For similar reasons, there exists a unique $\delta'' \in (\delta', 2(\bar{c}_P + \bar{c}_R))$ such that $\phi(x_1^*)|_{\delta=\delta''} = k_R$.

If $\delta < \delta'$, then the no learning equilibrium exists and the probability of war is zero and therefore constant in δ. If $\delta \in (\delta', \delta'')$. then the mixed equilibrium exists with a positive probability of learning, which is increasing over the range (δ', δ''). If $\delta \in (\delta'', 2(\bar{c}_P + \bar{c}_R))$, then the always learn equilibrium exists with a positive probability of learning, which is increasing over the range $(\delta'', 2(\bar{c}_P + \bar{c}_R))$. Moreover, the probability of war is continuous at δ'' due to the same logic in Claim A.3.2.

A.8 Proof of Proposition 7

Claim A.7. *If $\delta > 2(\bar{c}_P + \bar{c}_R)$ and $\lambda = 0$, then $\phi(x_1^*)$ is strictly increasing in δ.*

Proof. If $\delta > \delta^\dagger$, then $\underline{w}_R < x_1^* < x_0^*$, so $\phi(x_1^*) = \phi^-(x_1^*)$. We want to compute the total derivative:

$$\frac{d\phi^-}{d\delta}(x_1^*) = \frac{\partial\phi^-}{\partial\delta}(x_1^*) + \frac{\partial\phi^-}{\partial x}(x_1^*)\frac{\partial x_1^*}{\partial\delta}. \tag{A.5}$$

In a similar vein, we know the component expressions:

$$\begin{aligned}
\frac{\partial\phi^-}{\partial\delta}(x_1^*) &= \left(\frac{x_1^* - \underline{w}_R}{2\delta}\right)\left(1 - \frac{x_1^* - \underline{w}_R}{\delta}\right) \\
\frac{\partial\phi^-}{\partial x}(x_1^*) &= \frac{x_1^* - \underline{w}_R}{\delta}, \text{ and} \\
\frac{\delta x_1^*}{\partial\delta} &= -\frac{1}{2(2-\lambda)}.
\end{aligned}$$

Substituting these into Equation A.5 gives us

$$\frac{d\phi^-}{d\delta}(x_1^*) = \frac{x_1^* - \underline{w}_R}{2\delta}\left[\left(1 - \frac{x_1^* - \underline{w}_R}{\delta}\right) - \frac{1}{2-\lambda}\right].$$

Thus, the sign of the total derivative is pinned down by

$$\begin{aligned}
\frac{d\phi^-}{d\delta}(x_1^*) > 0 &\iff 1 - \frac{x_1^* - \underline{w}_R}{\delta} > \frac{1}{2-\lambda} \\
&\iff \frac{x_1^* - \underline{w}_R}{\delta} < 1 - \frac{1}{2-\lambda} \\
&\iff \frac{2(\bar{c}_P + \bar{c}_R) + \delta(1-\lambda)}{2\delta(2-\lambda)} < \frac{1-\lambda}{2-\lambda} \\
&\iff \frac{\bar{c}_P + \bar{c}_R}{\delta} + \frac{1-\lambda}{2} < 1 - \lambda \\
&\iff \frac{\bar{c}_P + \bar{c}_R}{\delta} < \frac{1-\lambda}{2} \\
&\iff \delta > \frac{2(\bar{c}_P + \bar{c}_R)}{1-\lambda}.
\end{aligned}$$

When $\lambda = 0$, the last expression is $\delta > 2(\bar{c}_P + \bar{c}_R)$, which holds by assumption. □

Claim A.8. *Assume $\lambda = 0$. If $\frac{\bar{c}_R + \bar{c}_R}{4} < k_R < \bar{\phi}_1(\bar{\delta}_0)$, then there exists δ^* such that (i) $2(\bar{c}_P + \bar{c}_R) < \delta^* < \bar{\delta}_0$ and (ii) $\phi(x_1^*)_{\delta=\delta^*} = k_R$.*

Proof. Recall that δ is bounded below at $2(\bar{c}_P + \bar{c}_R)$ and bounded above at $\bar{\delta}_0$. Moreover,

$$\phi(x_1^*)|_{\delta=2(\bar{c}_P+\bar{c}_R)} = \frac{\overline{c}_R + \overline{c}_R}{4} < k_R$$

and

$$\phi(x_1^*)|_{\delta=\bar{\phi}_1(\bar{\delta}_0)} > k_R.$$

Because ϕ is strictly increasing (Claim A.7), there exists a unique $\delta^* \in (2(\bar{c}_P + \bar{c}_R), \bar{\delta}_0)$ such $\phi(x_1^*)_{\delta=\delta^*} = k_R$. □

The remaining of the proof follows from Claim A.8, the fact that the equilibrium transitions from the partial learning equilibrium to the full learning equilibrium at δ^*, and at δ^* there is a discontinuous decrease in the probability of war due to the same logic in Claim A.3.3.

A.9 Proof of Proposition 8

Claim A.9. *If $\delta > 2(\bar{c}_P + \bar{c}_R)$ and $\lambda = 1$, then $\phi(x_1^*)$ is strictly decreasing in δ.*

Proof. Follows from the proof in Claim A.7 and the fact that $\lim_{\lambda\to 1} \frac{2(\bar{c}_P+\bar{c}_R)}{1-\lambda} = \infty$. □

Claim A.10. *Assume $\lambda = 1$. If $\delta > 2(\bar{c}_P + \bar{c}_R)$, then the probability of war is strictly increasing in δ in the mixed strategy equilibrium, i.e., and $\phi(x_1^*) < k_R < \phi(x_0^*)$.*

Proof. Propositions 1-3 imply that the probability of war is

$$\Pr(war) = \sigma_R^* \underbrace{[\sigma_P^*(x_0^*)(1 - G(x_0^*)) + (1 - \sigma_P(x_0^*))(1 - G(x_1^*))]}_{\Pr(war|l=1)} + (1 - \sigma_R^*) \underbrace{(1 - \sigma_P(x_0^*))}_{\Pr(war|l=0)}.$$

Recall that $G(x_0^*) = \frac{1}{2}$. Further assuming $\lambda = 1$ gives us

$$\begin{aligned} G(x_1^*) &= \frac{\bar{c}_P + \bar{c}_R}{\delta}, \text{ and} \\ \sigma_P(x_0^*) &= \frac{k_R - \phi^-(x_1^*)}{\phi^-(x_0^*) - \phi^-(x_1^*)} = \frac{4(\bar{c}_P + \bar{c}_R)^2 - 8\delta k_R}{(2(\bar{c}_P + \bar{c}_R) - \delta)(2(\bar{c}_P + \bar{c}_R) + \delta)}. \end{aligned}$$

Substituting these expressions in $\Pr(war \mid l = 1)$ gives us

$$\Pr(war \mid l = 1) = \frac{\bar{c}_P + \bar{c}_R - 4k_R + \delta}{2(\bar{c}_P + \bar{c}_R) + \delta}.$$

Finally, when $\lambda = 1$, R's learning probability is

$$\sigma_R^* = \delta \frac{8(\bar{c}_P + \bar{c}_R)}{(2(\bar{c}_P + \bar{c}_R) + \delta)^2}.$$

Putting all these expressions together gives us

$$\Pr(war) = \delta \frac{4(\bar{c}_P + \bar{c}_R) - 8k_R + \delta}{(2(\bar{c}_P + \bar{c}_R) + \delta)^2}.$$

Next we want to differentiate this expression with respect to δ:

$$\begin{aligned}\frac{\partial \Pr(war)}{\partial \delta} &= \frac{\delta}{(2(\bar{c}_P + \bar{c}_R) + \delta)^2} + \frac{4(\bar{c}_P + \bar{c}_R) - 8k_R + \delta}{(2(\bar{c}_P + \bar{c}_R) + \delta)^2} - 2\delta \frac{4(\bar{c}_P + \bar{c}_R) - 8k_R + \delta}{(2(\bar{c}_P + \bar{c}_R) + \delta)^3} \\ &= \left(\frac{1}{(2(\bar{c}_P + \bar{c}_R) + \delta)^2} \right) \underbrace{\left(\delta + 4(\bar{c}_P + \bar{c}_R) - 8k_R + \delta - 2\delta \frac{4(\bar{c}_P + \bar{c}_R) - 8k_R + \delta}{2(\bar{c}_P + \bar{c}_R) + \delta} \right)}_{\equiv X}\end{aligned}$$

Focus Expression X. The sign of X determines the sign of $\frac{\partial \Pr(War)}{\partial \delta}$. Recall that $\delta > 2(\bar{c}_P + \bar{c}_R)$. Thus, $2\delta > 2(\bar{c}_P + \bar{c}_R) + \delta$. As such X is increasing in k_R. We can establish a lower bound on X as

$$X > X|_{k_R=0} = \delta + 4(\bar{c}_P + \bar{c}_R) + \delta - 2\delta \frac{4(\bar{c}_P + \bar{c}_R) + \delta}{2(\bar{c}_P + \bar{c}_R) + \delta} = \frac{8(\bar{c}_P + \bar{c}_R)^2}{2(\bar{c}_P + \bar{c}_R) + \delta} > 0. \qquad \square$$

To complete the proof of the proposition, note that in this case δ is bounded below by $2(\bar{c}_P + \bar{c}_R)$, where

$$\phi(x_0^*)_{\delta=2(\bar{c}_P+\bar{c}_R)} = \phi(x_0^*)_{\delta=2(\bar{c}_P+\bar{c}_R)} = \frac{\bar{c}_P + \bar{c}_R}{4}.$$

Let $\bar{\delta}_1$ denote the theoretical limit of δ implied by Assumptions 1 and 2. Specifically, $\bar{\delta}_1 = 2\min\{x_0^*, (1 - x_0^*), \bar{p}, (1 - \bar{p})\} = 2\{x_0^*, 1 - \bar{p}\}$. Let $D_1^+ = [2(\bar{c}_P + \bar{c}_R), \bar{\delta}_1)$.

Consider two cases. First, assume $k_R \geq \frac{\bar{c}_P + \bar{c}_R}{4} = \phi(x_0^*)|_{\delta=2(\bar{c}_P+\bar{c}_R)}$. If $k_R \geq \phi(x_0^*)|_{\delta=\bar{\delta}_1}$, then, because $\phi(x_0^*)$ is strictly increasing in δ, $k_R \geq \phi(x_0^*)$ for all $\delta \in D_1^+$, so the equilibrium is always the pure strategy equilibrium. In this case, the probability of war is zero and constant. If $k_R < \phi(x_0^*)|_{\delta=\bar{\delta}_1}$, then there exists a unique $\delta' \in D_1^+$ such that $\delta < \delta'$ implies $k_R > \phi(x_0^*)$ and the equilibrium is the no learning equilibrium. If $\delta > \delta'$, $k_R < \phi(x_0^*)$, which means $k_R > \phi(x_1^*)$ because $\phi(x_1^*)$ is decreasing in δ (Claim A.9) and $k_R \geq \frac{\bar{c}_P + \bar{c}_R}{4}$. Thus, the partial learning equilibrium dominates with a positive probability of war. Moreover, Claim A.10 shows that the probability of war is increasing in δ.

Second, assume $k_R < \frac{\bar{c}_P + \bar{c}_R}{4}$. If $k_R \leq \phi(x_1^*)|_{\delta=\bar{\delta}_1}$, then $k_R \leq \phi(x_1^*)$ for all $\delta \in D_1^+$ because $\phi(x_1^*)$ is strictly decreasing in δ (Claim A.9). Thus, the equilibrium is always full learning equilibrium, in which case the probability of war is positive and strictly increasing in δ (Claim A.5). If $k_R < \phi(x_1^*)|_{\delta=\bar{\delta}_1}$, then there exists $\delta' \in D_1^+$ such that $\delta < \delta'$ implies that $k_R < \phi(x_1^*)$ so equilibrium is the full learning equilibrium and the probability of war is increasing in δ. For $\delta > \delta'$, $k_R > \phi(x_1^*)$. Because $k_R < \frac{\bar{c}_P + \bar{c}_R}{4} < \phi(x_0^*)$, the equilibrium is the partial learning equilibrium, and the probability of war is strictly increasing (Claim A.10). At $\delta = \delta'$, $k_R = \phi(x_1^*)|_{\delta=\delta'}$ and there is a discontinuous increase in the probability of war due to the same logic in Claim A.3.3.

B Additional figures and tables

Figure B.1: Uncertainty and the probability of war.

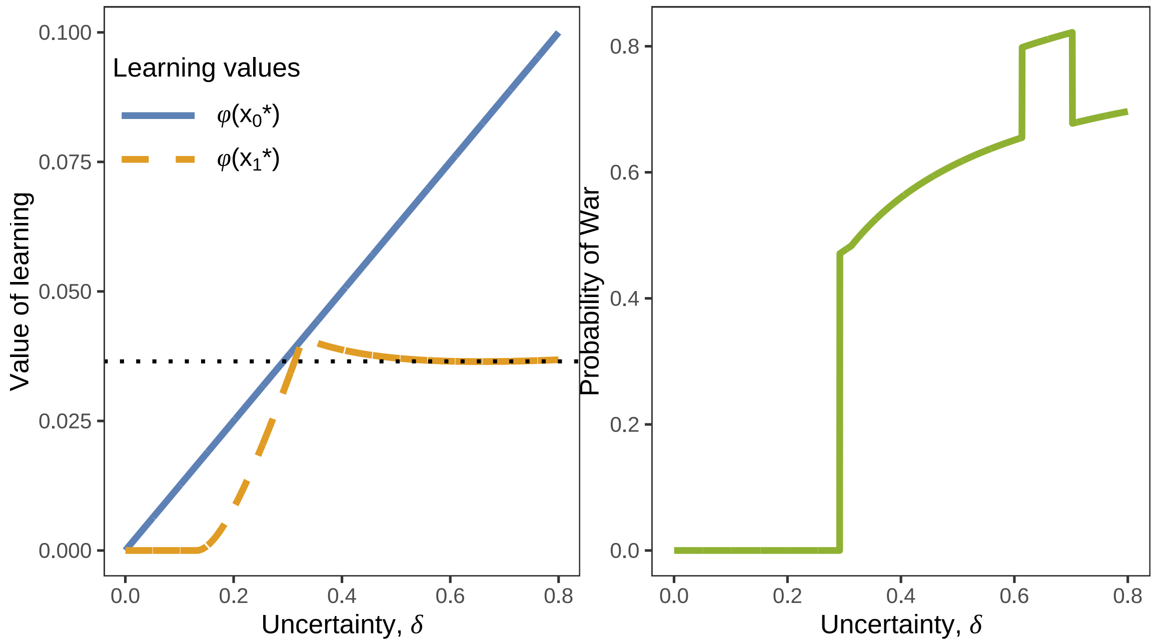

Notes: Left panel graphs the equilibrium values of learning as a function of uncertainty, δ, where we assume $\bar{p} = 0.5$, $c_P = \frac{9}{128}$, $c_R = \frac{3}{32}$, and $\lambda = 0.5$. Fixing learning costs $k_R = 0.0365$, the right panel graphs the resulting probability of war as a function of δ.