Learning to Fight: Information, Uncertainty, and the Risk of War *

Mark Fey†

March 6, 2015

Abstract

Uncertainty is a leading rationalist explanation for war. As a corollary, scholars have reasoned that if disputing countries acquire new information, say through observing the outcome of battles, then this will reduce uncertainty. This will, in turn, make peace more likely. We show that while this does hold in the very long run, in the short and medium run it need not hold. Specifically, we show that it is possible for an uncertain side to observe a finite number of noisy signals about relative strength and yet have the risk of war be higher after these signals than with no signals at all. Thus, we show that rational learning can lead to a higher risk of war. This can occur because acquiring new information not only reduces uncertainty, but it also shifts the expected outcome of the war which leads to a shift in the demands made.

*I thank Brad Smith for research assistance and Hein Goemans and Brenton Kenkel for useful conversations on this topic.

†Department of Political Science and Wallis Institute of Political Economy, University of Rochester. email: mark.fey@rochester.edu

1 Introduction

In order to understand the puzzle of why countries fight costly wars, researchers have focused on the role that uncertainty and the incentive to misrepresent private information plays in the occurrence of war (Fearon, 1995). In fact, recent work has shown that the presence of private information can make the risk of war unavoidable (Fey and Ramsay, 2011). Given these results, scholars have drawn the conclusion that reducing uncertainty, particularly by means of increasing information, reduces the chance of war. For example, Kydd (2010) writes "If uncertainty leads to cooperation failure, then information can lead to conflict resolution."

While this view that reducing uncertainty through information plays an important role in several areas of international conflict, such as third-party mediation, one of the most important is in understanding how wars end. This view is summarized by Filson and Werner (2002), who write "If private information and incentives to misrepresent that information often lead to war, then fighting the war can lead to peace by revealing the information that was previously concealed." Similarly, Werner and Yuen (2005) argue that "…consistent battle outcomes provide consistent information and encourage the belligerents' expectations about the future of war to converge. As a result, all parties likely agree about the probable consequences of renewing the conflict and are less likely to do so." Perhaps the clearest expression of this view is given by Reiter (2003) in his review article on the bargaining model of war:

> [A] central role of combat in the bargaining model of war is the reduction of uncertainty. As discussed, a common bargaining model explanation of the outbreak of war is uncertainty about the outcome of a hypothetical war. Combat can reduce uncertainty by providing information about the actual balance of power …The outcome of combat is observed by both sides and should cause their expectations to converge regarding the likely outcomes of future combat. This increases the likelihood of reaching an agreement that both sides prefer over continued fighting.

In a nutshell then, the argument about the role of combat in ending wars is the following. Initially, the two sides have uncertainty about the which side is likely to be victorious. Fighting provides information to both sides about their relative strengths. This reduction in uncertainty makes a peaceful agreement more likely than at the beginning of fighting. This formulation shows that there are two parts to this argument. The first claim is that information that is revealed about relative strength reduces uncertainty. The second claim is that reducing uncertainty makes peace more likely.

In this paper, we argue that the while both claims hold in the very long run, in the short and medium term both claims can be invalid. First, while it is true that eventually, after enough information is acquired, all uncertainty must be resolved, it is also easy to see that it is possible to receive a large but finite number of noisy but informative signals and have uncertainty increase as more signals are received. For example, if one side is quite certain it will prevail in fighting (and thus has low uncertainty) and then receives unexpected bad news, it will become more uncertain about its prospects of victory. [1]

Given this, our paper focuses on the second claim and gives a series of results that even if new information does reduce uncertainty, this does not necessarily make peace more likely. Again, in the very long run as uncertainty goes to zero, the chance of war must go to zero. But away from the limit, with a large but finite number of noisy but informative signals, reducing uncertainty may in fact increase the risk of war. In fact, we show that there are parameters for which, with expected probability close to one, receiving k signals generates a higher probability of war than receiving no signals at all.

We establish these results in the context of a standard ultimatum bargaining game with one-sided incomplete information. Specifically, we suppose that the side that makes the offer is uncertain about the likelihood of victory. We compare the offer that is made with no new information and the resulting risk of war with the offer that is made after receiving one or more noisy signals about the probability of winning and the resulting risk of war. While the obvious interpretation of these noisy signals is that they are the results of skirmishes between the two sides, we can also take a broader view of these signals. They could represent observations on the conduct of a conflict elsewhere, such as when various European countries sent military observers to the view the fighting in the American Civil War or when the air battles in the 1982 Lebanon War were viewed as providing information as to the relative capabilities of American and Soviet aircraft. Another interpretation of these signals could be pre-war espionage or information gained through war games and training exercises.

To establish our results that reducing uncertainty may make war more likely, we begin with case of one signal and show that for a large range of parameter values there is probability 1/2 that the signal will result in an offer with a higher risk of war than without the signal. We then look at the case of a large but finite number of signals. We show that this perverse result can still occur. In fact, we show that there exists parameter values such that we can have a probability close to one of receiving signals that make war more likely, even if the number of signals is large. Finally, we give some concrete lower bounds on the set of parameters and the

[1]Complete details formalizing this argument are presented in the Appendix.

probability of the perverse result.

In these results, for simplicity we suppose that the signals do not change the underlying probability of victory. But if we view these signals as representing actual battles, instead of mere skirmishes, it could be that the results of the battles not only provide information about the probability of victory, but can also change the value of the probability. Models of this sort include Smith (1998), Filson and Werner (2002), Smith and Stam (2004), and Slantchev (2003). To address these issues, in the final section we reconsider the effect of learning when battles can affect the prospects of victory and show that our results continue to hold.

How do these results occur, particularly given some contrary findings in the literature? For example, Reed (2003) shows in a similar model that reducing the variance of the belief of a country, holding the mean fixed, reduces the probability of war. Where this paper differs is that it gives a foundation for the change in variance by providing a model of noisy but informative signals. In particular, because the signals are informative, they not only reduce the variance of the belief but also change its mean. This in turn changes the demand that a country makes based on its updated posterior. This new offer is determined by the standard "risk-reward tradeoff," but this tradeoff is determined by the hazard rate of the posterior instead of its variance. Thus, as we show, the signal can simultaneously reduce variance and change the hazard rate in a way that makes war more likely.

In some ways, then, the analysis here harkens back to some old ideas in the conflict literature, namely that battle outcomes might not lead to peace because the winner will raise its demands in response. This argument appears in Wittman (1979), who says "If one side is more likely to win at war, its peaceful demands increase; but at the same time the other side's peaceful demands decrease. Thus we do not know whether [a peaceful settlement] is more or less likely." However Wittman does not present a fully rational model with incomplete information and noisy signals as is given here.

Perhaps the mostly closely related paper in the literature is Arena and Wolford (2012). The paper contains a model in which an uninformed country chooses a level of armaments and a level of intelligence gathering, which is followed by a noisy signal about the opponent's strength. The accuracy of this signal is determined by the level of intelligence gathering. The authors demonstrate that it is possible for a country to rationally choose a level of intelligence gathering that makes it possible to receive a signal leading to a higher risk of war than with no intelligence gathering. While this finding is similar to what is found in the current paper, there are also several differences in the paper. In the current paper we do not make the information available to the two sides a strategic choice, instead we view this information as being revealed on the battlefield. We also do not have a tradeoff between armaments and intelligence gathering. In this way, we show that reducing uncertainty can lead to a higher risk of war as a direct consequence of rational updating and thus it does not emerge only as a byproduct of strategic choices on arms and intelligence.

Our results also speak to the empirical literature that has attempted to test the relationship between battlefield outcomes and the end of wars. Ramsay (2008) uses data on battles fought in wars to test hypotheses drawn from existing models. He finds little for support for the hypothesis relating battle outcomes to information exchange and war termination. One other hand, Kaplow and Gartzke (2014) offer a new dataset that operationalizes military uncertainty, aiming to test the hypothesis that a reduction in uncertainty leads to a reduction in conflict. They find preliminary support for this hypothesis. In light of our theoretical results, these mixed empirical findings should not be surprising. We show that in some situations, battles can provide information to hasten peaceful settlements, but in other situations they can have the reverse effect. Thus an empirical test that looks to find an average effect of battle information may find nothing, or find inconsistent results.

2 The Basic Model

The baseline model for our analysis and comparison is the standard crisis bargaining ultimatum game (Fearon, 1995). In this game there is a resource of unit size under dispute by two countries. Country 1 makes a proposal of a division of the resource where it keeps a share $x \in [0, 1]$, leaving $1 - x$ for the other country. Country 2 can then accept or reject this offer. If accepted there is a peaceful settlement of the dispute and the payoffs are equal to the shares x and $1 - x$. If the proposal is rejected, war occurs. In the case of war, country 1 wins the resource with probability p and country 2 wins with the complementary probability, $1 - p$. To simplify notation, we sometimes write p_i for the probability that country i wins a conflict, where $p_1 = p$ and $p_2 = 1 - p$. Each side to the dispute pays a cost $c_i > 0$ if war occurs. Thus, the expected payoff to war to country i is $p_i - c_i$.

It is well know that because war is costly, there exists a set of peaceful agreements that both countries prefer to fighting. In the complete information environment the probability of the dispute turning to war is zero. In the unique subgame perfect equilibrium to this crisis bargaining game country 1 demands the largest share that country 2 will accept, which is $x = p + c_2$. This is accepted by country 2, which means that war does not occur in equilibrium with complete information.

With uncertainty, on the other hand, it is possible for war to occur in equilibrium. We consider the case of one-sided incomplete information in which there is uncertainty about the probability of winning. As in the model of Reed (2003), in this model country 2 knows the true probability of winning the war, while country 1 does not. Although most existing models have uncertainty about the costs of war, we are interested in how fighting reveals information and thus it is more natural to have the information about the likelihood of victory instead of the costs of war.

In our model, types are continuous. Specifically, the probability p that 1 wins the war is drawn from a continuous, strictly increasing cumulative distribution function $F(p)$ on $[0, 1]$ with density $f(p)$. The only condition that we put on the shape of F is that is has an increasing hazard rate. That is, the hazard rate $h(p) = f(p)/(1 - F(p))$ is strictly increasing in p.[2] While there is uncertainty about relative power (as measured by p), for simplicity we assume there is complete information about the costs of war c_i, as in Reed (2003). In addition, to avoid trivial cases, we assume that $c_1 + c_2 < 1$.

To find the equilibrium of this basic model, we begin by noting that if country 1 makes the offer $(x, 1 - x)$, country 2 will reject the offer x when $p < x - c_2$ and accept the offer when $p > x - c_2$. From this we see that the offer $(c_2, 1 - c_2)$ will be accepted with probability one, so clearly country 1 will never make an offer with $x < c_2$. Thus, assuming that $x \in [c_2, 1]$, it follows that the utility to country 1 of making an offer x, which we denote $U(x)$, is given by

$$U(x) = \int_0^{x-c_2} (p - c_1) f(p)\, dp + \int_{x-c_2}^{1} x f(p)\, dp.$$

The first derivative is obtained via Leibniz's rule and is given by

$$U'(x) = (-c_1 - c_2) f(x - c_2) + 1 - F(x - c_2).$$

This gives a first order condition which identifies an interior maximizer.[3] Specifically, an interior maximum x^* is a solution to

$$h(x^* - c_2) = \frac{1}{c_1 + c_2}.$$

[2]Many standard distributions have this property, such the uniform distribution, the normal distribution, etc.

[3]To confirm the solution to the first order condition is a maximum, we check the second order condition. We have $U''(x) = (-c_1 - c_2) f'(x - c_2) - f(x - c_2) = -((1 - F(x - c_2))/f(x - c_2)) f'(x - c_2) - f(x - c_2)$. Some algebra reveals that this yields $U''(x) = -(1 - F(x - c_2))^2 h'(x - c_2)/f(x - c_2)$ which is always negative as $h(x)$ is strictly increasing.

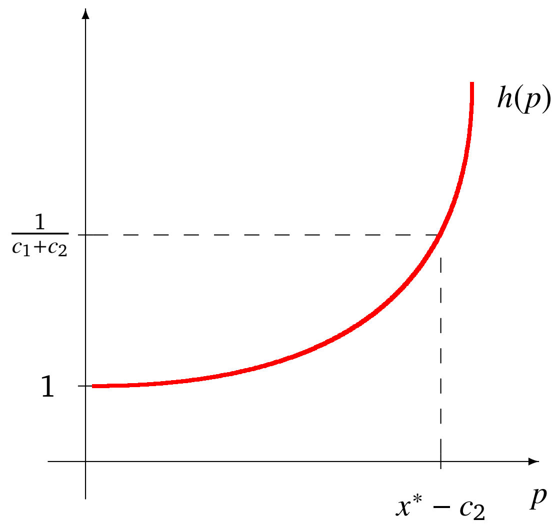

Figure 1: Equilibrium in the Basic Model

Note that the left hand side is strictly increasing on $[c_2, 1]$ and the right hand side is a constant. Therefore, if x^* is an interior maximum, it is unique. In addition, this implies that if $h(0) = f(0) \geq 1/(c_1 + c_2)$, then the optimal offer is $x^* = c_2$ and if $h(1 - c_2) \leq 1/(c_1 + c_2)$, then the optimal offer is $x^* = 1$. Of course, if neither of these boundary conditions are satisfied, then the optimal offer is interior. This establishes the following proposition:

Proposition 1. *Suppose the hazard rate $h(p)$ is strictly increasing. Then the ultimatum game has a unique equilibrium offer which is given by $x^* = c_2$ if $f(0) \geq 1/(c_1 + c_2)$, $x^* = 1$ if $h(1 - c_2) \leq 1/(c_1 + c_2)$, and by the unique solution to*

$$h(x^* - c_2) = \frac{1}{c_1 + c_2},$$

otherwise.

As an example of finding the unique equilibrium offer, consider the case of F uniform on $[0, 1]$. It is easy to see that the hazard rate $h(p) = 1/(1 - p)$ and therefore the solution will be interior, as $c_1 + c_2 < 1$. Therefore x_U^* satisfies

$$h(x_U^* - c_2) = \frac{1}{1 - (x_U^* - c_2)} = \frac{1}{c_1 + c_2},$$

and thus $x_U^* = 1 - c_1$. This is illustrated in Figure 1. Recall that country 2 rejects the offer x_U^* if $p < x_U^* - c_2$, so the probability of war in this example is $F(x_U^* - c_2) = 1 - c_1 - c_2$.

3 Uncertainty and the Risk of War

3.1 Adding Information to the Basic Model

We now turn to the question of what happens if country 1 receives new information about its prospects of victory. Specifically, we will augment the basic model by a process that generates noisy but informative signals about the true value of p. The source of these noisy signals could be skirmishes, observations of other conflicts, or even espionage. In order to avoid complications and to focus our analysis on what country 1 learns, we suppose that the observed signals do not change the underlying parameters of the game.

For simplicity, throughout this paper we suppose that country 1's prior is given by a uniform distribution on $[0, 1]$. The sequence of play in this section is as follows. First, country 1 receives one or more noisy signals about the true value of p. Using this information, country 1 updates its belief about p and proceeds to makes a take-it-or-leave-it offer to country 2, which country 2 either accepts or rejects, with payoffs as given for the basic model.

Formally, country 1 receives a binary signal $s \in \{g, b\}$. We interpret g as "good news" and b as "bad news". The signal is noisy but informative, meaning that good news is more likely when the true value of p is higher. In particular, we suppose that the probability that the signal is good news is equal to p, the probability of victory in the conflict. We denote this signal function as $\phi(g \mid p) = p$. Of course, this means that probability the signal is bad news is equal to $1 - p$. Country 1 then uses Bayes' Rule to form its posterior belief about the value of p and uses this posterior to choose its demand.

3.2 Results for the Short Term

We begin with the case of country 1 observing a single noisy signal.

So what is country 1's updated belief about the value of the p after observing the noisy signal? To calculate this, we use a continuous version of Bayes' Rule. This gives the country 1's updated belief $f(p \mid s)$, given the signal s as follows:

$$f(p \mid s) = \frac{\phi(s \mid p) f(p)}{\int_0^1 \phi(s \mid p') f(p') \, dp'},$$

where $f(p)$ is the prior of country 1.[4] We now use this expression to calculate the posterior

[4] Notice that the denominator of this expression gives the unconditional expected probability of receiving the noisy signal s.

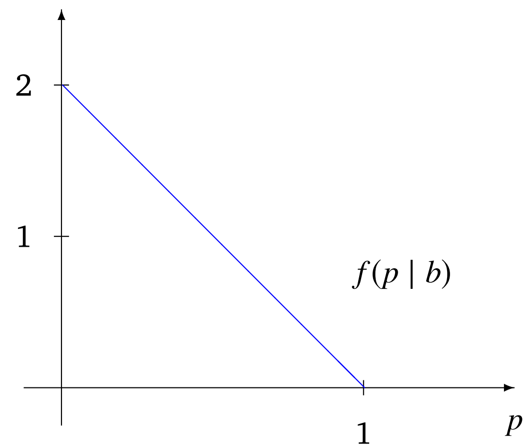

Figure 2: Belief after bad news

belief of country 1 after receiving the signal b as follows:

$$f(p \mid b) = \frac{(1-p)\cdot(1)}{\int_0^1 (1-p')\cdot(1)\,dp_1'} = \frac{1-p}{1/2} = 2(1-p).$$

This posterior belief is illustrated in Figure 2. As expected, after bad news country 1's posterior belief shifts to the left, reflecting the fact that country 1 thinks it is more likely that p is low than high. Note also that the variance of this posterior is equal to 1/18 which is smaller than the variance of the uniform prior, which is 1/12.

Country 1's uses its updated belief to decide on its ultimatum offer. It is easy to check that the hazard rate of the posterior is given by

$$h(p \mid b) = \frac{2(1-p)}{1-\int_0^p 2(1-p')\,dp'} = \frac{2(1-p)}{(1-p)^2} = \frac{2}{1-p}.$$

Clearly, this is strictly increasing and we can thus apply Proposition 1 to see that the unique equilibrium offer is given by $x_B^* = c_2$ if $c_1 + c_2 \geq 1/2$ and is given by the solution of

$$2/(1-(x-c_2)) = 1/(c_1+c_2)$$

if $c_1 + c_2 < 1/2$. The latter is easily solved to give $x_B^* = 1 - 2c_1 - c_2$. Comparing this to the optimal offer in the case of uniform prior, we see that after bad news country 1 makes a more generous offer to player 2, reflecting the fact that country 1 thinks it is more likely that country 2 has a high reservation value. The probability of war after bad news is given by

$$F(x_B^* \mid B) = (x_B^* - c_2)(2-(x_B^*-c_2)) = (1-2c_1-2c_2)(1+2c_1+c_2) = 1-4(c_1+c_2)^2.$$

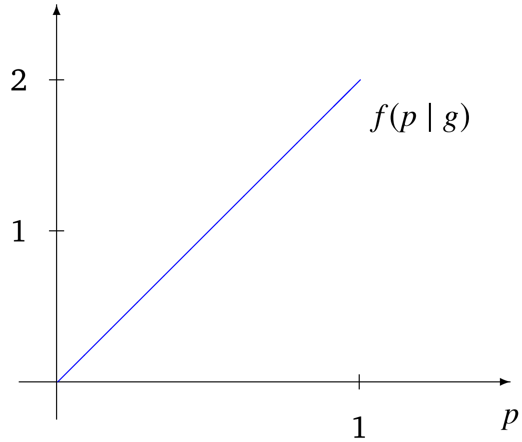

Figure 3: Belief after good news

if $c_1 + c_2 < 1/2$. Comparing this probability of war to that given for the case of a uniform prior in the previous section, we see that for this range of costs, war is more likely after receiving bad news when

$$1 - 4(c_1 + c_2)^2 > 1 - (c_1 + c_2),$$

which simplifies to $c_1 + c_2 < 1/4$. In other words, if the total social cost of war is less than one fourth of the value of the prize, receiving bad news leads country 1 to make an offer with a higher probability of war than if no information was received.

Does this result depend critically on the fact that country 1 receives bad news? To find out, we consider the other case in which country 1 instead receives good news. We proceed just as we did for the case of bad news. After receiving a good signal (which happens with probability p), the posterior belief of country 1 is given by:

$$f(p \mid g) = \frac{p \cdot (1)}{\int_0^1 p' \cdot (1)\, dp_1'} = \frac{p}{1/2} = 2p.$$

This posterior belief is illustrated in Figure 3. Here, we get a similar result as the case of bad news: country 1's posterior belief shifts to the right, reflecting the fact that country 1 thinks it is more likely that p is high than low. The variance of this posterior is again equal to 1/18 which is smaller than the variance of the uniform prior, which is 1/12.

In this case, the hazard rate of the posterior is given by

$$h(p \mid g) = \frac{2p}{1 - \int_0^p 2p'\, dp'} = \frac{2p}{1 - p^2}.$$

It is straightforward to verify that this is strictly increasing and we can thus apply Proposition 1 to find the equilibrium offer. We see that the unique equilibrium offer is given by the solution of

$$2(x - c_2)/(1 - (x - c_2)^2) = 1/(c_1 + c_2)$$

if $c_1 > c_2^2/2(1 - c_2)$ and is equal to 1 otherwise. Solving and combining these two cases, we see that $x_G^* = \min(\sqrt{1 + (c_1 + c_2)^2} - c_1, 1)$. Comparing this to the optimal offer in the case of uniform prior, we see that after good news country 1 makes a more demanding offer to player 2, reflecting the fact that country 1 thinks it is more likely that country 2 has a low reservation value. The probability of war after good news is given by

$$F(x_G^* \mid G) = (x_G^* - c_2)^2.$$

Again, comparing this probability of war to that given for the case of a uniform prior, we see that, assuming $x_G^* < 1$, war is more likely after receiving good news when

$$(\sqrt{1 + (c_1 + c_2)^2} - c_1 - c_2)^2 > 1 - (c_1 + c_2),$$

which simplifies to $c_1 + c_2 > 3/4$. In other words, if the total social cost of war is greater than three fourths of the value of the prize, receiving good news leads country 1 to make an offer with a higher probability of war than if no information was received.[5]

Finally, we can ask what is the ex ante probability of receiving each of these signals? With a uniform prior, the probability of receiving a good signal is simply

$$\int_0^1 (p)(1)\, dp = 1/2.$$

Likewise the ex ante probability of receiving a bad signal is also 1/2. Thus our results can be summarized in the following proposition:

Proposition 2. *If either $c_1 + c_2 < 1/4$ or $c_1 + c_2 > 3/4$ and $c_1 \geq 1/4$, then with probability 1/2 country 1 will receive a signal that makes more war likely than with no signal.*

In other words, there is a broad range of parameters for which reducing uncertainty has a significant chance of making peaceful settlements less likely.

[5]If $x_G^* = 1$, then this conclusion holds if $c_1 \geq c_2 - c_2^2$. Note that a sufficient condition for this is $c_1 \geq 1/4$.

3.3 Results for the Medium Term

Now that we have established how learning can affect the prospect of war after receiving one signal, it is natural to consider what effect learning multiple signals has. Perhaps this perverse result will lessen as more information is accumulated. Here we examine the case of $k > 1$ noisy signals and consider how likely it is for these signals to lead to an increased chance of war.

As before, we suppose that country 1's prior is given by a uniform distribution on $[0, 1]$. The sequence of play is the same as before, except country 1 first receives k noisy signals about the true value of p before making its offer. We assume that each signal is binary, as before, with each signal independently equal to g with probability p. To calculate the posterior belief of country 1, suppose it receives w good signals and l bad signals, with $w + l = k$. The probability is receives these signals is equal to $p^w(1-p)^l$ and therefore the posterior belief of country 1 after receiving these signals is:

$$f(p \mid (w,l)) = \frac{p^w(1-p)^l \cdot (1)}{\int_0^1 (p')^w(1-p')^l \cdot (1)\, dp'_1} = \frac{\Gamma[w+l+2]}{\Gamma[w+1]\Gamma[l+1]} p^w(1-p)^l,$$

where $\Gamma[n] = (n-1)!$. This shows that after observing w good signals and l bad signals, country 1's posterior belief is given by a beta distribution with parameters $w+1$ and $l+1$. As an example, this posterior belief is illustrated in Figure 4 for the case of 4 good signals and 1 bad signals. Here, country 1's posterior belief is unimodal and concentrated around its mean of 4/5, reflecting the effect of the new information on country 1's thinking about its likelihood of victory. Note also that as country 1's posterior is given by a beta distribution with parameters $w+1$ and $l+1$, it is true that the variance of this posterior is always strictly less than the variance of the uniform prior.

We now proceed to give a general argument. Fix a number of signals $k > 1$, with w good signals and l bad signals. As long as $l > 0$, the density of the beta distribution goes to zero as p goes to 1. Moreover, we know that the hazard rate of the beta distribution is always strictly increasing and it follows from L'Hôpital's rule that $h(p \mid (w,l))$ goes to infinity as p goes to 1. To continue, we employ the following construction. Let $\hat{p}$ be the smallest value of p such that $f(p \mid (w,l)) < 1$ for all $p > \hat{p}$. Let $\hat{C} = 1/h(\hat{p} \mid (w,l))$, the total war cost that solves the first order condition. Note that $\hat{C} < 1$ as $h(\hat{p} \mid (w,l)) > 1$. For every pair of costs (c_1, c_2) such that $c_1 + c_2 < \hat{C}$, the solution of $h(p) = 1/(c_1+c_2)$ will have $p > \hat{p}$ and thus $f(p \mid (w,l) < 1$. Moreover, for every cost $c_1 \in (0, \hat{C})$, let $\hat{c_2}(c_1)$ be the solution to $h(1-c_2) = 1/(c_1+c_2)$. Clearly

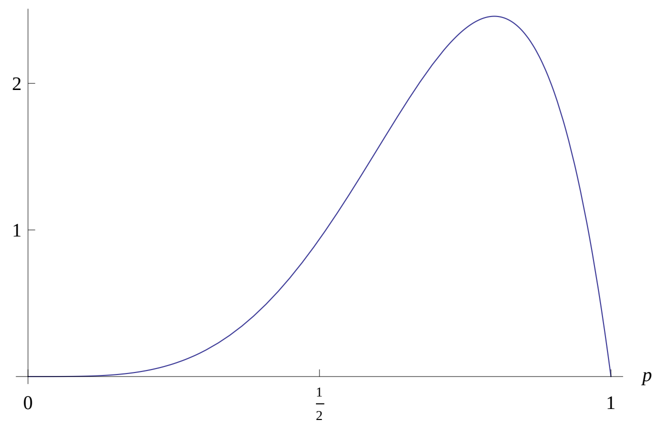

Figure 4: Belief after 4 good signals and 1 bad signal

$\hat{c_2}(c_1) > 0$ for every $c_1 \in (0, \hat{C})$. So by construction, every pair (c_1, c_2) that satisfies $c_1 \in (0, \hat{C})$ and $c_2 \in (0, \hat{c_2}(c_1))$ will have the following property: the solution p^* to $h(p) = 1/(c_1 + c_2)$ will have $p^* \leq 1 - c_2$ and $f(p^*) < 1$. This implies that setting $x_k^* = p^* + c_2$

$$f(x_k^* - c_2 \mid (w, l)) = \frac{1 - F(x_k^* - c_2 \mid (w, l))}{c_1 + c_2} = \frac{1 - F(x_k^* - c_2 \mid (w, l))}{1 - F(x_U^* - c_2)}$$

where the last equality comes from the fact that $F(x_U^* - c_2) = 1 - c_1 - c_2$ for the uniform distribution. But $f(p^*) < 1$, which implies that $1 - F(x_k^* - c_2 \mid (w, l)) < 1 - F(x_U^* - c_2)$ so that $F(x_k^* - c_2 \mid (w, l)) > F(x_U^* - c_2)$. In words, this means that the probability of war is higher after receiving these k signals than with no signals.

Note that this argument works for an arbitrary value of k, no matter how large. This establishes the following proposition:

Proposition 3. *For every set of w good signals and l bad signals with $l > 0$, there exists a set of cost pairs (c_1, c_2) of positive measure such that, with positive probability, the probability of war after receiving this set of signals is higher than receiving no signals.*

Of course, as this proposition stands, it could be that the set of costs that give this perverse result is small and the probability of receiving a given set of signals is also small. We explore these issues in the next two propositions.

First, we can strengthen Proposition 3 by carefully examining the set of cost pairs identified in the proof for each specific set of signals. We see that for every set of signals with $l > 0$

bad signals, the set of cost pairs is described as every pair (c_1, c_2) that satisfies $c_1 \in (0, \hat{C})$ and $c_2 \in (0, \hat{c_1})$. If we takes the minimum of the $\hat{C}$ values and the minimum of the $\hat{c_2}(c_1)$ functions across the k possible signals with $l > 0$, we still end up with a set of cost pairs of positive measure. But now notice that the expected probability of receiving one of these k possible signals is equal to one minus the expected probability of receiving k good signals and 0 bad signals, which is given by

$$1 - \int_0^1 p^k(1)\,dp = 1 - \frac{1}{k+1} = \frac{k}{k+1}.$$

We thus have established the following proposition:

Proposition 4. *Fix a number of signals k. There exists a set of cost pairs (c_1, c_2) of positive measure such that the probability of war after receiving this set of signals is higher than receiving no signals with probability greater than or equal to $k/(k+1)$.*

In other words, even if the number of signals is large, there exists parameters of the model such that fighting becomes more likely after learning with probability close to one.

In this proposition, we still have the possibility that the set of cost pairs could be very small. To give a more concrete analysis, we can develop some lower bounds on the set of costs and the probability of learning leading to a higher chance of war. To begin, fix the number of signals k and consider the situation in which there are 0 good signals and k bad signals. After observing these signals the posterior belief of country 1 takes on a particularly simple form:

$$f(p \mid (0,k)) = \frac{\Gamma[k+2]}{\Gamma[1]\Gamma[k+1]} p^0 (1-p)^k = (k+1)(1-p)^k.$$

From this it follows that the hazard function is given by

$$h(p \mid (0,k)) = \frac{(k+1)(1-p)^k}{1 - [1 - (1-p)^{k+1}]} = \frac{k+1}{1-p}.$$

Based on the hazard function, we can see that the optimal offer of country 1 will be interior as long as $c_1 + c_2 < 1/(k+1)$. In particular, the optimal offer of country 1 is always strictly less than 1.

Following the same notation as in the argument for Proposition 3, the solution of $f(p \mid (0,k)) = 1$ is denoted

$$\hat{p} = 1 - \left(\frac{1}{k+1}\right)^{\frac{1}{k}}.$$

Likewise, let

$$\hat{C} = \frac{1}{h(\hat{p} \mid (0,k))} = \left(\frac{1}{k+1}\right)^{1+\frac{1}{k}}.$$

As the boundary constraint of $x^* = 1$ never binds, it follows that every cost pair (c_1, c_2) such that $c_1 + c_2 < \hat{C}$ will give a solution p^* to $h(p) = 1/(c_1 + c_2)$ with $f(p^*) < 1$. Therefore by the argument for Proposition 3, every such cost pair will have a higher ex ante risk of war after receiving k bad signals than with no signals. Finally, we can calculate the expected probability of receiving k bad signals. It is simply $\int_0^1 (1-p)^k(1)\,dp = 1/(k+1)$. We can thus give the some lower bounds in the following proposition:

Proposition 5. *Fix a number of signals k. If the set of cost pairs satisfies $c_1+c_2 < (1/(k+1))^{1+(1/k)}$, then the expected probability of receiving signals such that the probability of war after receiving this set of signals is higher than receiving no signals is at least* $1/(k+1)$.

As an example of this proposition, suppose country 1 receives 5 signals. Then if $c_1 + c_2 < 6^{-6/5} \approx .12$, the probability of receiving signals that lead to a higher risk of war than with no signals is at least $1/6$.

3.4 Results for the Very Long Term

We have shown that for each fixed number of signals k, there exists a set of costs such that, with high probability, war is more likely after receiving signals than before. Here we show that as the number of signals goes to infinity, this result no longer holds. That is, we show that for every parameter value, if the number of signals is sufficiently large, then the probability of war goes to zero.

Specifically, here we fix the costs of war c_1 and c_2. As long as country 1's prior has full support, its posterior will be asymptotically normal (that is, after a sufficiently large number of signals) with mean equal to the true value of p and variance going to zero in the limit. This is a consequence of the Bernstein-von Miller Theorem in Bayesian analysis. In particular, we know the following about the posterior $f(p \mid k)$, given the true value p, which we denote p^*. In the limit as $k \to \infty$, $F(p^* \mid k) = 1/2$ almost surely and for all $p < p^*$, $F(p \mid k) \to 0$ almost surely. This implies that for sufficiently large k, the hazard rate $h(p) \in (f(p \mid k), 2f(p \mid k))$ almost surely. In addition, the hazard rate $h(p \mid k)$ is almost surely strictly increasing for sufficiently large k. This implies that there exists a number K such that for all $k > K$, almost surely there exists $p' < p^*$ that solves $h(p'(k) \mid k) = 1/(c_1 + c_2)$. But then $F(p') \to 0$ as $k \to \infty$. This establishes the following proposition:

Proposition 6. *Fix a pair of costs (c_1, c_2). The expected probability of war after receiving k signals goes to zero as k goes to infinity.*

Thus, while there are cost parameters such that after a fixed number of signals, it is very likely that war is more likely than without signals, eventually this effect disappears. When a sufficiently large number of additional signals are received, the probability of war goes to zero.

4 Learning From Battles

Up to this point, we have assumed that the noisy signals provide information about the relative power of the two sides, but don't affect it. However, particularly if we view the signals as the observed results of battles between the two sides, it may be that the signals alter the true value of p. That is, it is not unreasonable to thing that the outcome of battles changes the relative likelihood of winning the war.

Do our results that informative signals can lead to a greater risk of war extend to this case? In this section we show that they do. We maintain the assumption that the signals are either good or bad and occur with probability p and $1 - p$, respectively. In addition, though, we now assume that after a bad signal occurs, the true value of p is shifted down to a value of βp, for some $\beta \in (0, 1)$. Likewise, after a good signal occurs, the value of $1 - p$ is shifted to a value of $\beta(1 - p)$. In other words, the losing side of the signal has their probability of victory reduced by a factor β.

Formally, we suppose that country 1 begins with a uniform prior, as before, and observes a signal. This signal is used to form a posterior $f(p \mid s)$ in the usual way. Then, understanding that the signal has shifted the true value of p, country 1 shifts its posterior to reflect this change, leading to a new posterior, which we denote $\tilde{f}(p \mid s)$. For example, if a bad signal is received and thus the value of p is shifted to βp, the posterior $f(p \mid s)$ is replaced by

$$\tilde{f}(p \mid s) = \frac{f(\frac{p}{\beta} \mid s)}{\beta},$$

which is distributed on $[0, \beta]$.

Now, continuing with the case of a bad signal, the posterior $f(p \mid b) = 2(1 - p)$ from the

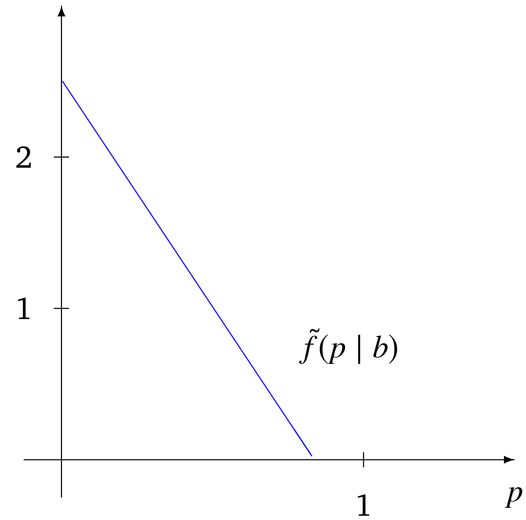

Figure 5: Shifted belief after bad news

previous section. Therefore the shifted posterior is given by

$$\tilde{f}(p \mid b) = \frac{2(1 - \frac{p}{\beta})}{\beta}.$$

This shifted posterior belief is illustrated in Figure 5. Here we see that the posterior has the same shape as before, only shifted to left by a factor β.

Country 1's uses its shifted belief to decide on its ultimatum offer. The hazard rate of the shifted posterior is given by

$$\tilde{h}(p \mid b) = \frac{\frac{2}{\beta}(1 - \frac{p}{\beta})}{1 - \int_0^p \frac{2}{\beta}(1 - \frac{p'}{\beta})\, dp'} = \frac{2(\beta - p)}{(\beta - p)^2} = \frac{2}{\beta - p}.$$

Clearly, this is strictly increasing and we can thus apply Proposition 1 to see that the unique equilibrium offer is given by $x_B^* = c_2$ if $c_1 + c_2 \geq \beta/2$ and is given by the solution of

$$2/(\beta - (x - c_2)) = 1/(c_1 + c_2)$$

if $c_1 + c_2 < \beta/2$. The latter is easily solved to give $\tilde{x}_B^* = \beta - 2c_1 - c_2$. We see that in this case the optimal offer is shifted to the left by an amount $1 - \beta$ as compared to the offer in the unshifted case. The probability of war after bad news in this case is given by

$$F(x_B^* \mid B) = \frac{\tilde{x}_B^* - c_2}{\beta}\left(2 - \frac{\tilde{x}_B^* - c_2}{\beta}\right) = \frac{(\beta - 2c_1 - 2c_2)(\beta + 2c_1 + c_2)}{\beta^2} = 1 - \frac{4(c_1 + c_2)^2}{\beta^2}.$$

if $c_1 + c_2 < \beta/2$. Comparing this probability of war to that given for the case of a uniform prior, we see that for this range of costs, war is more likely after receiving bad news when

$$1 - \frac{4(c_1 + c_2)^2}{\beta^2} > 1 - (c_1 + c_2),$$

which simplifies to $c_1 + c_2 < \beta/4$. Thus, comparing this result to the unshifted case, we see that while the range of costs is smaller, it still possible that receiving bad news leads country 1 to make an offer with a higher probability of war than if no information was received.

5 Multiple Rounds of Bargaining

The structure of the bargaining game we have analyzed so far is quite simple. Indeed, country 2 plays no part in the bargaining until the very end. But it may be more realistic to imagine bargaining as an ongoing process in which offers and battles are interspersed. In particular, one could imagine that a very weak type of country 2 would be eager to settle early in the conflict in order to prevent her weakness from being revealed on the battlefield.

In order to improve our understanding of this process, in this section we extend our earlier results to a model in which there are two rounds of bargaining. That is, we consider a model in which country 1 makes an initial offer which country 2 can accept and end the game. However, if country 2 rejects the initial offer, country 1 receives a noisy signal of country 2's strength, as before, and then makes a final offer to country 2. Other than this additional offer, all other details of our earlier model will still hold here. We are interested in whether our results from our previous analysis carry over to this setting.

5.1 Baseline With No Information

In order to create a proper comparison, we first examine a baseline model in which there are two offers, but no information gathering by country 1. We can then compare the results to the model with learning by country 1. Specifically, then, we consider a model in which country 1 makes an initial offer $(x_1, 1 - x_1)$. If country 2 accepts, this division is implemented, and if country 2 rejects, then country 1 makes a second offer $(x_2, 1 - x_2)$. Again, acceptance results in payoffs $(x_2, 1 - x_2)$ to the players while rejection results in war. We do not discount payoffs in the second round. As before, country 1 is uncertain about p, which is known to country 2. We suppose that the prior distribution of p is uniform on $[0, 1]$.

This game is very similar to one analyzed in Fey, Meirowitz and Ramsay (2013) and the result essentially the same. Recall that $x_U^* = 1 - c_1$ is the equilibrium offer made in the baseline ultimatum game. As we show in the next proposition, in the game with two offers and no learning the unique equilibrium outcome is that country 1 makes an offer in the first period that all types of country 2 reject and country 2 makes the offer $(x_U^*, 1 - x_U^*)$ in the second period. Thus, the outcome of the game with two offers is identical to the game with only a single offer.

Proposition 7. *The game with two offers and no signals has a unique equilibrium outcome in which all types of country 2 reject the first offer, the second offer is $x_2 = x_U^*$ and country 2 rejects this offer if $p < x_U^* - c_2$ and accepts this offer if $p > x_U^* - c_2$.*

Proof. We assume that after an initial offer $(x_1, 1 - x_1)$, country 2 plays a cutpoint strategy such that all types with $p < \bar{p}(x_1)$ reject the first offer and all types with $p > \bar{p}(x_1)$ accept the first offer. This means that country 1's belief before making the second offer is uniform on the interval $[0, \bar{p}]$. Therefore, if $\bar{p} \geq c_1 + c_2$, then country 2 will make an offer $x_2 = \bar{p} - c_1$ and all types $p \in [\bar{p} - c_1 - c_2, \bar{p}]$ will accept and if $\bar{p} < c_1 + c_2$, then country 2 will make an offer $x_2 = c_2$ and all types $p \in [0, \bar{p}]$ will accept. We note two things from this. First, in either case, $x_2 \leq x_U^* = 1 - c_1$. Second, in either case, some types of country 2 are rejecting the first offer and accepting the second offer, so it must be that $1 - x_2 \geq 1 - x_1$.

Now suppose that $\bar{p}(x_1) < 1$. This means that some types of country 2 are accepting the first offer instead of rejecting the first offer and accepting the second offer. This requires that $1 - x_1 \geq 1 - x_2$. We conclude that either $\bar{p}(x_1) = 1$ or $x_2 = x_1$. It is easy to check that $\bar{p}(x_1) < 1$ if and only if $x_2 < x_U^*$.

From these facts, we conclude that $x_1 < x_U^*$ implies $x_2 = x_1$ and $x_1 \geq x_U^*$ implies $\bar{p}(x_1) = 1$ and $x_2 = x_U^*$. In other words, if country 1 offers $x_1 < x_U^*$, then its payoff is the payoff from offering x_1 in the one-shot ultimatum game. On the other hand, if country 1 offers $x_1 > x_U^*$, then its payoff is the payoff from offering x_U^* in the one-shot ultimatum game. As the optimal offer in the one-shot ultimatum game is x_U^*, then in this game country 1 must make an offer $x_1 > x_U^*$ in equilibrium, which is rejected by all types of country 2, followed by $x_2 = x_U^*$. $\square$

Somewhat surprisingly, allowing country 1 to make a second offer after rejection does not change the outcome relative to a model with no initial offer.

5.2 Learning In Between Offers

Now that we have established the outcome without the possibility of learning, we now consider what happens when country 1 can observe a noisy signal in the event its first offer is rejected.

The structure of this signal is the same as before: country 1 observes a good signal with probability p and a bad signal with probability $1 - p$. After observing this signal, country 2 makes an offer $(x_2, 1 - x_2)$. Note that a key feature of this model is that now country 2 has the opportunity to accept the initial offer and end the game, forestalling the acquisition of additional information by country 1.

As before, we suppose that country 1 has a uniform prior on p and also, for simplicity, we assume that $c_1 = c_2 = c$. In addition, we assume that c is small enough that all offers are interior. Finally, as in the previous part, we suppose that after an initial offer $(x_1, 1 - x_1)$, country 2 plays a cutpoint strategy such that all types with $p < \bar{p}(x_1)$ reject the first offer and all types with $p > \bar{p}(x_1)$ accept the first offer.

Given the cutpoint $\bar{p}$, we can proceed as before and calculate the optimal offers in the second period following a good signal or a bad signal. Assuming the solution is interior, this gives an offer $x_G^*(\bar{p}) = \sqrt{4c^2 + \bar{p}^2} - c$ after a good signal and an offer $x_B^*(\bar{p}) = 1 - \sqrt{4c^2 + (1 - \bar{p})^2} - c$ after a bad signal.

Now given an initial offer $(x_1, 1 - x_1)$, the type $\bar{p}$ must be indifferent between accepting $1 - x_1$ and rejecting the offer and waiting for the signal to be revealed and the resulting offer to be made. We will suppose that the type $\bar{p}$ of country 2 will accept the resulting offer and later check that this is the case in equilibrium. This means that we have the following indifference condition:

$$\begin{aligned} 1 - x_1 &= \bar{p}(1 - x_G^*(\bar{p})) + (1 - \bar{p})(1 - x_B^*(\bar{p})) \\ x_1 &= \bar{p}x_G^*(\bar{p}) + (1 - \bar{p})x_B^*(\bar{p}) \end{aligned}$$

Define $p_G^* = x_G^* - c$ and $p_B^* = x_B^* - c$. Then we can further see that the types of country 2 are expected to behave in the following way: types $p < p_G^*$ will reject the initial offer and both possible offers in the second round, types $p \in [p_B^*, p_G^*]$ will reject the initial offer and accept x_B^* but reject x_G^*, types $p \in [p_G^*, \bar{p}]$ will reject the initial offer and accept both possible second round offers, and finally types $p > \bar{p}$ will accept the initial offer. This allows us to write the expected utility of country 1 for some offer x_1 as

$$\begin{aligned} U_1(x_1) = \int_0^{p_B^*} (p - c)\, dp + \int_{p_B^*}^{p_G^*} [p(p - c) + (1 - p)(x_B^*)]\, dp \\ + \int_{p_G^*}^{\bar{p}} [p(x_G^*) + (1 - p)(x_B^*)]\, dp + \int_{\bar{p}}^{1} x_1\, dp. \end{aligned}$$

Thus, country 1 faces the problem of maximizing $U_1(x_1)$ subject to the indifference condition above. This can be simplified by instead choosing a $\bar{p}$ that generates an x_1 that maximizes $U_1(x_1)$.

The result of this maximization is the following proposition:

Proposition 8. *The game with two offers and one noisy signal has a unique equilibrium outcome in which all types of country 2 reject the first offer and play in the second period is as in the game with one noisy signal and one offer.*

To establish this proposition, we look at $U_1(x_1(\bar{p}))$, where the value of $x_1(\bar{p})$ is given by the above indifference condition. Differentiating this with respect to $\bar{p}$ and doing algebra, we obtain

$$\frac{dU_1(x_1(\bar{p}))}{d\bar{p}} = \frac{1-\bar{p}}{\gamma\beta}\left[\gamma((1-\bar{p})^2+\beta^2)-\beta\gamma+\bar{p}^2\beta+\gamma^2\beta\right],$$

where $\gamma = \sqrt{4c^2+\bar{p}^2} > \bar{p}$ and $\beta = \sqrt{4c^2+(1-\bar{p})^2} > 1-\bar{p}$. In particular we see that $\gamma+\beta > 1$. Now the expression in brackets can be written as

$$\gamma((1-\bar{p})^2+\beta^2)-\beta\gamma+\bar{p}^2\beta+\gamma^2\beta = \gamma\beta(\beta+\gamma-1)+\gamma(1-\bar{p}^2)+\beta\bar{p}^2.$$

As $\gamma+\beta > 1$, each of these terms is positive and so $U_1(x_1(\bar{p}))$ is always positive. This implies that the optimal initial offer for country 1 leads to $\bar{p} = 1$. But this implies that country 1's belief before receiving the noisy signal is uniform on $[0,1]$, which is identical to the model analyzed in the earlier section. Thus, behavior of both sides after the noisy signal is the same as in the earlier model with one noisy signal and just one offer. This establishes the proposition.

Again, somewhat surprisingly, even though country 2 can accept the initial offer and prevent the other side from observing the signal, it never does so in equilibrium. This occurs because country 1 understands that only the weakest types of country 2 would be willing to do so, and so it has an incentive to make more and more demanding offers.

As well, this proposition shows that, at least in this simple model, multiple rounds of bargaining have no effect on the likelihood that information will lead to more, rather than less, fighting. In particular, for the same ranges of cost presented in the earlier section, the chance of war increases with ex ante probability one half.

6 Conclusion

In this paper we have considered the effect of learning about the likelihood of victory on the risk of war. Although the conventional wisdom holds that acquiring information reduces the risk of war, we show that this is not necessarily true. Even with multiple signals, there is a significant chance that the risk of war increases relative to the baseline in which there is no learning.

Our results suggest that there is no straightforward relationship between uncertainty and the risk of war, at least in the short and medium term. This casts doubt on empirical approaches on war that simply include a measure of uncertainty as an explanatory variable. But one commonality in our results is that low costs correspond to a higher chance that information leads to more fighting. Future work can perhaps further explore this connection and thus offer a more nuanced method of empirically understanding how uncertainty and war relate.

Finally, the framework we have analyzed could be generalized in future work. Initial work suggests that our results will continue to hold for a more general class of prior distributions and a more general signal function. It may also be possible to consider more general models with multiple rounds of bargaining, although this must also be left to future work.

Appendix: Information and Uncertainty

The Basic Model With Two Types

Here we specify uncertainty with two types. The probability p that country 1 will win the war is known to country 2 but not to country 1. That is, p is private information to country 2. Let $p = p_W$ with probability π and $p = p_S$ with probability $1 - \pi$, where $0 < p_S < p_W < 1 - c_2$. The costs of war to both sides is common knowledge.

The unique equilibrium to this model is well known. First, it is clear that country 1 will either demand $x = p_S + c_2$ or $x = p_W + c_2$, which makes the strong type or the weak type, respectively, indifferent between accepting $1-x$ and fighting. If country 1 demands $x = p_S + c_2$, then both types accept which gives a payoff of $p_S + c_2$ to country 1. On the other hand, if country 1 demands $x = p_W + c_2$, then the weak type accepts and the strong type fights in response. This give country 1 an expected payoff of $\pi(p_W + c_2) + (1 - \pi)(p_S - c_1)$. From this point, letting $\pi^* = (c_1 + c_2)/(p_W - p_S + c_1 + c_2)$, some simple algebra reveals that the optimal choice for country 1 is to demand $x = p_S + c_2$ if $\pi < \pi^*$ and to demand $x = p_W + c_2$ if $\pi > \pi^*$. Formally, we have the following proposition

Proposition 9. *In the baseline model with two types of country 2, there is a unique equilibrium. Country 2 accepts the offer $(x, 1 - x)$ if and only if $x \leq p + c_2$. If $\pi < \pi^*$, then country 1 offers $x = p_S + c_2$ and both types accept. If $\pi > \pi^*$, then $x = p_W + c_2$. In this case, the weak type accepts and the strong type rejects.*

New Information and Uncertainty

We now show, using the two type model, that acquiring new information can increase uncertainty. Recall that π is the belief of country 1 that country 2 is the weak type and thus $1 - \pi$ gives the belief that country 2 is the strong type. As is standard, we measure uncertainty by the variance of this distribution. As this is a Bernoulli distribution, it is elementary that its variance is given by $\pi(1 - \pi)$. This is illustrated in Figure 6. As we can see, this variance is maximized at $\pi = 1/2$ and strictly decreasing away from this point.

Now suppose country 1 receives a binary signal $s \in \{g, b\}$ as described in section 3.1. The probability of receiving a good signal is equal to the true value of p. Thus, starting from the prior probability q, we find country 1's updated belief π that country 2 is the weak type after receiving a good signal:

$$\pi_g = \frac{p_W q}{p_W q + p_S(1 - q)},$$

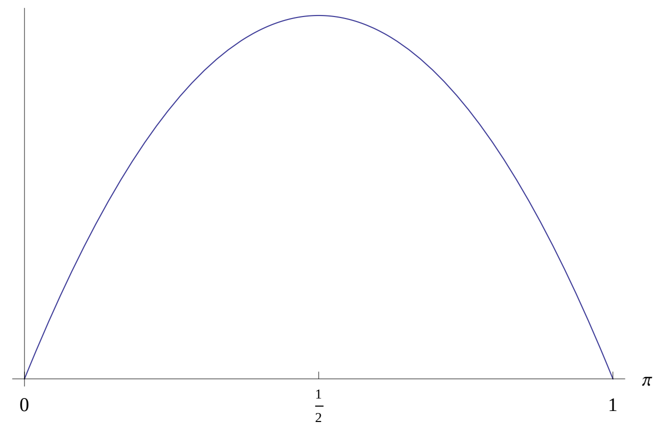

Figure 6: Variance of belief π

Likewise, country 1's updated belief after receiving a bad signal is given by

$$\pi_b = \frac{(1-p_W)q}{(1-p_W)q + (1-p_S)(1-q)},$$

It is clear that $0 < \pi_b < q < \pi_g < 1$.

As an example of a signal creating more uncertainty, suppose $q = 3/4$, $p_S = 3/5$, and $p_W = 4/5$. The variance of the prior distribution is $(3/4)(1/4) = 3/16$.

Then after observing a bad signal,

$$\pi_b = \frac{(1/5)3/4}{(1/5)3/4 + (2/5)(1/4)} = \frac{3}{5}.$$

The variance of the posterior is is $(3/5)(2/5) = 6/25$. This is larger than $3/16$ so observing this signal causes uncertainty to increase. This also follows immediately from Figure 6, as π_b is closer to $1/2$ than q is. Note also that in this example, the ex ante probability of a bad signal is $(1-p_W)q + (1-p_S)(1-q) = (1/5)3/4 + (2/5)(1/4) = 1/4$. So with probability $1/4$, country 1 receives a signal that increases uncertainty.

We can extend this observation to the case of k signals, composed of w good signals and l bad signals. Fix $q > 1/2$. Then the posterior belief π_k is given by

$$\pi_k = \frac{(p_W)^w(1-p_W)^l q}{(p_W)^w(1-p_W)^l q + (p_S)^w(1-p_S)^l(1-q)},$$

It is not difficult to check that $\pi_k \geq 1/2$ after k bad signals as long as $(1-p_S)/(1-p_W) \leq$

$(q/(1-q))^{1/k}$. When this condition on the parameters is satisfied, then uncertainty increases after receiving k bad signals.

References

Arena, Philip and Scott Wolford. 2012. "Arms, Intelligence, and War." *International Studies Quarterly* 56(2):351–365.

Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(3):379–414.

Fey, Mark, Adam Meirowitz and Kristopher W Ramsay. 2013. "Credibility and Commitment in Crisis Bargaining." *Political Science Research and Methods* 1(01):27–52.

Fey, Mark and Kristopher W. Ramsay. 2011. "Uncertainty and Incentives in Crisis Bargaining: Game-Free Analysis of International Conflict." *American Journal of Political Science* 55(1):149–169.

Filson, Darren and Suzanne Werner. 2002. "A Bargaining Model of War and Peace." *American Journal of Political Science* 46:819–838.

Kaplow, Jeffrey M. and Erik Gartzke. 2014. "Knowing Unknowns: The Effect of Uncertainty in Interstate Conflict.".

Kydd, Andrew H. 2010. "Rationalist approaches to conflict prevention and resolution." *Annual Review of Political Science* 13:101–121.

Ramsay, Kristopher W. 2008. "Settling It on the Field Battlefield Events and War Termination." *Journal of Conflict Resolution* 52(6):850–879.

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

Reiter, Dan. 2003. "Exploring the bargaining model of war." *Perspective on Politics* 1(01):27–43.

Slantchev, Branislav L. 2003. "The Principle of Convergence in Wartime Negotiation." *American Political Science Review* 97(4):621–632.

Smith, Alastair. 1998. "Fighting battles, winning wars." *Journal of Conflict Resolution* 42(3):301–320.

Smith, Alastair and Allan C Stam. 2004. "Bargaining and the Nature of War." *Journal of Conflict Resolution* 48(6):783–813.

Werner, Suzanne and Amy Yuen. 2005. "Making and keeping peace." *International Organization* 59(02):261–292.

Wittman, Donald. 1979. "How Wars End." *Journal of Conflict Resolution* 23(4):743–763.