A Little Bit of Cheap-Talk is a Dangerous Thing.

Michael F. Joseph
Yale University
michael.joseph@yale.edu

September 16, 2019

Abstract

I study a bargaining model of war where states contest two issues, and the state receiving offers has private information about her value for each issue. When the state making offers starts the game sufficiently uncertain about the receiver's favorite issue, he makes a balanced offer (i.e. distributed evenly across both issues). A balanced offer is less efficient than an offer concentrated on the receiver's most valuable issue would have been, but avoids guessing which issue the receiver values the most. I then allow players to address uncertainty about the receiver's preferences through cheap, private diplomacy. The receiver uses diplomacy to truthfully reveal her favorite issue leading to a Pareto-improving offer concentrated on the receiver's favorite issue. However, diplomacy has a dark-side: when it causes the equilibrium offer to switch from a balanced offer to a concentrated offer, it induces an equilibrium that (weakly) carries a larger risk of war. Diplomacy increases the risk of war in this case because the state making offers confronts a risk-return trade-off over how much more the receiver values one issue over the other. When the cost of war is low he prefers to accept some risk, rather than make a balanced offer, to capitalize on the efficiency gains from a concentrated offer. The results track with Sino-American and Sino-Soviet pre-crisis diplomacy.

Keywords: Diplomacy, crisis bargaining, motives, game theory, great power war

Note: A version of the Paper is Forthcoming in the Journal of Politics.

Over the past two decades the United States has expressed both optimism and concern about how China communicates its intentions. Based on China's diplomacy, American leaders understand that some concessions (e.g Taiwan) clearly fit within China's core interests and others do not (e.g a Ugandan colony). However, ambiguity remains. The United States warns that imprecision is a source of tension in Sino-American relations that may lead to accidental conflict. As Deputy Secretary of State Robert Zoellick warned, "Many countries hope China will pursue a peaceful rise, but none will bet their future on it." Zoellick noted that "China should openly explain its defense spending, *intentions*, doctrine and military exercises to ease concerns." Yet China refuses to distinguish between which core interests are most important or even provide an exhaustive list of what those core interests are.

Chinese leaders were not always so mysterious in the way they expressed their intentions. In the 1960s, Mao expressed his dissatisfaction with the distribution of territory along the Sino-Soviet border. Wanting to avoid conflict with a nuclear-armed ally, Khrushchev sent senior diplomats to learn more about China's interests and negotiate a settlement. During negotiations, the Chinese explained that the Eastern region of Xinjiang and the Zhenbao Island were of greatest importance. Consistent with Mao's priorities, Khrushchev offered Mao control over the Zhenbao Island and large parts of Eastern Xinjian in 1966. Despite precise information about China's preferences, the offer was insufficient. The Chinese military invaded parts of Xinjiang resulting in between 130 and 800 battle deaths.

How does private diplomacy influence state strategy? For a long time, scholars thought that the incentive to misrepresent intentions was overwhelming (Waltz 1979; Carr 1964). Thus, rational states could not reveal information about their preferences either at all (Mearsheimer 2001) or at least without costly signals (Kydd 2005) or audience costs (Fearon 1994). Recent research finds that states can credibly transmit some information using cheap-talk (Sartori 2002; Kurizaki 2007; Kurizaki and Whang 2015; Trager 2011, 2013). One insight is that states have heterogeneous preferences that lead them to value some concessions more than others (Moravcsik 1998). For example, China would (probably) not accept territorial control over Uganda instead of Taiwan even though Uganda is larger. China's historical context and interests imply China *prefers* Taiwan over Uganda. Research shows that the incentive to misrepresent does not apply to variation of this sort (Trager 2011; Chakraborty and Harbaugh 2010; Battaglini 2002). States making offers want to reveal the order of their preferences to minimize the offers they are forced to give. States receiving offers want to receive valuable concessions first. Thus, at the very least, states can communicate information about their preference order. Consistent with this reasoning both the United States and the Soviet Union thought China's diplomatic statements provided vital information about the order of China's preferences.

While China's persuasive diplomacy matches the rational diplomatic mechanism, the patterns of war and peace that follow do not. We think that diplomacy reduces uncertainty in international relations, which leads to more efficient bargains and less war (Kurizaki 2007; Kurizaki and Whang 2015; Sartori 2002; Trager 2011). But in the 1960s, China and the Soviet Union bargained over a few well-defined territorial issues along their borders. China provided detailed information about its preferences and war occurred anyway. In the modern world, Sino-American diplomacy covers dozens of issues including territory, trade, global order and international norms. China has refused to reveal detailed information about its preferences despite constant warnings that this ambiguity may ignite conflict. Yet war has not occurred. If diplomacy works, why do less detailed messages in complex bargaining scenarios induce peace, but more detailed messages in simple settings induce war?

This paper explores the dark-side of diplomacy. In it, I study a spatial bargaining model of war (cf Fearon 1995) between two states—an Offerer (of bargains) and a Sender (of messages)—that bargain over two issues. Rather than vary the Sender's value for these issues relative to the cost of war, I analyze cases where the Sender holds different values for each issue under dispute (a-la Trager 2011; Battaglini 2002). Variation in the Sender's value for each issue creates opportunities for more efficient offers: if the Offerer can concentrate on the Sender's favorite issue, he can produce a small, high-valued (to the Sender) concession and keep more of the surplus for himself. The problem is that the Sender has private information about her valuation. Thus, the Offerer is uncertain about which issue the Sender cares about the most, and how much more the Sender values one issue over the other. I contrast equilibrium behavior in two worlds: one where costless diplomacy is allowed, and another where it is not.

I show that when the Offerer is very uncertain about the Sender's favorite issue[1] he makes a balanced offer evenly distributed across both issues. This offer is safer because it avoids confronting uncertainty about the Sender's relative preferences. But it is inefficient because it concedes part of the issue that the Sender cares less about. When the game starts with high uncertainty about the Sender's favorite issue and there is no opportunity to communicate, the Offerer makes this balanced offer.

I then allow the Sender to send a costless diplomatic message about her preferences. I find a unique informative equilibrium where the Sender credibly reveals her favorite issue but transmits no information about how much she cares about one issue relative to the other. This message entices the Offerer to concentrate on the Sender's favorite issue, leading to a Pareto-improving offer. However, the Offerer is still uncertain about how much he can exploit the Sender's heterogeneous preferences because he is uncertain about how much the Sender values one issue over the other. This creates a risk-return trade-off because types that value one issue much more than the other will accept a very small offer. But types that have approximately equal values require larger offers. To capitalize on the Sender's different values for each issue, the Offerer must accept some risk that the offer will be insufficient. Under a balanced offer, she would avoid this type of risk.

It follows that when a diplomatic message induces the Offerer to switch from a balanced to a concentrated offer it raises the risk of war. This happens because the concentrated offer must confront uncertainty over the Sender's relative value between two issues, but a balance offer does not. In cases where diplomacy induces a shift in strategy, it increases both players'

[1]That is, there is a roughly equal chance that the Sender values issue 1 more than issue 2

expected value but also increases the risk of war they accept.

To better understand the conditions under which effective diplomacy raises or reduces the risk of war, I study an extension that alters the distribution from which the Sender's values are drawn. In the baseline model, the Sender is equally likely to value issue A more than B. In the extension, the Sender is more likely to value issue A over issue B. I show that when the cost of war is low and the Offerer is sufficiently confident that issue A is the Sender's favorite issue, the Offerer makes a concentrated offer even absent diplomacy. In this case, including diplomacy reduces the chance of war because it helps the Offerer target his concentrated offer but does not change the Offerer's strategy—the Offerer always concentrates his offer on one issue. However, when uncertainty about the Sender's relative value between issues is high, the Offerer makes a balanced offer to avoid confronting risk. In this case, diplomacy induces a different offering strategy that raises the risk of war because it entices the Offerer to confront a risk-return trade-off that he would have avoided if he did not know where to concentrate his offer. This demonstrates diplomacy's affect on risk-taking depends on the Offerer's preferred strategy before communication takes place. At times of high prior uncertainty, when coordination is most needed, diplomacy's dark side emerges. However, in cases where states have reasonably good information about each other, then diplomacy has pacific effects that are consistent with existing findings.

To understand how more complex foreign relations between great powers can influences the result, I study the case where states bargain over many issues simultaneously. More complexity makes diplomacy much more effective than it is in the two-issue bargaining game. However, the Sender always does better in scenarios where uncertainty is high because the Offerer is less willing to take risks. Thus, the Sender prefers simple settings where states bargain over fewer issues. The reason is that the Offerer will exploit additional information to make low-ball offers. The Sender wants to provide just enough information to entice the Offerer away from a balanced offer — but not enough information to induce a low-ball offer.

My theory advances research that questions the assumed relationship between uncertainty and war (Fey and Ramsay 2011; Leventŏlu and Tarar 2008; Arena and Wolford 2012; Bils and Spaniel 2016; Debs and Weiss 2016) by identifying a new mechanism for war that paradoxically comes from greater access to information. It also extends research on cheap-talk diplomacy by showing that incentives to communicate are mixed in ways that can lead to war even if offers are more efficient (Trager 2013; Kurizaki 2007; Sartori 2002).

1 Motives, Bargaining and War

In standard bargaining models of war, states are assumed to value each piece of the pie uniformly (Fearon 1995). This simplification helps scholars analyze relationships between power and resolve, or shifting power and war (Powell 1999). Yet growing evidence locates the source of international conflict in claims over territories for historical and cultural reasons (Moravcsik 1998; Jackson and Morelli 2011). States fight to unify their ethnic group (Goemans and Schultz 2013), restore borders (Carter and Goemans 2011), security, or one of many other principles that motivate their foreign policies. In the modern world maintaining large militaries and taking territory has a high financial cost. This means that foreign policy conquest rarely increases a state's economic welfare (Brooks 1999); especially when states benefit significantly from international trade (Keohane 2005). As a result, states are only willing to contest territories that satiate the particular principles they hold high.

Economists have shown that heterogeneous preferences provide greater opportunity for coordination and communication under uncertainty in models of spatial bargaining (Chakraborty and Harbaugh 2003; Jackson, Simon, Swinkels, and Zame 2002), expert-decision-maker interactions (Chakraborty and Harbaugh 2010; Crawford and Sobel 1982) and auctions (Kim and Kircher 2015). Recently, political scientists have applied these insights to expand upon the bargaining model of war. Different studies emphasize different ways that state preferences may vary. Some focus on preference divergence between two states (Spaniel and Bils 2017; Trager 2013). Others emphasize how one state may hold different values for different issues (Trager 2011).[2] Although each study is different, the common result is that cheap-talk helps states communicate their preferences leading to more peace.

In what follows, I borrow from all of these approaches to operationalize heterogeneous preferences. My goal is to appropriately characterize variation in motives in a context that captures the dynamics of inter-state bargaining described by informal studies of conflict. States: (1) bargain over multiple issues simultaneously; (2) value some issues more than others; (3) are uncertain about what their rivals' values the most, and by how much; (4) can make partial or full concessions over any single issue or group of issues in an effort to sustain peace; and (5) lose something they value when they make a concession.

I depict this idea in Figure 1 in the context of Sino-American bargaining. Each panel represents a different "type" of China. The x-axis is a list of four issues China and the United States may dispute.[3] The y-axis is China's value for each issue. If China values restoring its borders intensely, then it has extreme value for concessions in Taiwan relative to all other issues. If China truly values regional hegemony in East Asia and the Pacific, then it places roughly equal value over concessions in that region. All of those concessions are more valuable than the others—but no particular concession stands out as vital. Preference order is captured by comparing values across panels (or different types of China). If China secretly wants to capture natural resources, then concessions in Central Asia are more valuable than Taiwan. However, the reverse is true if China wants to restore borders.

My operationalization of heterogeneous preferences is most similar to Trager (2011). However, our findings are different because we make different assumption about which bargaining settlements are allowed. Like Trager, I assume that states bargain over multiple issues and hold different values for each issue. But in Trager's model, each issue cannot

[2]All of these cases assume that states have at least some conflict in preferences. Analyses of states with either compatible or conflicting preferences are outside my scope (Kydd 2005; Jervis 1978).

[3]Although I depict each value function separately, states may hold any combination of these principles. This facilitates wide variation in the plausible value functions. The list is not exhaustive.

be divided. Rather the Offerer chooses to concede (or not) issues in full. This assumption creates conditions where both states prefer war to a peaceful settlement because the Offerer is not allowed to make compromises over each issue. In contrast, I focus on a continuous bargaining space and continuous valuations for each issue. In my model there is always a settlement that both states prefer to war.

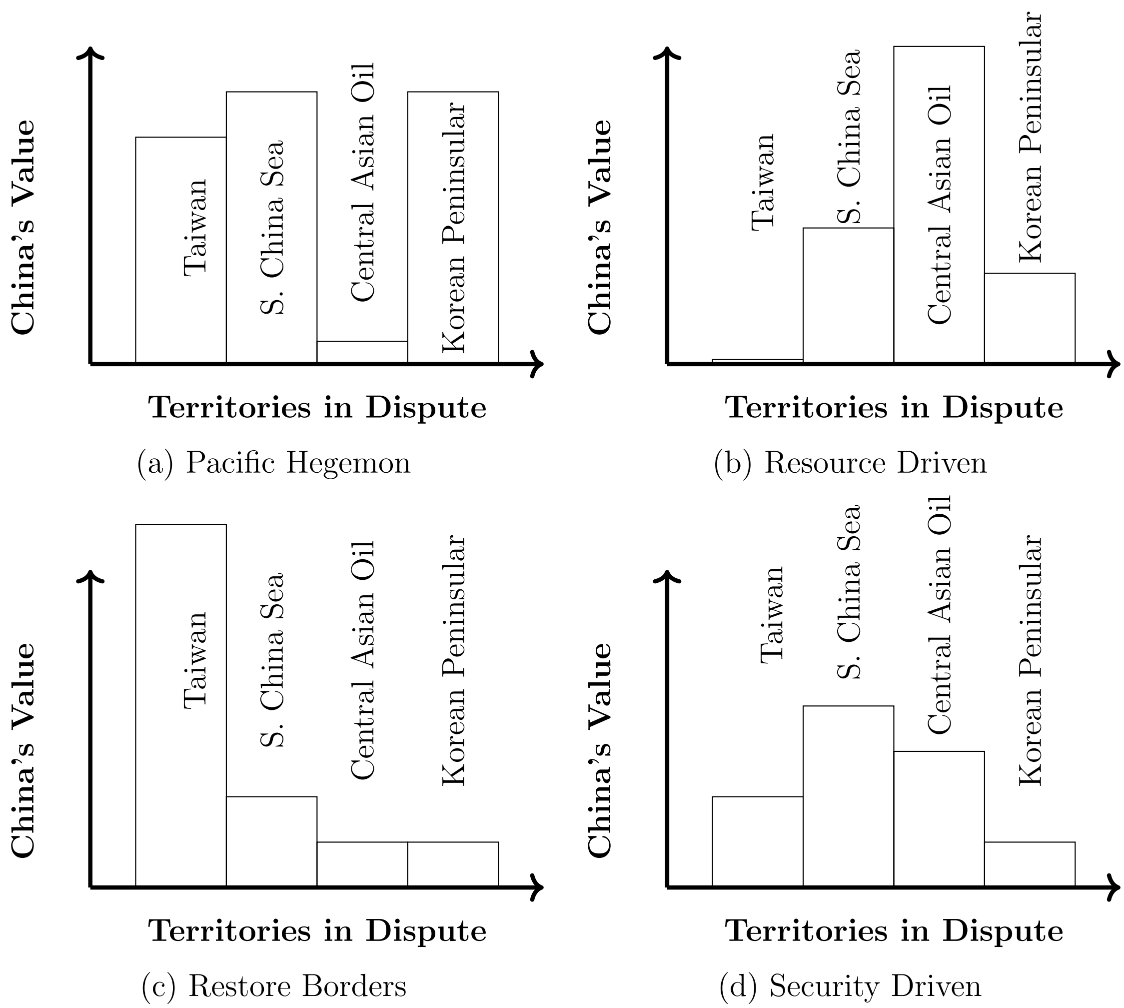

Figure 1: Hypothetical Chinese interests given different foreign policies

I study pre-crisis diplomacy: a period of cheap diplomatic exchange in which states can signal private information to each other. Unlike costly signals and audience costs (Fearon 1994; Kydd 2005), these messages do not exogenously alter the payoffs for each player. That period is followed by a crisis where one state is forced to make concessions to another.

2 A Bargaining Model of diplomacy-induced risk

I develop a bargaining model of war between two players—an Offerer of bargains (O, male) and a Sender of messages (S, female)—who bargain over multiple issue. I assume O is uncertain about S's value for each issue and S can signal her valuation for each issue through costless messages. I claim that an equilibrium exists where S credibly signals information about her value leading to Pareto-improving offer that raise the risk of war. To demonstrate the core mechanism the main model focuses on the case where: O and S bargain over only two issues; and S's value for both issues is drawn independent and identically distributed from a common distribution. I then extend the model in two ways. First, I adjust O's prior beliefs such that S is more likely to value one issue over the other. Second, I increase the number of issues states bargain over.

2.1 Main Model

I model an interaction between two players, O and S, that bargain over two divisible issues, n_1 and n_2. These issues might represent two different territories, or preferences for two different global norms. I refer to an arbitrary issue as n_j where $j \in \{1, 2\}$. In single-issue bargaining games, an offer is a single partition of one issue. In simple terms, there is one bucket, and O decides how much to fill it up. In this game, O may want to offer different amounts of each issue. That is, each issue n_1, n_2 is its own bucket and O chooses how much to fill each bucket. O's offering strategy $q(q_1, q_2)$ includes two different offers $q_1, q_2 \in [0, 1]$ where subscripts correspond to issues n_1, n_2.

I assume that S has heterogeneous preferences by assigning her different values for the issues in dispute: θ_1, θ_2. Subscripts make clear that θ_1 is S's value for n_1. S's utility from accepting an offer is $U^S(accept, q_1, q_2) : \frac{1}{2}(\theta_1 q_1 + \theta_2 q_2)$.[4] Through this utility function, the θ_j act as S's value for each issue where higher draws of θ_1 imply S has more value for n_1. For simplicity, I assume O values both issues the same. O's total value from an offer that is

[4] Dividing by $\frac{1}{2}$ ensures that D's value for the pie is 1. Later I will consider a bargain over $J >> 2$ issues. In that extension dividing by J is necessary to keep the total utilities similar.

accepted is $U^O(accept|q_1, q_2) : 1 - \frac{q_1+q_2}{2}$. Since O is indifferent between offers against these issues, he cares only about the total length of both offers. To focus on this length, I define $t = \frac{q_1+q_2}{2}$ and re-write O's utility as $U^O(accept|q_1, q_2) : 1 - t$.

To create uncertainty about S's preferences, I assume θ_1, θ_2 are drawn i.i.d from a common distribution $\Theta \sim Unif\,[0, 1]$. S's type is the realization of θ_1, θ_2. The total set of types is of size Θ^2 and is uniform in two important senses: every type is equally probable; and both n_j have an equal probability of realizing any value in Θ. To be clear, each issue is just as likely as the other to be more valuable to S: $pr[\theta_1 > \theta_2] = pr[\theta_2 > \theta_1] = 1/2$.[5] Of course, in any realization of S's type one of these issues will be more valuable than the other. I refer to S's most valuable issue $\max[\theta_1, \theta_2] = \theta_A$ and S's least valuable issue $\min[\theta_1, \theta_2] = \theta_B$ with corresponding q_A, q_B, n_A, n_B.

I allow S to send messages about her type $m(\hat{\theta}_1, \hat{\theta}_2)$ where S can choose from messages $\hat{\theta}_1, \hat{\theta}_2 \in \{\Theta, \emptyset\}$.[6] The message is cheap because sending any message does not exogenously affect payoffs. I say O's initial belief about S's type is $\gamma(\Theta^2)$ and O's posterior belief after it receives a message is $\gamma^\dagger(\Theta^2, m())$. A message is informative if $\gamma(\Theta^2) \neq \gamma^\dagger(\Theta^2, m())$.

The game proceeds as follows. Nature draws θ_1, θ_2 and reveals it privately to S. S signals a type $m(\hat{\theta}_1, \hat{\theta}_2)$ to O. O processes the signal and makes a take-it-or-leave-it offer $q(q_1, q_2)$. S either accepts the offer or rejects it in favor of war.

I assume that war is a costly lottery where the winner takes both issues in full. S wins the lottery with probability p, and O wins with probability (1-p). Following Powell (2006), war damages the pie and is modeled as an attenuation factor $d \in (0, 1)$ (d for damage). Therefore, $EU^S(War) = pd\frac{\theta_1 q_1 + \theta_2 q_2}{2}$ and $EU^O(war) = (1 - p)d$. For ease of exposition we write $pd = \alpha$ and $\beta = (1 - p)d$. I focus on the case where S is not too powerful: $\alpha < 1/2$.[7]

I model the cost of was as a factor (pd) rather than a constant $(p - d)$ to focus on uncer-

[5]Nature can draw $\theta_1 = \theta_2$ but the expectation of this outcome is 0.

[6]To be clear, S's message claims her value for n_j is $\hat{\theta}_j$. Her actual value is θ_j.

[7]When α is large, there is an additional condition on equation 6 that does not change my findings.

tainty about the Sender's relative preferences between different issues and not uncertainty about the Sender's value relative to the cost of war. Uncertainty about relative preferences produces complex strategic incentives for the Offerer because the Offerer must consider which issue the Sender values the most, how much more the Sender values one issue over the other. It is these dimensions of uncertainty, not uncertainty about the value of bargains relative to the cost of war that drives my result.[8]

2.1.1 Analysis

I argue that diplomacy can provide enough information to induce Offerers to switch from an offer distributed across two issues to a an offer concentrated on a single issue. When it does, it raises the risk of war. To establish this claim, I structure the analysis in three sections. First, I define two different offering strategies that O might plausibly select from in equilibrium and contrast S's minimum demands when faced with these different offers. Second, I introduce uncertainty about S's type. Since the purpose is to understand the consequences of allowing diplomatic communication, I start with a counter-factual: what happens if S has private information about her type and is not allowed to send messages? I then allow S to send diplomatic messages and solve for an equilibrium where S sends a credible message about her type. Third, I make an inference based on the following comparison: when diplomacy is not allowed O makes an offer that leaves both players with

[8]In Appendix C I introduce uncertainty about preferences relative to the cost of war by modeling the cost of war $p - d$. My main finding is the same: diplomacy can cause O to switch from a balanced to a concentrated offer. When it does, O accepts a larger risk of war. There is one difference: when the cost of war is modeled as $p - d$, diplomacy is prohibited, and d is small, O's equilibrium offer accepts a risk of war (but the risk of war is still larger when I include diplomacy). This is not true when I use the functional form pd. Given this difference, I only claim that diplomacy can raise (rather than create) a risk of war. The results would be different if I assumed S's value for each issue was known, and S instead had different costs of war (d_1, d_2). When O is uncertain about S's values he is uncertainty about S's value for two offers of the same length. This creates opportunities for inefficient peaceful settlements. When O is uncertain about S's different costs of war he can always make the most efficient offer. Thus there is only uncertainty about S's value relative to the cost war.

less (or equal) expected utility but carries a lower (or equal) chance of war than when diplomacy is allowed.

To begin, I define two offering strategies. I define a balanced offer where O offers equal parts of both issues: $\tilde{q}_1 = \tilde{q}_2 = t/2$. I notate it $\tilde{q}$ with corresponding $\tilde{t}$. Notice that when O makes a balanced offer, S's minimum demand does not depend on her type. S accepts a balanced offer if and only if $\frac{\tilde{t}(\theta_1+\theta_2)}{2} \geq \frac{\alpha(\theta_1+\theta_2)}{2}$, or $\tilde{t} \geq \alpha$.

I define a concentrated offer ($\hat{q}$, with corresponding $\hat{t}$) as one where O makes an offer against only one issue: $\hat{q}_j = \hat{t}, \hat{q}_k = 0$. Notice S's minimum demand now depends on her type. Suppose, O makes a concentrated offer against n_1, S only accepts it if: $\frac{\hat{q}_1\theta_1}{2} \geq \frac{\alpha(\theta_1+\theta_2)}{2}$, or $\hat{q}_1 \geq \frac{\alpha(\theta_1+\theta_2)}{\theta_1}$.

Figure 2 summarizes these two offering strategies. The top row describes a balanced offer, the bottom describes a concentrated offer. Column (a) reports notation and analytical values for each strategy. Column (b) and (c) depict four examples of S's minimum demand given different draws of θ_1, θ_2 and O's choice of offering strategy. In column (b) θ_1 is much larger than θ_2. In (c) θ_1 and θ_2 are closer together. Each picture represents the offer O must make to leave S indifferent with her war payoff: $\alpha \times \frac{\theta_1+\theta_2}{2}$.

Looking across the first row, the balanced offer that meets S's minimum demands does not change even though S's relative value between the two issues varies. The reason is that S's high value for one issue balances out her low value for the other. Looking across the second row, the concentrated offer that leaves S indifferent with her minimum demand varies with S's relative value between issues. When θ_1 is much larger than θ_2, S will accept a smaller concentrated offer ($\hat{t}$ is less) than she would have accepted if θ_A and θ_B were closer together. At one extreme, if $\theta_B = 0$, then S's minimum demand is: $\hat{t} = \hat{q}_A = \alpha$. At the other extreme, when S values both issues the same ($\theta_A = \theta_B$), S's minimum demand is $\hat{q}_A = 2\alpha$. Of course, this assumes that the concentrated offer targets S's favorite issue (n_A). If O made a concentrated offer against S's least favorite issue (n_B) then the concentrated offer that met S's minimum demand would be larger than 2α.

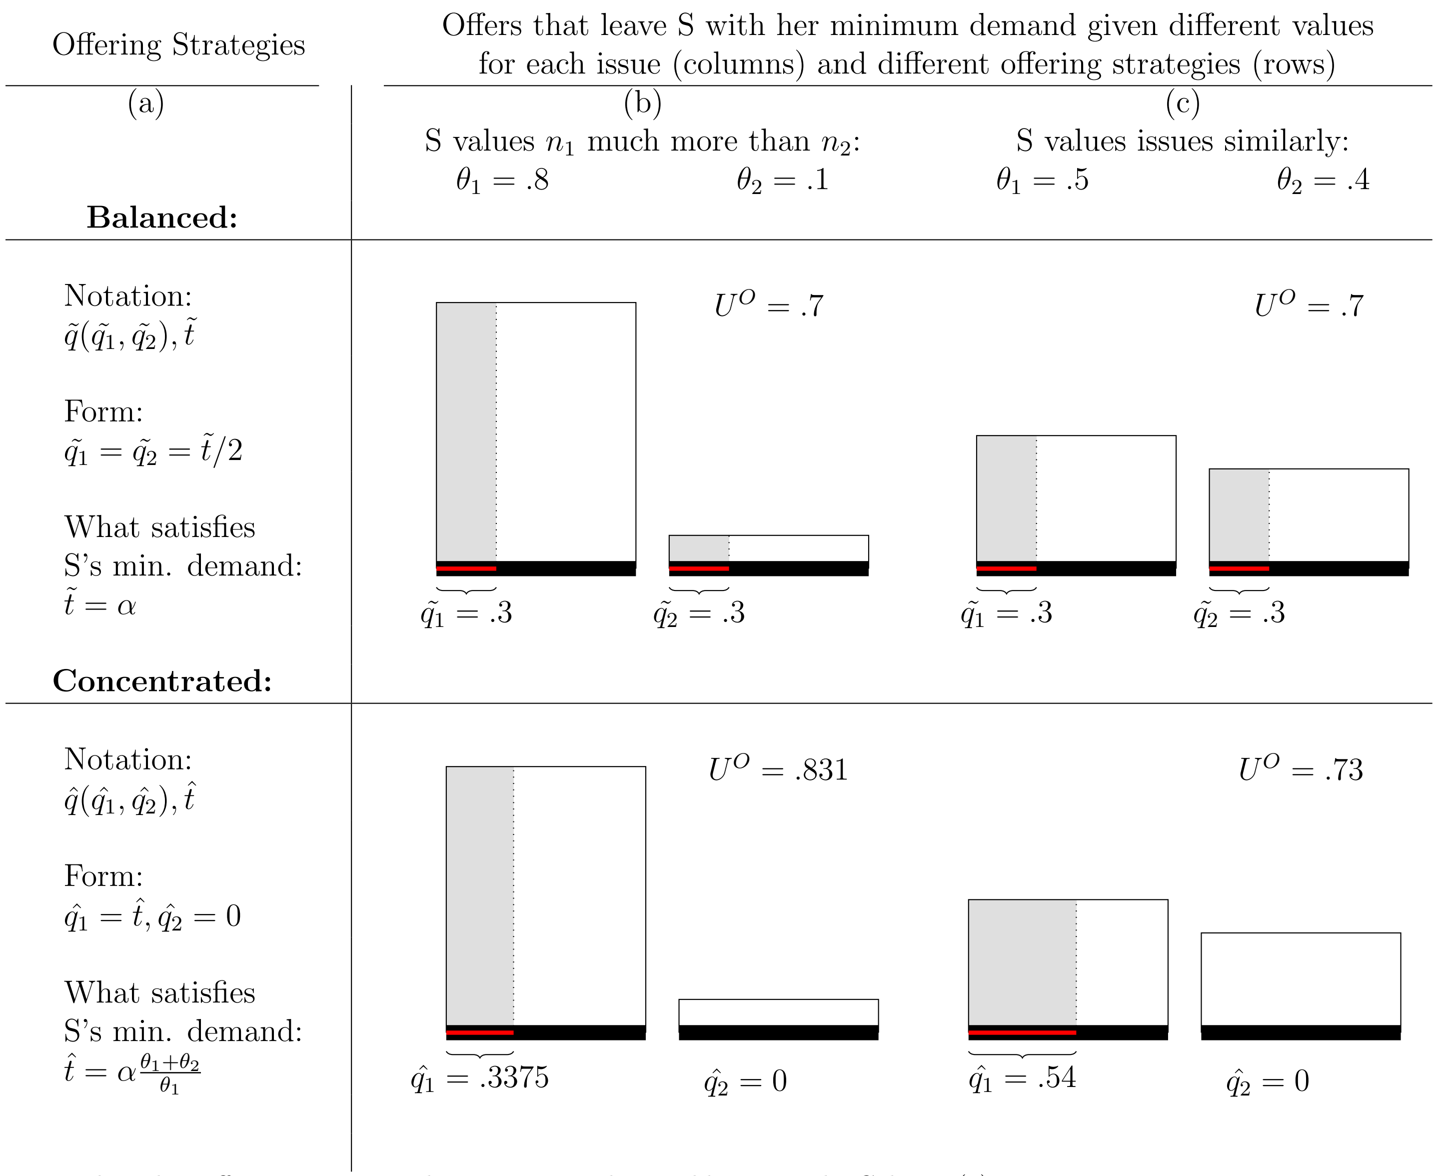

Rows describes offering strategies that appear on the equilibrium path. Column (a) summarizes notation for these offers. Columns (b) and (c) depict the offers that meet S's minimum demand given different draws of θ_j, and different offering strategies. In each picture, the thick black lines represent the two issues in dispute n_j. The area above intervals represents S's value (θ_j) for each issue. The red lines mark the length of O's offers against each issue. The area shaded gray is S's utility from accepting the offers. U^O is O's utility if S accepts. Assumes $\alpha = .3$.

Figure 2: S's minimum demand with different offering strategies

O wants to make the smallest (in length) offer that S will accept to extract as much of the surplus as possible. Clearly, when O is completely informed of S's type he makes a concentrated offer $\hat{q}_A = \frac{\alpha(\theta_1+\theta_2)}{\theta_1}, \hat{q}_B = 0$ that matches S's minimum demand. But uncertainty complicates O's decision. When S's type is private, O does not know S's minimum demand from a concentrated offer. If O makes a concentrated offer, he must guess which issue S values the most (to determine where to concentrate his offer) and by how much (to determine how large an offer to make). If O makes a balanced offer he need not confront uncertainty about S's type because O's best balanced offer is the same no matter S's value for each issue.

First, I analyze a counter-factual case where S cannot send messages. I define the balanced offer that maximizes O's expected utility given O's beliefs as $\tilde{q}^*$ with corresponding $\tilde{t}^*$, and the concentrated offer that maximizes O's expected utility as $\hat{q}^*$ with corresponding $\hat{t}^*$.

Proposition 2.1 *Suppose a counter-factual setting where diplomacy is impossible (i.e S cannot signal), then there is a unique equilibrium where O makes a balanced offer that sets $\tilde{t}^* = \alpha$ leaving no chance of war.*

The proof is in Appendix A.1. When O was informed about S's type he always made a concentrated offer. Proposition 2.1 argues that this is not true when S has private information and cannot communicate. I show why this difference arises referencing Figures 2 and 3. Since S's only decision is to accept or reject the offer, I focus on O's offering decision taking as given S will reject any offer below her minimum demand.

O's equilibrium offer depends on a three part optimization problem. First, O identifies the balanced offer that maximizes his expected gains ($\tilde{q}^*$). Second, *O* identifies the concentrated offer that maximizes his expected gains ($\hat{q}^*$).[9] Third, O determines whether his expected utility is larger from offering $\hat{q}^*$ or $\tilde{q}^*$.

Since O's best balanced offer does not depend on S's type, the optimal $\tilde{q}^*$ always sets $\tilde{t} = \alpha$ with O offering an $q_1 = q_2 = \alpha$ of each issue. O receives exactly $U^O(\tilde{q}^*) = 1 - \alpha$.

[9] To be clear, I do not restrict O's offer to one issue. Rather, the most efficient offer that is not balanced must be concentrated in one issue. Thus, I focus on it.

In contrast, when O makes a concentrated offer he must confront two types of uncertainty: which issue S values the most; and how much S values θ_A relative to θ_B. O's value from making an offer can be summarized as:

$$EU^O(\hat{q}, \gamma^\dagger(\Theta^2)) : (1 - \hat{t})\psi_t\omega + (1 - \psi_t\omega)\beta \tag{1}$$

In this equation, the first term is O's payoff from making an offer of length $\hat{t}$ against a single issue, multiplied by the proportion of types that an offer of length $\hat{t}$ will satisfy (ψ_t). ω is the probability that O selected the most valuable issue to offer against. The second term is the probability that O either did not select the largest issue to offer against, or that the offer was not large enough given the ratio between the two issues $(1 - \psi_t\omega)$ and multiplied by O's war payoff (β).[10]

I separate two dimensions of uncertainty to illustrate that there are two reasons a concentrated offer can fail. First, ψ represents uncertainty about the relative value between θ_A, θ_B. O manages this kind of uncertainty by varying the size of his offer $\hat{t}$. ω represents uncertainty about which issue S values the most. O manages this uncertainty by choosing an issue to offer against.

I depict how these different dimensions of uncertainty affect O's offer in Figure 3. Both panels plot all possible realizations of θ_1, θ_2 along the x and y axes respectively. The dashed line marks types $\theta_1 = \theta_2$. Above that line $\theta_2 > \theta_1$.

Panel (a) emphasizes the different levels of risk that O accepts from varying $\hat{t}$ given a concentrated offer against n_1. Each colored line assumes a different value for $\hat{t}$ and demarcates the types that are indifferent between accepting and rejecting that offer. All types below a colored lines accept that $\hat{q}_1$.

Figure 3 clarifies O's risk return trade-off from increasing the size of the offer $\hat{t}$. When $\hat{q}_1$ is high more types accept that offer but this leaves O with less of the surplus. Notice that

[10]The model allows for O to offer a partially concentrated offer that includes uneven but positive offers against n_1 and n_2. However, this never arises in equilibrium.

Plots depict all possible draws of θ_1, θ_2 and which types accept a concentrated offer targeted on issue n_1.

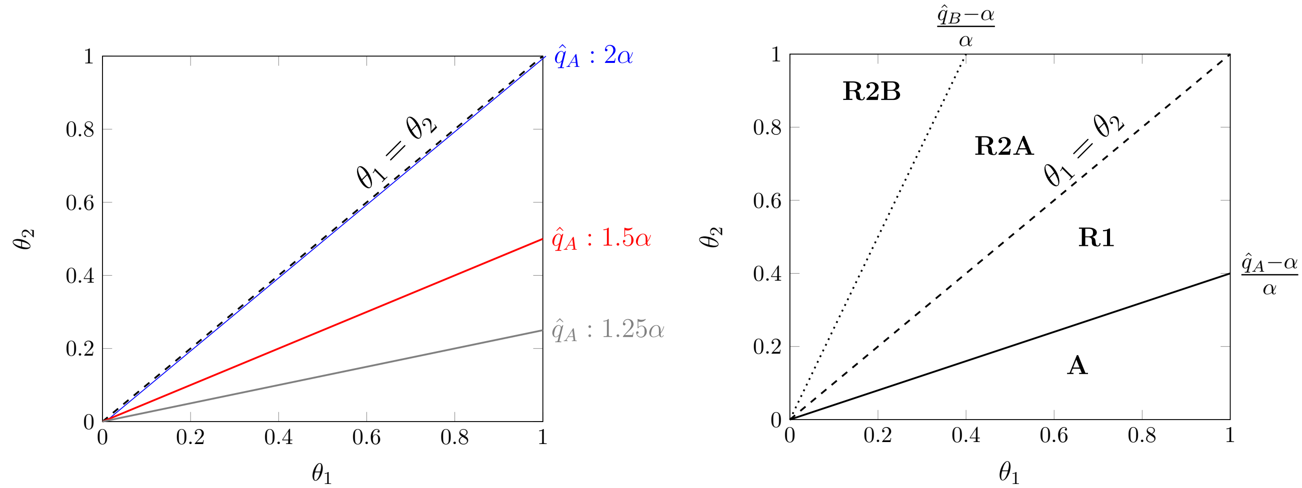

(a) Who accepts given varying lengths $\hat{t}$.

Proportion of types that accept concentrated offers of different sizes. Each colored curve depicts an offer of a different lengths $\hat{t}_A$. Area under these curves are types that accept this offer.

(b) Who accepts $\hat{q}_1$ and why.

This plot demarcates spaces based on an arbitrary offer $\alpha < \hat{q}_1 < 2\alpha$. Types of S in space **A** accept an offer $\hat{q}_1$. Types in each **R** space reject an offer $\hat{q}_1$. But they reject for different reasons.

Figure 3: Risk a concentrated offer fails give S's type

the blue line plots a concentrated offer $\hat{q}_1 = 2\alpha$. All types $\theta_1 \geq \theta_2$ accept this offer. But this is only half of all feasible types. This is a surprisingly large offer for only half of the types to accept. If O made a balanced offer of the same length: $\tilde{q} = \tilde{q}_2 = \alpha$, $\tilde{t} = \alpha$ all types would accept it. O does so poorly from a concentrated offer because all types who value $\theta_2 > \theta_1$ reject it. In effect, raising a concentrated offer only reduces the risk created by (ψ). It does not reduce the risk O must face from guessing where to concentrate his offer (ω).

Panel (b) emphasizes the different reasons that S might reject an offer $\hat{q}_1$ concentrated in n_1. The different **R** spaces identify the different reasons that S would reject the offer $\hat{q}_1$. Types in space **R1** value $\theta_1 > \theta_2$. They reject the offer because of the risk-return trade-off created by ψ. All the types in **R2** would reject the offer because n_1 is not their favorite issue. Thus, the offer fails because of the risk created by ω. In particular, the space **R2B** marks all the types that would have accepted an offer of length $\hat{t}$ if it was concentrated on n_2 not n_1. Because O has no information about S's type beyond his priors, an identical problem emerges of O concentrates his offer in n_2.

Figure 3 illustrates the limits of a concentrated offer when diplomacy is impossible because O cannot identify which issue is more valuable. Given that both issues are drawn from a common distribution, $\omega = 1/2$, any concentrated offer $\hat{t} \leq 2\alpha$ has at least a 50% chance of failure because there is a 50% chance that O has made the offer against n_B. The risk that O picks the wrong issue to offer drives O to balance because balanced offers are always accepted no matter S's type. This guarantees both players: $U^O : 1 - \alpha$, $U^S : \alpha$. This value is always more than the expected value of a concentrated offer when $\omega = 1/2$. Since the uncertainty created by ω dominates O's strategic problem, O prefers to balance, rather than deal with the risk return trade-off created by ψ.[11]

I now study the case where S can signal her type using a cheap, private message $m(\hat{\theta}_1, \hat{\theta}_2)$. I focus on a messaging strategy where all types of S choose a signal that mixes over the values that are consistent with their preference ordering. Under this messaging strategy, any type $\theta_1 \geq \theta_2$ chooses a message at random from the set of types $\theta_1 \geq \theta_2$. I write this message as $m(A = 1)$ because it has the effect of signaling S's most valuable issue is θ_1. All types $\theta_2 > \theta_1$ chooses a message at random from the remaining set. I write this message as $m(A = 2)$. In general, I call this messaging strategy $m(A)$.[12] Finally, I say a concentrated offer is a gamble if it accepts a risk of war.

Proposition 2.2 *A cheap-talk equilibrium that uses messaging strategy $m(A)$ always exists:*

- *S observes her type and sends a randomly chosen message from the set $m(A = 1)$ if $\theta_1 \geq \theta_2$ and $m(A = 2)$ otherwise.*

[11] This no diplomacy, counter-factual produces the same result as the babbling equilibrium when diplomacy is allowed. In the babbling equilibrium, all S choose a message at random. O makes a balanced offer.

[12] As with all cheap-talk models several messaging strategies can survive equilibrium. I rule out most off-path messages because they have been shown not to alter the results in similar models (Chakraborty and Harbaugh 2010). In Appendix A.2.1 I consider other messages. I show a message similar to $m(A)$ with out of equilibrium messages is robust. However, S cannot send a credible message that reveals how much S values issue A relative to B because S faces incentives to over-state her interests (a-la Fearon 1995).

- *O processes that signal and updates her beliefs* $\gamma^\dagger(.) \neq \gamma(.)$ *such that* $pr(\theta_1 \geq \theta_2 | m(A = 1)) = 1$ *and* $pr(\theta_2 > \theta_1 | m(A = 2)) = 1$.
- *O offers a concentrated offer* $\hat{q}^*$ *contained only in issue* n_A *which:*
  - *is a gamble of length* $\hat{t}^* = \frac{2-2\beta+\alpha}{4}$*, with positive probability of war* $\psi^* = \frac{2-2\beta-\alpha}{2\alpha}$ *if* $\psi^* < 1$*; and*
  - *is of length* $\hat{t}^* = \alpha$*, with no risk of war otherwise.*
- *S accepts that offer if* $\frac{\theta_A \times \hat{t}^*}{\theta_A + \theta_B} \geq \alpha$ *and rejects otherwise.*

The proof is in Appendix A.2. Here I provide an intuition for why diplomacy leads O to shift from a balanced offer to a concentrated offer. Recall that when diplomacy was absent, O did not make a concentrated offer because the risk that O's concentrated offer targeted n_B was 50%. But the message $m(A)$ ensures that S credibly reveals her most valuable issue causing $\omega = 1 | \gamma^\dagger(\Theta^2, m(A))$.

Referring to Figure 3(b), the signal wipes out all of the types **R2**, halving the set of feasible types and ensuring that O's concentrated offer targets n_A. Assuming S sends a message $m(A = 1)$ the proportion of types that accept a concentrated offer $\hat{q}_1$ (which is $\psi\omega$) are marked by regions **A**/ (**A** + **R1**).

At minimum, O is now indifferent between a concentrated offer $q_A = 2\alpha, q_B = 0$ and a balanced offer $q_A = \alpha, q_B = \alpha$ because all feasible types accept both offers. Effectively, wiping out all risk created by ω. Now O's choice focuses on the risk return trade-off created by ψ. Analogous to the offering strategy studied by Fearon (1995), when the cost of war is sufficiently low, O prefers to make an offer $\hat{q}_A < 2\alpha$ because the gains from making a smaller offer outweigh the risk that some types will reject it (and the consequences of that rejections).

In Figure 4 I plot O's expected utility from the equilibrium concentrated offer and the probability of war. The dashed vertical line reflects the point where O's optimal concentrated offer risks war. Four facts about the results are interesting. First, a credible signal and concentrated offer is always an equilibrium strategy. Second, when O's cost of war is low (β is high), the equilibrium concentrated offer takes on a risk of war $0 < \psi^* < 1$. Third, when and only when there is a positive risk of war, gambling is strictly Pareto improving on a balanced offer for both players. Fourth, without signaling, O would always make a balanced offer. Signaling induces a gamble equilibrium in this range with a positive probability of war. Therefore, in this range, cheap-talk diplomacy alters strategic behavior in a way that raises the risk of war.

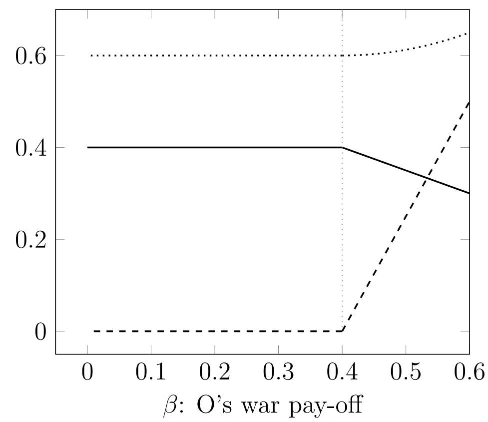

Results hold $\alpha = .4$. β is constrained by $0 < \alpha + \beta < 1$. Dashed: pr(War) $1 - \psi^*$. Dotted: O's expected utility $EU^O(\hat{q}^*)$. Solid: Optimal offer $\hat{q}^*$. Vertical dotted line marks where concentrated offer risks war.

Figure 4: Comparative Statics

By realizing heterogeneous preferences it is possible that states increase the amount of value they can extract from the pie beyond what Fearon originally envisioned. This is only possible if states can coordinate their beliefs about which concessions are the most valuable. Thus, in a world with persuasive diplomacy, both players do better because each receives their preferred concessions. But this increased expected utility comes with increased risk. Diplomacy induces the Offerer to confront uncertainty about the relative value between different issues that it otherwise would not have if it had made offers over equal parts of the pie. When O's cost of war is low enough, O trades-off the size of the concession against the possibility that the offer will not be large enough and war will ensue.

O could always avoid war with a balanced offer. But both players expected benefit is higher when war is possible. O benefits because the concentrated offer allows O to extract the additional surplus created by heterogeneous preferences. S benefits because O is forced to take a risk. O cannot low-ball all types of S, and so the types who value one issue much more than the other do better than they would have if they had received a balanced offer. Of course, no types of S can do worse because war is always an option.

One may question these inferences for two reasons. First, the counter-factual result follows because O was very uncertain about S's favorite issue. Yet in realistic settings O may have more information and choose a concentrated offer even absent diplomacy. Second, I only consider two issues. In the real world, states bargain over multiple issues simultaneously. Intuitively, increased uncertainty and complexity should make signaling more difficult. Thus, it is unclear if my results hold in more complicated bargaining settings. I now turn to these issues.

2.2 Non-Uniform Priors

In many real-world settings O has strong priors about what issues S prioritizes based on O's understanding of S's history, culture, defense posture and commercial ties. For example in 1969, before Kissinger's visit to China, the CIA assessed that the "primary objectives of the present regime in Peking include treatment as a major world power and as a primary source of revolutionary leadership; accommodation of its policies by other Asian states; and control of Taiwan."[13] If a crisis had erupted before Kissinger went to China it is likely that the United States would have used this assessment to concentrate their offer on these issues even absent pre-crisis diplomacy. That negotiating strategy would have entailed certain risks. After all, the CIA's assessment underestimated China's interests in both controlling Tibet and recognition as a nuclear power. These interests were explained to Nixon during his 1971 visit. The CIA's partially correct assessment could have led to a concentrated but insufficient offer. But following Nixon's visit, the US was more likely to tailor their offer to China's core interests raising the chance of a peaceful settlement.

In this section, I alter my baseline model to understand the effects of non-uniform priors about S's preferences. The analysis clarifies how my theory fits with traditional theories that find diplomacy induces peace (cf Trager 2011). I find that when O has strong priors about

[13] Summary of the CIA Response to NSSM 14. National Archives, RG 59, S/S Files: Lot 80 D 212, NSSM 14. Date still classified.

which issue S values the most, he prefers to make a concentrated offer over a balanced offer even when diplomacy is absent. In these conditions, diplomacy decreases the risk of war. Yet when O has non-uniform priors that only weakly suggest one issue is more valuable than the other, or the cost of war is sufficiently high, O offers a balanced offer even when absent diplomacy. In these cases diplomacy induces risk-taking.

To model non-uniform priors, I use a different method for allocating values to issues. In the main model I drew two values θ_1, θ_2 that corresponded to two issues n_1, n_2. In this game, I draw two values θ_y, θ_z i.i.d from a common Θ. The subscripts y, z emphasize that these draws do not automatically correspond to the two issues n_1, n_2. Rather, I assume that the highest draw ($\max\{\theta_y, \theta_z\}$) is assigned to n_1 with probability ξ and the lowest draw is assigned to n_1 with probability $1-\xi$. The remaining draw is assigned to n_2. Without loss of generality, I assume that S is more likely to value n_1 over n_2 by setting $1/2 < \xi < 1$.[14] I assume ξ is common knowledge to both players but the two draws, θ_y, θ_z and their allocation to n_1, n_2 are only observed by S. Otherwise, I hold the baseline model the same. These assumptions capture the strategic dynamic of strong priors. O believes that S probably values one issue more than the other. However, O is uncertain if his assessment is correct. Even if it is, O does not know exactly how much more S values one issue over the other.

Proposition 2.3 *Suppose a counter-factual setting where diplomacy is impossible. In equilibrium, S makes a concentrated offer if* $\xi > \frac{(1-\alpha-\beta)8\alpha}{(2-2\beta-\alpha)^2}$ *and makes a balanced offer otherwise. All equilibrium concentrated offers are of length* $t^* = \frac{2-2\beta+\alpha}{4}$ *(same as prop. 2.2), and accept a positive probability of war:* $(1-\xi)\frac{2-2\beta-\alpha}{2\alpha}$ *(differing from prop. 2.2 only by a factor* $(1-\xi)$*).*

The result follows almost immediately from the proof of proposition 2.3. To see it, notice two facts. First, in the game with non-uniform priors and no diplomacy, O's risk return trade-off from a concentrated offer is:

$$EU^O(\hat{q}|\text{diplomacy is impossible}) : (1-t)\psi_t\xi + (1-\psi_t\xi)\beta \equiv (1-t-\beta)\psi_t\xi + \beta \qquad (2)$$

The only difference between equations 1 and 2 is that $\omega = 1/2$ is replaced by $\xi > 1/2$.

[14] The main model is equivalent to the case where $\xi = 1/2$.

Second, O sets the length of t^* based on his beliefs about how much S values her favorite issue relative to her least favorite issue. Thus, once O decides to make a concentrated offer against one issue, O sets t^* based solely on O's beliefs about ψ. Since O's beliefs about ψ are the same in both games, O makes an offer of the same length in both games. In the game with non-uniform priors, O must accept some additional risk that S will reject the offer because O may not have correctly selected S's favorite issue $(1-\xi)$. Thus, the length of the offer is the same, but the risk of war is higher.

In contrast, non-uniform priors have no effect on the behavior described in the credible cheap talk equilibrium:

Lemma 2.4 *In the game with non-uniform priors, the credible cheap-talk equilibrium described in proposition 2.2 is an equilibrium. It does not depend on O's prior (ξ).*

Lemma 2.4 follows immediately from proposition 2.2. S faces the same incentives to signal honestly. Once S signals her preferred issue, O updates his beliefs such that $\omega = 1$ in Utility Equation 1.

Figure 5 overlays equilibrium behavior for the games with uniform and non-uniform priors. Consistent with Figure 4, Figure 5 holds S's war pay-off constant at $\alpha = 0.4$. The y-axis plots ξ: O's prior beliefs that S values n_1 more than n_2. The x-axis plots O's war pay-off (high values imply the cost of war is low). The plots identifies three possible equilibrium marked **A**, **B** and **C**.

In area **A**, O's cost of war is so high that O does not risk war even when S uses diplomacy to reveal her preferences.[15] To the right of this area, O's cost of war is sufficiently low that he is willing to gamble if his priors about S's favorite issue are strong enough. In area **B** O will offer a gamble $t^* < \alpha$ only following a diplomatic message. In this range, diplomacy increases the probability of war by inducing a risk-taking strategy. In area **C**, O will gamble even if S does not signal her intentions. In this range, diplomacy decreases the probability of

[15]In the games where signaling is allowed, **A** marks the region where O makes a concentrated offer with no risk war. In the games where signaling is omitted, O offers a balanced offer in this region.

war because O's gamble is now always made against S's favorite issue. The plot shows that O is only willing to gamble absent diplomacy if O is confident that he knows S's preferred issue and the cost of war is low.

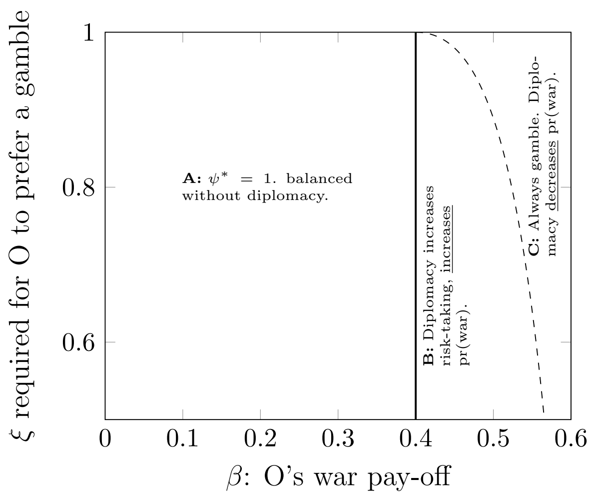

Plots equilibrium strategy for games with non-uniform and uniform priors. Assumes $\alpha = .4$. β is constrained by $0 < \alpha + \beta < 1$.

Figure 5: Equilibrium comparison: uniform and non-uniform priors

Implication 1 *When a state has strong priors about its rival's preference order diplomacy decreases the chance of war. When a state is very uncertain about it's rival's preference order, diplomacy induces a concentrated offer with a higher chance of war.*

This substantive implication is troubling. I hoped that diplomacy reduced the risk of war in cases where states had the least information about each other. Instead, diplomacy reduced the prospect of war only when states have quite a lot of information about each other. Based on these results, I expect that states with new foreign policies (say right after a regime change) are most likely to experience war induced by diplomacy. Whereas well established regimes with long-standing foreign policy agendas are likely to reduce the risk of war through diplomacy. This tracks well with the motivating example. China had only recently developed a foreign policy agenda before the Sino-Soviet dispute. Thus, the Soviets had very little information about what Mao would wan. By the mid-2000s, American policy-makers had developed strong priors about China's core interests.

2.3 The Consequences of Increasing Complexity

Sometimes states only bargain over a few issues at a time ($J = 2$), but other times bargaining is more complex. For example, China and the United States bargain over territorial control in Taiwan, the Scarborough Shoals, influence in Asia, currency and trade norms, institutional design, human rights, etc ($J > 2$). What effect does diplomacy have when states contest multiple issues and bargain over these issues simultaneously? I model increased complexity by increasing the number of issues states bargain over and randomly assigning values to them (draw J θs from Θ).

Intuitively, increased complexity should increase the amount of uncertainty that O has about S's preferences which should make gambling less attractive. When diplomatic signaling is impossible this is the case. However, I observe the opposite once I include diplomacy. Specifically, suppose S sends a message that distinguished between the group of her most valuable issues and the remaining less valuable issues. I'll show that even if S provides no information about the relative value between these groups of issues, O still learns a considerable amount about S's relative value.

The intuition is similar to models studied by Jackson and Sonnenschein (2007) and Jackson, Sonnenschein, and Xing (2015) who consider repeated negotiation models with pre-play communication.[16] When many issues are drawn from a common distribution, O learns much more about how much S values high valued issues relative to low-valued issues *on average*. As the number of issues increases, O exploits additional information about the relative average values for these issue sets to make smaller offers and extract more of the surplus. Jackson et al. (2015) assume repeated bargains between market actors that re-negotiate whether bargains succeed or fail. This creates an opportunity to learn from failure but does not well capture the problem of major war. Further, they focus on how the surplus is shared given different offering strategies and do not analyze differences in the risk of failure. I analyze a

[16] Similarly, pricing models find that monopolists more accurately assess buyers' valuations when they analyze groups not individual buyers (Adams and Yellen 1976; McAfee, McMillan, and Whinston 1989).

counter-factual case where communication is absent to emphasize how diplomacy affects O's offering strategy and the risk of war.

In this section I describe the equilibrium intuitively and visualize results. I show that as the number of issues being bargained over increases, O makes an equilibrium offer that risks war under more and more conditions. However, the hazard of war in each case is reduced.

Before I begin, I re-write the notation from the basic model to fit the $J >> 2$ issue case. I now assume O and S bargain over n_j issues where S's value for each issue is θ_j drawn i.i.d from a common Θ. An offer consists of J distinct offers $q\{q_j\}$. S's utility from exception this offer is $\frac{1}{J}\sum_{i=1}^{J} q_i\theta_i$. O's utility from having an offer accepted is still $1-t$ but now $t = \frac{1}{J}\sum_{i=1}^{J} q_i$.

The balanced offer easily translates to the $J > 2$ case. O makes an offer such that every $q_j = t/J$. It is more difficult to generalize O's concentrated offer to the multi-issue case. The reason is that O's concentrated offer will depend on the information contained in S's message (see Jackson et al. 2015, for discussion). Here I focus on a specific combination of message and concentrated offer. The original messaging strategy $m(A)$ had the effect of identifying S's most valuable issue. I generalize that idea to the multi-issue case as a message that identifies S's L most valuable issues (where L is a positive integer). To do so, divide up the set of J issues into two sub-sets $A\{n_l\}$ containing exactly L issues and $B\{n_b\}$ containing $J-L$ issues. The issues in $A\{n_l\}$ are more valuable because for any n_l in A and n_b in B, $\theta_l \geq \theta_b$. A messaging strategy $m(A\{n_l\})$ is one were S identifies a set of L most valuable issues. Like the two-issue game, this strategy has the effect of distinguishing between S's L most valuable issues and $J-L$ least valuable issues but providing no information about the relative values between different issues.[17]

[17] More specifically, partition the type-space Θ^J into sub-sets of types where each subset shares a common $A\{n_l\}$. Map the feasible set of messages $\Theta^J, \emptyset$ one-to-one and onto these subsets such that each subset is assigned at least 1 message. In $m(A\{n_l\})$, each subset mixes over their assigned messages. There are thus no-off path messages and every type can distinguish their L most valuable issues.

I re-define O's concentrated offer as one where O identifies L issues $L\{n_l\}$ to make offers against and $J-L$ issues to discard. Critically, O's concentrated offer targets the same number of issues as S's message. Against the L issues, O offers $\hat{q}_l = t/L$ of each issue and nothing of the remaining $J-L$ issues.[18] S's utility from this offer is $\frac{1}{J}\sum_{l=1}^{L} q_l\theta_l$.

Other messaging strategies can trigger different equilibrium offers. For example, an equilibrium exists where S sends a message that is a precise ordering of its issues. O offers some number of these issues in full, then makes an offer with a risk return trade offer over only 1 issue. I chose my combination of messages and offers for three reasons. First, the message $m(A\{n_l\})$ most closely reflect the two issue message $m(A)$ allowing for a closer comparison. That is, the message distinguishes between high and low valued issues but provides no further information. Second, the message closely reflects how China describes its core interests. China claims to value a number of issues and territories as vital to its national interest but refuses to specific which ones are most important.

Third, although there are many messaging strategies to consider, they must all be consistent with the main conclusion I draw from the equilibrium I focus on: credible messages do not reduce the risk of war and may raise it. The reason is that O makes a balanced offer when diplomacy is prohibited with no risk of war. To the extent that some of these messages can force O to confront a risk return trade-off, it must be that some of them lead to an additional risk of war. As a result, it is interesting that even one combination of messages and offers can lead to an equilibrium that raises the risk of war.

In Appendix B.1 I complete a similar analysis to the two-issue game. I sketch out the equilibrium to extend Propositions 2.1 and 2.2 to the multi-issue case. As with before, O faces the same three-part optimization problem. O identifies his best balanced ($\tilde{q}^*$) and concentrated ($\hat{q}^*$) offer, then chooses the strategy that maximizes his utility.

In both cases, O's best balanced offer directly extends from the basic game such that $\tilde{q}^*_j = \alpha/J$. The total length of this offer $\tilde{t}^* = \alpha$ O's utility is $1-\alpha$, S's utility is $\frac{1}{J}\sum_{i=1}^{i=J}\tilde{q}^*_j\theta_j =$

[18]For simplicity, I restrict my attention to the case where $\alpha < 1/2$ and $L \geq J/2$.

α. I re-define O's utility from a concentrated offer as:

$$EU^O(\hat{q}) : (1-t)\psi + (1-\psi)\beta. \tag{3}$$

O's decision to increase t by ϵ implies O increases the amount offered against each issue by ϵ/L. Since O offers equal parts of all issues in $A\{n_l\}$, ψ is the probability that the mean value $\mu(n_l) = \sum_1^L \theta_l/L$ is x times larger than the mean of the discarded intervals: $\psi = pr(\mu(n_l) \geq x\mu(n_b))$. Critically, this implies that the precise value of each issue n_l does not matter. All that matters is the mean of $A\{n_l\}$ ($\mu(n_l)$) relative to the mean of $B\{n_b\}$.[19]

First, I consider the case where signaling is impossible and O's offering rule relies on his prior beliefs: $\gamma^\dagger(\Theta^J) = \gamma(\Theta^J)$. Under this restriction, I consider a set of randomly chosen issues $\bar{Z}$ that O makes a concentrated offer against and call the remaining issues $\underline{Z}$. The expected mean value of $\bar{Z}$ issues $\mu(\bar{Z})$ is defined via the Bates Distribution over $[0,1]$ with $\bar{Z}$ draws.[20] This PDF is continuous and integrable. By the independence of draws, I similarly characterize $\mu(\underline{Z})$. $\psi_Z = \psi|\gamma^\dagger(\Theta^J)$ is the ratio of these two distributions.

I present the simulated distributions for $\mu(\bar{Z}), \mu(\underline{Z})$ and ψ_Z in Figure 6. The simulated results hold $\bar{Z}$ constant as 1/4 of the total number of issues. The lighter distributions reflect 8 total issues ($\bar{Z} = 2$), and the darker distribution reflect 100 total issues ($\bar{Z} = 25$) under dispute. The first row depicts the distribution of mean values for the $\bar{Z}$ issues selected. The second row depicts the mean of the $\underline{Z}$ issues discarded, and the bottom row depicts the distribution of ratio of the means defined by ψ_Z.

The figure illustrates the challenges O faces in making an offer when there are many issues and no communication. Consider O's best response as J approaches infinity. In the limit, the variance surrounding these distributions tends towards 0. O grows increasingly certain that the average value of issues in the set he offers against and the discarded sets is very close. However, O never learns which one is larger than the other on average. In finite

[19]The specific distribution ψ is conditional on $\gamma^\dagger(.)$. I distinguish between these as they become relevant.

[20]See Appendix B.2 for more information about the Bates Distribution.

samples one set of issues will be larger than the other (i.e. the probability the means are equal is 0). Thus, O learns that the mean of both groups are close, but cannot determine which group is larger. Thus increasing J amplifies the effect that $\omega = 1/2$ had in the main model.

Next, I consider how a credible message $m(A\{n_a\})$ of a fixed number of concessions L, effects ψ. I define $\psi_M = \psi|\gamma^{\dagger}(\Theta^J, m(A\{n_a\}), \hat{q})$. An equilibrium exists where S observes her type then sends a message $m(A\{n_a\})$. O updates beliefs $\gamma^{\dagger}(m(A\{n_l\}), \Theta^J) : pr(\hat{\theta}_l > \hat{\theta}_b) = 1 \forall n_l, n_b$, and offers equal parts of each issue in $A\{n_a\}$. This can carry a risk of war: $\psi_M^* < 1$.

Notice that O knows that all θ_l are larger than the remaining θ_b. Further, O knows there are L issues in $A\{n_l\}$, and $J - L$ issues remaining. Although O does not know the precise value for any of the L issues in $A\{n_l\}$, he knows that each is more valuable than the discarded intervals.

I cannot define the distribution ψ_M explicitly for an arbitrary J. I assume that it is continuous with a unique maximum value, positive on $[1, \infty)$. I also assume the distribution tends to 0 as x approaches infinity.[21] In Appendix B.3 I analyze results in the limit by considering the function that is achieved if J reaches infinity. I define w (for width) as a fixed ratio L/J to show the following.

Lemma 2.5 *Consider the equilibrium that follows from a signal $m(A\{n_l\})$, for a fixed $\alpha, \beta, \gamma^{\dagger}(\Theta^J, m(A\{n_l\}))$, L defined as a fixed proportion of $J : L/J = w$, and a parameter J. As J approaches infinity, $\hat{t}_l^*$ approaches $\frac{\alpha(w-1)^2}{2-w}$ from above, $EU^S(\hat{q}_l^*)$ approaches $EU^S(War)$ from above, $EU^O(\hat{q}_l^*)$ approaches $1 - \frac{\alpha(w-1)^2}{2-w}$ from below.*

The logic relies on two points. O's ideal offer trades off the risk of war against the size of the offer. Thus, O's decision to set $\hat{t}^*$ does not depend on how much S will profit, but the number of types that accept a particular $\hat{q}^*$. Second, like the counter-factual case without signaling, as J increases, the variance around the distribution of ψ_M approaches 0. O becomes more confident he understands what S's average value $\mu(n_l)$ is relative to $\mu(n_b)$.

[21]Simulations for large J and analytical solutions for small J are consistent with these conjectures.

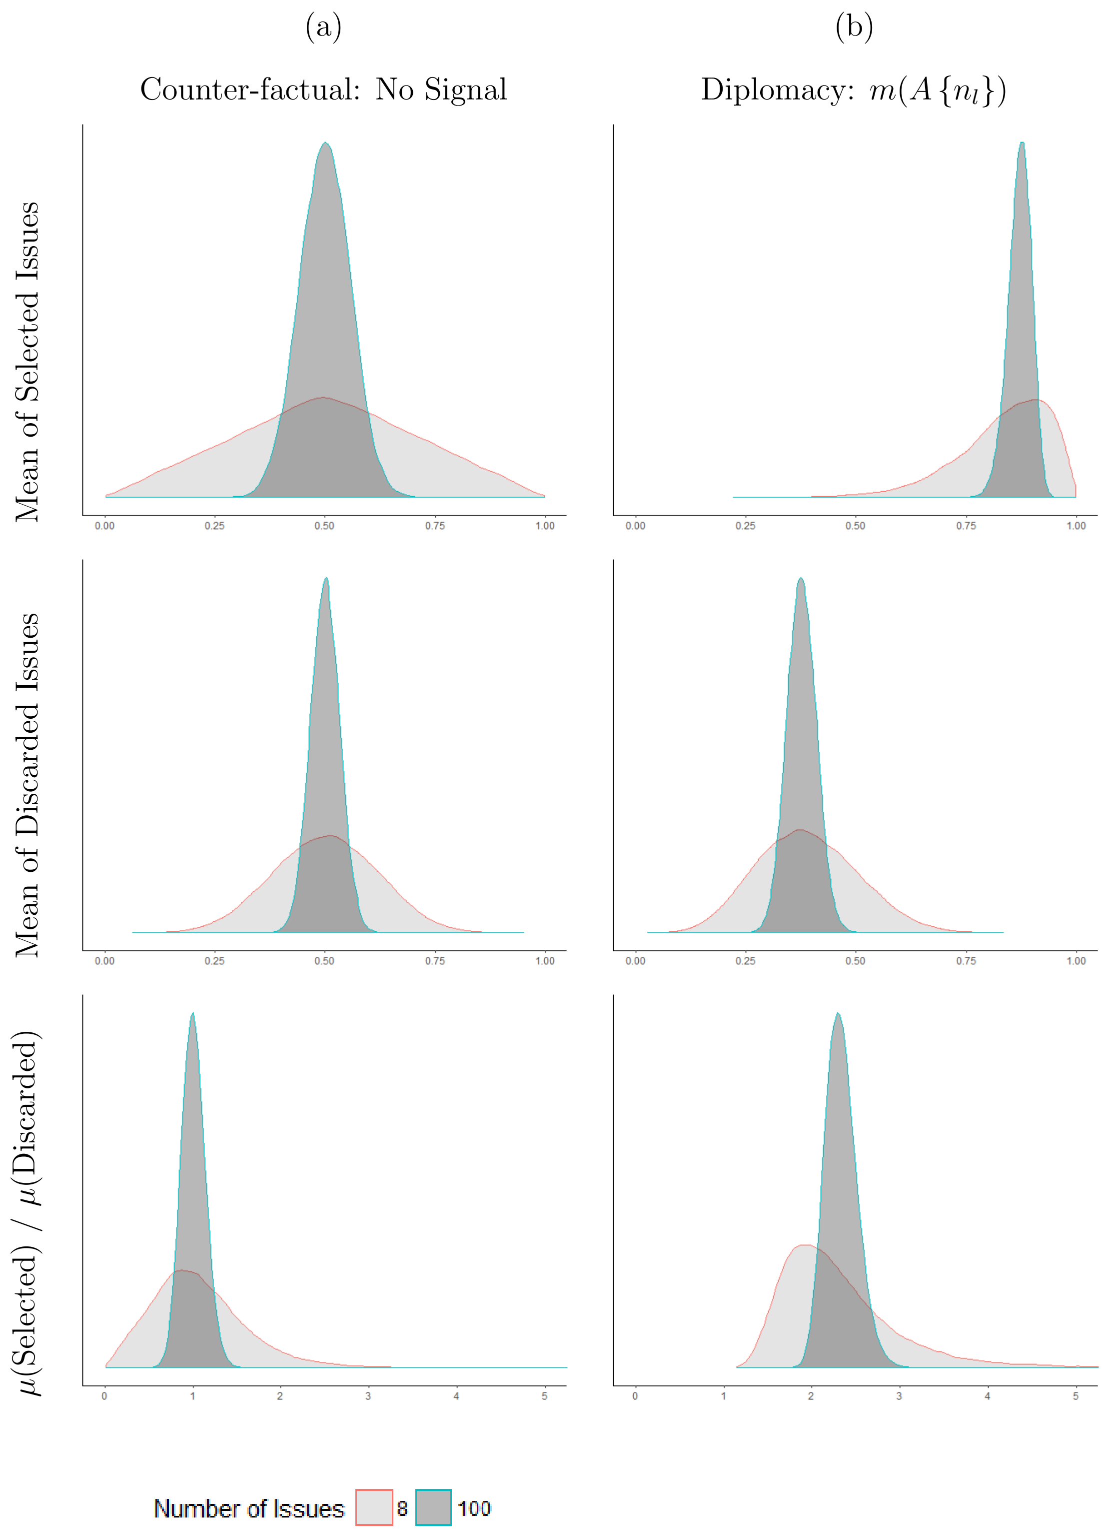

Assumes 1/4 of total issues are selected and 3/4 are discarded. Ratio constrained between 0 and 5. Results of simulations with 500,000 observations.

Figure 6: Expected Values of Collections of Issues with Different Information

I visually represent the effect of a message $m(A\{n_l\})$ on O's beliefs in the second column of Figure 6. The rows depict the simulated distributions for $\mu(n_l), \mu(n_b)$ and ψ_M respectively. Unlike, the counter-factual case, $\mu(n_l)$ does not converge to $\mu(n_b)$. Rather, $\mu(n_l)/\mu(n_b)$ converges to a number larger than 1: $\frac{\alpha(w-1)^2}{2-w}$. Thus, as J approaches infinity, O grows extremely confident that $\mu(n_l)$ is exactly that much larger than $\mu(n_b)$.

This increased certainty further implies:

Lemma 2.6 *Suppose for a fixed β, α, and fixed ratio w, $EU^O(\hat{q}|\psi_M = 1) > EU^O(\hat{q}|\psi^* < 1)$. There must be a $J' > J$ such that $EU^O(\hat{q}|\psi^* < 1) > EU^O(\hat{q}|\psi^* = 1)$ for all $J'' > J'$*

The proof is in Appendix B.4. In the two interval case I saw that when O's cost of war is very high, O sometimes prefers to make a concentrated offer that did not risk war. Lemma 2.6 implies that for costs of war so high that O is unwilling to risk war if he is bargaining over just two issues, as the bargaining space grows in complexity (by including additional issues), O eventually selects an equilibrium offer that accept a positive risk of war.

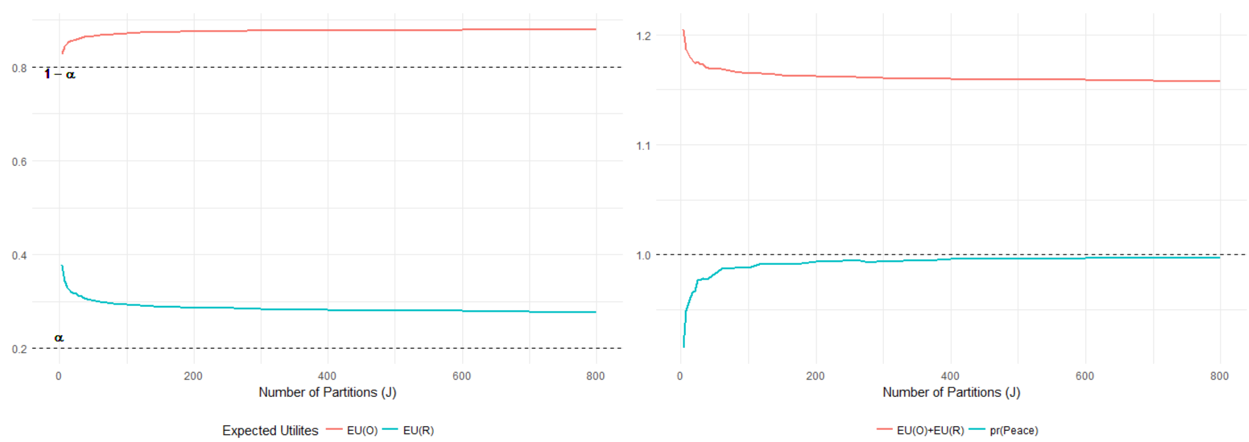

Figure 7: Comparative Statics

Figure 7 plots equilibrium results from a concentrated offer as bargaining grows in complexity.[22] Two features of these results are surprising. First, in cases where O risks war, the hazard of war attenuates as bargaining grows in complexity (but always remains positive).

[22]Since I cannot find a closed form analytical solution for ψ_M, the results are simulated at different J.

The reason is that O gains more information about the relative difference between the average value of issues $\mu(n_l)$ verses the average value $\mu(n_b)$. This implies, that O can be more confident that smaller offers will satisfy more types. O takes advantage of this by making smaller offers. While the hazard of war attenuates in cases where O is willing to risk war, the number of cases where O is willing to risk war increases for the exact same reason. Thus, war becomes possible under more conditions, but less likely in each particular condition.

Second, there are clear winners and losers from increased complexity. Specifically, O can extract more of the surplus because he is more confident that he understands the relative value between $A\{n_l\}$ and $B\{n_b\}$. S's expected utility diminishes as the offer approaches the bare minimum that each type will accept.

Implication 2 *States receiving concessions do better in expectation either when they bargain over a few issues at a time; or when there is more uncertainty about the relative value between issues.*

The implication is consistent with variation in Chinese diplomatic strategy and war onset. As discussed, modern Sino-American diplomacy is extremely complex because there are multiple issues to consider. China has actively concealed information about its preferences and resists emphatically any attempts at a grand bargain that links multiple issues together in a single negotiation. The United States pushes for greater transparency and more precise diplomacy. The United States makes explicit that ambiguity risks conflict. Despite much rhetoric about the heightened risk that uncertainty brings, war has not materialized. This can be contrasted with Sino-Soviet diplomacy in the 1960s. Mao identified three issues under dispute, and made explicit his priority order in that set. Khrushchev's offer matched Mao's priority order. Yet it still wasn't enough. This very simple bargain ended in war.

3 Discussion

In the real world, states dispute a variety of issues but care about some more than others. This creates information problems based on uncertainty about: a rival's preferred issue; and how much a rival values one issue over another. Consistent with the existing literature, I showed that states can use diplomacy to reveal their preferred issues leading to more coordination and Pareto-improving offers. Inconsistent with the existing literature, I showed that when diplomacy overcame only risk about a rival's preferred issue, it sometimes enticed the Offerer to accept additional risk when managing uncertainty about relative values between the two issues. As a result, diplomacy induced a different kind of offer that raised the risk of war.

Extensions showed that the risk inducing effects of diplomacy arose in settings where states had the least information about each other's preferences. Further, in complex bargains where states bargained over multiple issues, diplomacy led to more conditions where states accepted some risk of war (although the amount of risk they accepted was less). These results helped clarify conditions under which diplomacy reduced or increased the risk of war. The answer hinges on whether states accept risk when diplomacy is absent. In cases where an Offerer prefers to confront a risk return trade-off even when communication is prohibited, diplomacy eased that risk. But sometimes when communication is prohibited the state making offers avoided the risk return trade-off altogether at a cost to efficiency. In this cases, diplomacy enticed the Offerers to accept additional risk, thus raising the risk of war.

The results also showed that when states used diplomacy to reveal their intentions they did best when they provided limited information. Although the number of issues is exogenously set in these games, I conjecture that the state receiving offers will provide as little information about the scope of her preferences as possible. The receiver's incentives to hide relative values between two issues is consistent with a common bargaining strategy in world politics. Belligerents regularly claim that their foreign policy demands are driven by broader goals (e.g. nationalism, security, status, etc.) and these broad goals imply several valuable concessions. But they then refuse to specify details, or make grand bargains. This trend extends beyond Sino-American Diplomacy: Hitler appealed to ethnic nationalism before making demands; the United States in the 1890s used the Monroe Doctrine to stake its claim over the entire Western Hemisphere; and Stalin used security to motivate his demands during conferences at Yalta and Moscow. All these broad appeals to values have four things in common. First, they identify a set of core interests that they claim are much more valuable than peripheral interests. Second, they imply that a whole range of concessions are clearly worthless. Third, by keeping the claims broad, they leave a lot of ambiguity about the relative importance of core interests to each other. Finally, these states all bargained over issues one at a time, rather then settling all their claims simultaneously. My theory suggests that states may be creating ambiguity about the preference order within the set of core interests to extract larger bargains at some risk of conflict.

The results imply that diplomats should divide their time between concealing the precise order of their preferences and discovering the order of their counterparts' interests. This stands in contrast to the conventional wisdom about signaling. Scholars expect that states with private information try desperately to communication credibly and fail because their adversaries do not believe them. Our results suggest that states with no information want their adversaries to reveal private information cheaply because they will believe it. This dynamic played out during the recent peace negotiations between the FARC and the Colombian Government. The Colombian government tried to compel the FARC to limit their demands to just six items to be negotiated. Instead the FARC "has been endlessly expanding the agenda and bringing myriad proposals and demands to the table," in an attempt to create some ambiguity about their specific demands.

The FARC example is just one case that extends beyond inter-state relations. This dynamic applies broadly to any bargaining scenario with heterogeneous preferences. Future work may consider the model's implications for other policy areas where actors spend time identifying key issues before bargaining takes place. This includes trade negotiations, civil war termination, contract negotiations between firms and plea bargaining in courts.

Acknowledgments

Michael Miller, Alexandre Debs, James Morrow, Anne Sartori, Mark Fey, Brenton Kenkel, Hein Goemans, Richard Jordan, Robert Trager, William Spaniel, James Fearon, Bradley Smith, Navin Bapat, two anonymous reviewers and the Editors at JOP, and participants from both the Watson Center Workshop at the University of Rochester, and the Formal Models of International Relations Conference (2017) provided excellent comments that improved the paper—Thanks!

References

Adams, W. J. and J. L. Yellen (1976, aug). Commodity Bundling and the Burden of Monopoly. *The Quarterly Journal of Economics 90*(3), 475.

Arena, P. and S. Wolford (2012, jun). Arms, Intelligence, and War. *International Studies Quarterly 56*(2), 351–365.

Battaglini, M. (2002, jul). Multiple Referrals and Multidimensional Cheap Talk. *Econometrica 70*(4), 1379–1401.

Bils, P. and W. Spaniel (2016, sep). Slow to Learn. *Journal of Conflict Resolution*.

Brooks, S. (1999, oct). The Globalization of Production and the Changing Benefits of Conquest. *Journal of Conflict Resolution 43*(5), 646–670.

Carr, E. H. (1964). *The Twenty Years' Crisis, 1919-1939.* Harper & Row.

Carter, D. B. and H. E. Goemans (2011, apr). The Making of the Territorial Order: New Borders and the Emergence of Interstate Conflict. *International Organization 65*(02), 275–309.

Chakraborty, A. and R. Harbaugh (2003, mar). Cheap talk comparisons in multi-issue bargaining. *Economics Letters 78*(3), 357–363.

Chakraborty, A. and R. Harbaugh (2010, dec). Persuasion by Cheap Talk. *American Economic Review 100*(5), 2361–2382.

Crawford, V. and J. Sobel (1982). Strategic Information Transmission. *Econometrica: Journal of the Econometric Society 50*, 1431–1451.

Debs, A. and J. C. Weiss (2016, apr). Circumstances, Domestic Audiences, and Reputational Incentives in International Crisis Bargaining. *Journal of Conflict Resolution 60*(3), 403–433.

Fearon, J. D. (1994, sep). Domestic Political Audiences and the Escalation of International Disputes. *American Political Science Review 88*(3), 577–592.

Fearon, J. D. (1995, jun). Rationalist Explanations for War. *International Organization 49*(03), 379.

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

Goemans, H. E. and K. A. Schultz (2013). The Politics of Territorial Disputes: A Geospatial Approach Applied to Africa.

Jackson, M. O. and M. Morelli (2011). The Reasons for Wars: An Updated Survey. *The Handbook on the Political Economy of War*, 34–57.

Jackson, M. O., L. K. Simon, J. M. Swinkels, and W. R. Zame (2002). Communication and Equilibrium in Discontinuous Games of Incomplete Information.

Jackson, M. O., H. Sonnenschein, and Y. Xing (2015, mar). The Efficiency of Bargaining with Many Items. *SSRN Electronic Journal*.

Jackson, M. O. and H. F. Sonnenschein (2007, jan). Overcoming Incentive Constraints by Linking Decisions. *Econometrica 75*(1), 241–257.

Jervis, R. (1978, jan). Cooperation Under the Security Dilemma. *World Politics 30*(02), 167–214.

Keohane, R. O. R. O. (2005). *After hegemony : cooperation and discord in the world political economy*. Princeton University Press.

Kim, K. and P. Kircher (2015, sep). Efficient Competition Through Cheap Talk: The Case of Competing Auctions. *Econometrica 83*(5), 1849–1875.

Kurizaki, S. (2007, aug). Efficient Secrecy: Public versus Private Threats in Crisis Diplomacy. *American Political Science Review 101*(03), 543–558.

Kurizaki, S. and T. Whang (2015, jul). Detecting Audience Costs in International Disputes. *International Organization 69*(04), 949–980.

Kydd, A. H. (2005). *Trust and Mistrust in International Relations*. Princeton, N.J.; Woodstock: Princeton University Press.

Leventŏlu, B. and A. Tarar (2008, sep). Does Private Information Lead to Delay or War in Crisis Bargaining? *International Studies Quarterly 52*(3), 533–553.

McAfee, R. P., J. McMillan, and M. D. Whinston (1989, may). Multiproduct Monopoly, Commodity Bundling, and Correlation of Values. *The Quarterly Journal of Economics 104*(2), 371.

Mearsheimer, J. J. (2001). *The Tragedy of Great Power Politics*. New York: Norton.

Moravcsik, A. (1998, oct). Taking Preferences Seriously: A Liberal Theory of International Politics: Erratum. *International Organization 52*(4), 229.

Powell, R. (1999). *In the Shadow of Power: States and Strategies in International Politics*. Princeton, N.J: Princeton University Press.

Powell, R. (2006, jan). War as a Commitment Problem. *International Organization 60*(01).

Sartori, A. E. (2002, feb). The Might of the Pen: A Reputational Theory of Communication in International Disputes. *International Organization 56*(1), 121–149.

Spaniel, W. and P. Bils (2017). Policy Bargaining and International Conflict. *Journal of Theoretical Politics*.

Trager, R. F. (2011, jul). Multidimensional Diplomacy. *International Organization 65*(03), 469–506.

Trager, R. F. (2013, nov). How the scope of a demand conveys resolve. *International Theory 5*(03), 414–445.

Waltz, K. N. (1979). *Theory of international politics*. Reading, Mass: McGraw-Hill.

Biographical Statement

Michael F. Joseph is the Henry Chauncey Jr. Postdoctoral Fellow at the International Security Studies Center at Yale University, New Haven, CT 06520.

Online Appendix for A Little Bit of Cheap Talk is a Dangerous Thing

A Appendix for Main Model

A.1 Proposition 2.1: Counter-factual Case

We'll now show that when signaling is impossible O never makes a concentrated offer. Since S cannot signal, her only decision is whether or not to accept $q(q_1, q_2)$. Recall, S accepts any offer $(\theta_1 - \alpha)q_1 + (\theta_2 - \alpha)q_2 \geq 0$. Since a balanced offer is not type dependent, and a balanced offer sets $q_1 = q_2$, and O values all concessions, then O's best balanced offer must leave S exactly indifferent is $q_1 = q_2 = \alpha$ leaving $U^O(\tilde{q}^*) : 1 - \alpha$.

O's expected utility from a concentrated offer is conditional on S's type and the length of the offer $\hat{q}$.

$$EU^O(\hat{q}) : (1 - \hat{t})\psi\omega + (1 - \psi\omega)\beta \equiv (1 - \hat{t} - \beta)\psi\omega + \beta \tag{4}$$

In this equation, the first term is O's payoff from making an offer of length $\hat{t}$ against a single issue, multiplied by the proportion of types that an offer of length $\hat{t}$ will satisfy (ψ). ω is the probability that O selected the most valuable issue to offer against. The second term is the probability that O either did not select the largest issue to offer against, or that the offer was not large enough given the ratio between the two issues $(1 - \psi\omega)$ and multiplied by O's war payoff (β).

When communication is prohibited, O's prior beliefs hold and $\omega = 1/2$. O's expected utility from a concentrated offer is thus: $EU^O(\hat{q}|\omega = 1/2) : \frac{(2-2\beta-\alpha)^2}{16\alpha} + \beta$. O only prefers a concentrated offer when: $EU^O(\hat{q}^*|\omega = 1/2) > U^O(\tilde{q}^*) \implies \frac{(2-2\beta-\alpha)^2}{16\alpha} + \beta > 1 - \alpha$. The root of this equation is only achieved when $\beta = 1$. Thus, O never makes an offer that accepts $\omega = 1/2$. It follows that O must make a balanced offer.

A.2 Proposition 2.2: Diplomacy Induces Risk-taking

I begin by optimizing O's expected utility from an offering strategy $\hat{q}^*$. For now, I'll assume the message is credible. Later, I'll show that S cannot profitably deviate from it. Let x be the ratio between the two issues such that $\theta_A = \theta_B x$. The probability that the most valuable interval θ_A is at least x times larger than θ_B is:

$$\psi_x : pr(\theta_A \geq \theta_B x) = \frac{1}{x} \tag{5}$$

Subscript ψ_x denotes the specific distribution of risk of war based on a message $m(A)$ and offering strategy $\hat{q}, \hat{t}$.

Next I compute the length of the offer $\hat{q}$ that S will accept:

$$x\hat{t} > \frac{\alpha(x+1)}{2}. \tag{6}$$

In this equation, I include subscript t_x to denote an offer of a particular length that will satisfy all types who satisfy the inequality in Equation 5 with respect to a particular ratio: x. Notice that Equation 6 is written in terms of the ratio between θ_A, θ_B but does not require O to know S's minimum demand. I solve for $x > \frac{\alpha}{2t_x - \alpha}$, and replace x in the optimal ψ_x defined in Equation 5: $\psi_x = \frac{2t_x - \alpha}{\alpha}$. Now sub ψ_x into $EU(O|q^*)$:

$$EU(O|\hat{q}, t_x) : (1 - t_x - \beta)\frac{2t_x - \alpha}{\alpha}\omega + \beta. \tag{7}$$

Since the signal about the largest interval is credible in this equilibrium, $\omega = 1$. However, I leave it in the equation to illustrate that it will not effect the size of O's offer. Since ψ_x is a probability, bounded by 0 and 1, I define a <u>constraint</u> on the corresponding values for q: $0 < \frac{2t_x - \alpha}{\alpha} \leq 1$.

To optimize his expected utility, O must determine the length of q_x, t_x that best balances the probability that the offer is too small against the size of the offer. Taking the derivative with respect to the length t_x:

$$\frac{\partial u}{\partial t_x} : \frac{\omega(2 - 4t_x - 2\beta + \alpha)}{\alpha}. \tag{8}$$

Solving for the FOC:

$$t_x^* = \frac{2 - 2\beta + \alpha}{4}. \tag{9}$$

Above I defined two critical factors in terms of $\hat{q}$ and x. In Equation 5 I defined ψ_x—the probability that S will accept a concentrated offer of a length t_x. I then defined a constraint in terms of $\hat{q}$ as the conditions under which S sets $\psi_x < 1$ and accepts risk of war. I can now redefine these parameters given the optimal t_x^*. Plugging t_x^* into the constraint, O sets $\psi_x^* < 1$ when: $\frac{2-3\alpha}{2} < \beta$. When this condition is met, O's best concentrated offer accepts a risk of war. I derive O's optimal risk of war, ψ_x^*, by plugging the value for t_x^* derived in Equation 9 into the inequality defined in Equation 6, and then solving for x. Per Equation 5, the reciprocal of this defines the risk of war O accepts in equilibrium:

$$\psi_x^* = \frac{2-2\beta-\alpha}{2\alpha} \leq 1. \tag{10}$$

So long as the constraint is met O sets $\psi_x^* < 1$. O's expected utility from the optimal gamble is:

$$EU(O|\hat{q}^*, \psi_x^* < 1) : \frac{(2-2\beta-\alpha)^2\omega}{8\alpha} + \beta. \tag{11}$$

When the constraint is not met, O sets $\psi_x^* = 1$. O's expected utility from this strategy is:

$$EU(O|\hat{q}^*, \psi_x^* = 1) : (1-\alpha-\beta)\omega + \beta. \tag{12}$$

So long as S complies with messaging strategy $m(A)$ truthfully, O weakly prefers a concentrated offer to a balanced offer. To see it is true, set $\omega = 1$. When $EU(O|\hat{q}^*, \psi_x^* < 1) \geq EU(O|\tilde{q}) : \frac{(2-2\beta-\alpha)^2\omega}{8\alpha} + \beta \geq 1-\alpha$. With algebra, the solution to this inequality is: $\frac{2-3\alpha}{2} \geq \beta$. This is the same condition as the constraint. Thus, under all conditions where a concentrated offer is a gamble that risks war, O strictly prefers a concentrated offer over a balanced offer. Clearly, when the constraint is not met, O is indifferent between making a concentrated offer that does not risk war and a balanced offer.

I'll now show that S will signal honestly. Since all off-path messages are covered, the only deviation I need to consider is that any type $\theta_1 > \theta_2$ sends some message in the set $m(A=2)$. Any such signal must lead to an offer $2\alpha \geq q_2$, $q_1 = 0$ which any type $\theta_1 > \theta_2$

rejects in favor of war. But war is S's min-max payoff. Thus, S must be at least indifferent between sending this off-path message. An analogous argument holds for all types $\theta_1 \leq \theta_2$. It follows that no types can profitably deviate from off-path messages. We've shown that O prefers a concentrated offer to any other strategy so long as S plays message $m(A)$ and derived the t_x^*, ψ_x^* reported in the proposition as the values that optimize O's utility from this strategy. We've also shown that S cannot profit from deviating from a message $m(A)$ if she expects a concentrated offer. This completes the proof.

A.2.1 Remarks on alternative messages

In proposition 2.2, I focused on a messaging strategy $m(A)$ which covered all off-path messages. Here I address two questions: Can S send a message in equilibrium that perfectly reveals her type? If I adjust $m(A)$ so that there are off-path messages, does the equilibrium still survive the intuitive criterion?

The messaging strategy $m(A)$ prevents S from communicating information about the relative differences between θ_1, θ_2. It is uncertainty over this ratio that created the risk return trade off O accepts in equilibrium. But if S could use a message that perfectly revealed her type, then diplomacy would not induce conflict because O could exploit that message to offer all types exactly their minimum demand. As a result, it is important to rule out this possibility.

Remark A messaging strategy $m(\hat{\theta}_1 = \theta_1, \hat{\theta}_2 = \theta_2)$ cannot survive in a PBE.

Suppose some type $\theta_1 > \theta_2$ follows this strategy. O processes the signal and offers $\hat{q}_1' = \alpha(\theta_1 + \theta_2)/\theta_1$. This leaves S with her war-payoff: $\alpha(\theta_1 + \theta_2)/2$. Since messages are costless, that type can deviate to a message: $m(\hat{\theta}_1 = \theta_1, \hat{\theta}_2 = \theta_2 + \epsilon)$ at no cost. If the message was credible, O would offer $\hat{q}_1 = \alpha(\theta_1 + \theta_2 + \epsilon)/\theta_1$ such that $U^S(\hat{q}_1) > U^S(\hat{q}_1')$. Thus, all types $\theta_1 \neq \theta_2$ can deviate for profit. It follows that this message cannot survive in equilibrium.

A second issue might be that by restricting all off-path messages, I have focused an equilibrium that would not survive if off-path beliefs were possible. To address this, I turn to a less restrictive version of $m(A)$. Define a messaging strategy $m(A')$ as one where all types $\theta_1 \geq \theta_2$ send a message $m(\hat{\theta}_1 = 1, \hat{\theta}_2 = 0)$ and all types $\theta_1 < \theta_2$ send a message $m(\hat{\theta}_1 = 0, \hat{\theta}_2 = 1)$. This strategy has the same effect as $m(A)$ because all types send a message that identifies their favorite issue but does not distinguish between the relative value of any issues. The core difference is that there are lots of off-path messages.

Remark Replacing the messaging strategy $m(A)$ in proposition 2.2 with messaging strategy $m(A')$ produces an equilibrium with the same behavior. The equilibrium survives the intuitive criterion.

Since the message has the same effect on O's beliefs in equilibrium as $m(A)$ then I need only consider S's incentives to deviate from it. Thus, I turn strait to the two steps of the intuitive criterion. The first step is to consider all types that could profit from deviating. We consider some arbitrary off-path message $m(AA)$. Since messages are costless and $t^* < 1$, all types face the same incentives to deviate to $m(AA)$. Further, so long as war is costly, all types could receive higher offers. Thus, all types could profit from a deviation to $m(AA)$.

Since all types could profit, the second step must consider the full type space. We've seen already that this type space produces a balanced offer that leaves S with her minmax. Thus, no type can profit. Since $m(AA)$ was arbitrary, a messaging strategy $m(A')$ must survive the refinement.

B Appendix for Complex Game

B.1 Equilibrium when J is large

I now write down the equilibrium for the complex case described in Section 2.3. For any J, L define $\psi_M = pr(\mu(n_l) \geq x\mu(n_b))$. We assume that this CDF is continuous supported on $[0, 1]$ and its derivative is fully supported on $[1, \infty)$ with no atoms. We further assume it has a unique maximum and is differentiable in all but finitely many places. We restrict our attention to cases where $L/J > \alpha$.

Proposition B.1 *A cheap-talk equilibrium always exists that satisfies the following properties:*

- *S observes her type and sends message $m(A\{n_l\}|\theta_j)$.*
- *O Processes that signal and updates her beliefs $\gamma^{\dagger}(.) \neq \gamma(.)$.*
- *O offers a concentrated offer $\hat{q}^*$ contained only in the intervals n_l. The offer is of length t: $\alpha \geq t > \alpha L/J$.*
- *S accepts that offer if $\frac{\mu(n_l)J+\mu(n_b)(J-L)}{J}\alpha \leq \mu(n_l)t$ and rejects otherwise.*

We have already shown that for any two issues n_l, n_b, if O's expectation $e(\theta_l) = e(\theta_b) = \Theta$, then O will offer equal parts of each (possibly nothing). Thus, we focus attention on whether O prefers a concentrated offer only in J or prefers to a balanced offer. S will accept an offer of length t for a fixed x when

$$x \geq \frac{(J-L)\alpha}{Jt - \alpha L} \tag{13}$$

We now sub equation 13 into ψ leaving the expected utility for O: $EU^O(t) = (1 - t - \beta)\psi(\frac{(J-L)\alpha}{Jt-\alpha L} + \beta$, with FOC: $\frac{\partial EU^O}{\partial t} = (1 - t - \beta)\psi'(\frac{(J-L)\alpha}{Jt-\alpha L}) - \psi(\frac{(J-L)\alpha}{Jt-\alpha L})$.

It must be that O's optimal t lies between $\alpha \geq t > \frac{\alpha L}{J}$. Consider, that at $\alpha = t$, the right hand side of equation 13 solves at 1. But since the elements in J are at least as large, this is the minimum possible ratio. Thus, setting $t = \alpha$ is the smallest amount that guarantees all types will accept. This implies that at $t > \alpha$, $\psi' = 0$. Thus, O cannot improve a concentrated offer by offering more. Further, O is indifferent between this concentrated offer and a balanced offer because: $(1-\alpha-\beta) * 1 + \beta = 1-\alpha$.

Turing to the leftmost inequality, $t = \frac{\alpha L}{J}$ is the smallest offer that will satisfy the type who draws $\mu(n_b) = 0$. If O offers this, then all types except $\mu(n_b) = 0$ will reject it (we use a strict inequality since at $t = \frac{\alpha L}{J}$ equation 13 is undefined). Since war is inefficient, this is never O's best offer. With t fixed between these values, it is clear that $1 \geq \psi(.) > 0, \psi'(.) > 0$. Therefore, there must be a unique maximum in this range. In the next section we demonstrate that under certain conditions O's optimal $t^* < \alpha$ and therefore O gambles in equilibrium at least sometimes.

Since there are no off-path messages, we need only consider if S can profitably deviate to a different message. Clearly not, the message would lead to an offer of the same total length against a set of issues that was on average less valuable. Thus, S will not deviate.

B.2 Description of the Bates Distribution

The Bates Distribution takes two values: the limits of a closed interval $[a, b]$ and the number of draws n. The variable is commonly written as x and the PDF is:

$$\sum_{i=0}^{i=n}(-1)^i\binom{n}{i}\left(\frac{x-a}{b-a}-i/n\right)^{Z-1} sgn\left(\frac{x-a}{b-a}-i/Z\right).$$

The mean is $\frac{1}{2}(a+b)$, variance $\frac{1}{12n}(b-a)^2$, and skewness 0. This PDF is continuous and integrable. It has spline with one knot for each n and so is differentiable in all but finitely many places.

In the multi-issue game without signaling, if O was to make some offer against $Z \subset \{n_j\}$ issues, the mean of those randomly selected Z issues is defined via the Bates Distribution over $[0, 1]$ with Z draws:

$$PDF = \begin{cases} \sum_{i=0}^{i=Z}(-1)^i\binom{Z}{i}\left(x-i/Z\right)^{Z-1} sgn\left(x-i/Z\right) & \text{if} \quad x \in [0,1] \\ 0 & \text{otherwise} \end{cases} \tag{14}$$

By the independence of draws, I can similarly characterize the remainder $J-Z$ intervals.

B.3 Lemma 2.5: Deriving the Limits.

We now derive the limits presented in Lemma 2.5 by considering the continuous case. We begin by noting that infinite draws on $U[0,1]$ can be re-ordered to $y = 1-x$ and supports S's a total value for controlling all issues: $U^S(q_j = 1 \forall j) : x - x^2/2$. As discussed, L is a proportion of J and can be written as a real number w. The ratio between intervals x converges to $\int_0^w (1-x) / \int_w^1 (1-x)$, equivalent to: $\frac{w(2-w)}{(w-1)^2} > 1$. Since, S's value for all issues is known in the limit, we find the optimal $\hat{q}^*$ using: $\frac{\hat{q}x}{w} = \alpha$. Subbing in x, achieves the value in the Lemma.

B.4 Proposition 2.6: O Accepts Risk of War in Equilibrium

We'll now establish Proposition 2.6 by contradiction. Suppose there exists a fixed, β, α there cannot be a $\hat{q}'$ and corresponding $\psi' < 1$ that O prefers to $\hat{q}$ and corresponding $\psi < 1$. Clearly, for any ϵ, sufficiently small, and a fixed J, $pr\left(1 < \frac{\mu(n_l)}{\mu(n_b)} < 1+\epsilon\right) > 0$. Thus, to avoid war, O must offer $q = (\alpha) \cap A\{n_l\}$. This implies, $EU^O = 1-\alpha$. Define q_ϵ as the q that satisfies all types with ratios $\frac{\mu(n_l)}{\mu(n_b)} > 1+\epsilon$, and corresponding ψ_ϵ for a fixed J. If O prefers not to risk war then, $1-\alpha > (1-q_\epsilon)\psi_\epsilon + (1-\psi_\epsilon)\beta$. Then, $\psi_\epsilon(J) < \frac{1-\alpha-\beta}{1-q_\epsilon-\beta}$. By definition, $q_\epsilon < \alpha$, since it must risk war. Thus, the right hand side is between 0 and 1. In this equation only ψ is a function of J. As $J \to \infty$, $\int_1^{1+\epsilon} \psi_M \to 0$. Thus, there must be some J' that satisfies, $\psi_\epsilon(J) < \frac{1-\alpha-\beta}{1-q_\epsilon-\beta} < \psi_\epsilon(J')$. But this implies that some risk of war is profitable for that J', a contradiction. This completes the proof.

C Appendix: Introducing uncertainty about preferences relative to the cost of war: war-payoffs $p-w$ not pw.

In this section, I consider an extension to the model by altering both player's cost of war to a form $p-w$ not pw. This introduces a new type of uncertainty that O must contend with: uncertainty about S's sum total preferences for the pie $(\theta_1 q_1 + \theta_2 q_2)$ relative to the cost of war (what is commonly referred to as uncertainty about resolve).

We'll show that adding this type of uncertainty does not change the basic result: if diplomacy induces O to alter his offering strategy, it also produces an equilibrium offer that (weakly) raises the risk of war. This happens even though O faces a risk-return trade-off about S's preferences relative to the cost of war leading O to sometimes accept a risk even in a balanced offer. However, even though O would accept this type of risk from a balanced offer it does not change my main result because O also accepts this type of risk from a concentrated offer. Thus, we can think about accepting risk about S's resolve as operating somewhat independently from risk about relative preferences.

The extension closely connects my model with Trager (2011). Consistent with Trager I study the case where S has discrete (e.g. high/low) values for each issue and alter the cost of war to a constant cost. Inconsistent with Trager I still assume issues are divisible, thus allowing for balanced offers (and other divisions of the pies). Building directly off Trager has the added advantage of contrasting my result with a recent attempt in the study of war and diplomacy to understand the bargaining implications of diplomacy with heterogeneous preferences.

C.1 Model adjustments

I change the model in the manuscript in two ways. First, I assume θ_1, θ_2 can take on only either a high/low values θ_H, θ_L where $1 \geq \theta_H > \theta_L > 0$. I assume these values are drawn i.i.d with a λ probability that $\theta_j = \theta_H$ and a $1 - \lambda$ probability that $\theta_j = \theta_L$. Thus, there are only 4 realizations of S's type, which I summarize in Figure 8. For the purpose of describing messages, I still refer to S's most valuable issue $\max[\theta_1, \theta_2] = \theta_A$ and S's least valuable issue $\min[\theta_1, \theta_2] = \theta_B$ with corresponding q_A, q_B, n_A, n_B.

The messages S can send are similarly altered to account for this difference in the plausible values for S's type. I allow S to send messages about her type $m(\hat{\theta}_1, \hat{\theta}_2)$ where S can choose from messages $\hat{\theta}_1, \hat{\theta}_2 \in \{\theta_H, \theta_L, \emptyset\}$.[23] The message is cheap because sending any message does not exogenous affect payoffs. I say O's initial belief about S's type is $\gamma(2 \times \{\theta_H, \theta_L\})$

[23] To be clear, S's message claims her value for n_j is $\hat{\theta}_j$. Her actual value is θ_j.

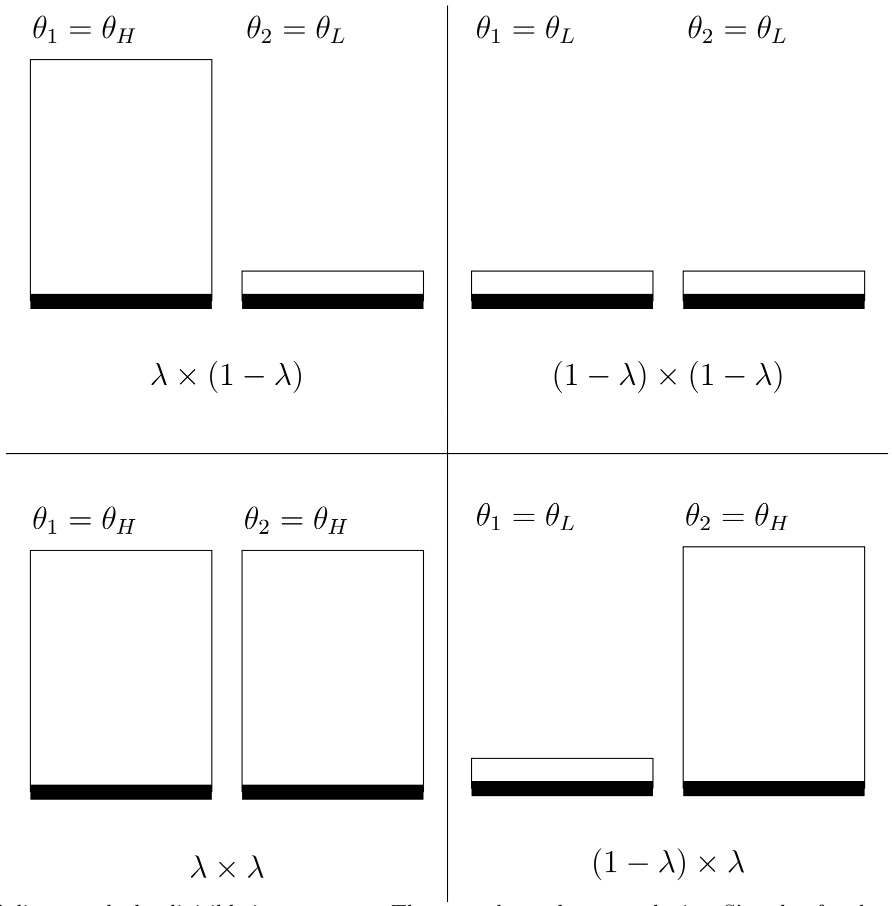

Note: Thick black lines mark the divisible issues n_1, n_2. The area above the curve depicts S's value for these issues conditions on draws θ_1, θ_2. The values for θ_H and θ_L are chosen arbitrarily for this example.

Figure 8: S's type-space and associated probabilities

and O's posterior belief after it receives a message is $\gamma^\dagger(2 \times \{\theta_H, \theta_L\}, m())$. A message is informative if $\gamma() \neq \gamma^\dagger()$.

Second, I alter each player's war payoff. I assume that players incur a constant cost from war. Therefore, $EU^S(War) = p\frac{\theta_1+\theta_2}{2} - w$ and $EU^O(war) = 1 - p - w := 1 - W$.[24]

Changing the functional form of S's war pay-off creates additional complications for O. In the baseline model, S's minimum demand only depended on the ratio of values between θ_1, θ_2. For example, any type that satisfied $\theta_1 = 2\theta_2$ held an identical minimum demand. This feature allowed O to make offers based only on O's expectations about the relative

[24]Later I will use the notation $1 - W$ to make comparisons between O's war payoff and O's pay-off from making an offer of length t that is accepted: $1 - t$.

value between θ_1, θ_2. This allowed O to always make a balanced offer that avoided war and exactly meet S's minimum demand. By changing S's war pay-off to a constant cost, I am forcing O to confront two kinds of uncertainty simultaneously. O must make offers based on S's relative value for the two issues, and S's total value relative to the cost of war. In this extension, the type $\theta_1 = \theta_2 = \theta_H$ has a different minimum demand to the type $\theta_1 = \theta_2 = \theta_L$

As a result of these modeling choices, we can summarize the different types of S along two main dimensions: her value for war, and her value for receiving different types of offers of the same length. As a result of these differences, we can think about three different types of S: highly resolved, weakly resolved and mixed.

First, there is one high-resolved type: $\theta_1 = \theta_2 = \theta_H$. This type has the highest expected value from war: $EU^S(war|\theta_1 = \theta_2 = \theta_H) : p\theta_H - w$. Since this type values both issues the same (θ_H), her value for any offer of a fixed length t is $\theta_H t$.

Second, there is one weakly resolved type: $\theta_1 = \theta_2 = \theta_L$. This type's expected value from war is the lowest: $EU^S(war|\theta_1 = \theta_2 = \theta_L) : p\theta_L - w$. Since this type values both issues θ_L, her value for any offer of a fixed length t is $\theta_L t$.

Finally, there are two types that prefer one issue over another (I call these mixed types): $\theta_1 = \theta_H, \theta_2 = \theta_L$ and $\theta_1 = \theta_L, \theta_2 = \theta_H$. These two mixed types have the same war pay-off: $EU^S(war|\theta_1 \neq \theta_2) : \frac{p(\theta_H + \theta_L)}{2} - w$, which is between the war payoffs of the high and weakly resolved types.

Unlike other types, a mixed type's value from accepting an offer of a fixed length t depends on how that offer is distributed between the high and low valued issues. For any offer $q(q_A = t_A, q_B = t_B)$, a mixed type receives $\frac{q_A\theta_H + q_B\theta_L}{2}$. They would extract less utility from an offer of the same length t that was $q(q_A = t_A - x, q_B = t_B + x)$. These types extract the most from an offer of length t if that offer is concentrated in their high-valued issue: $q(q_A = t, q_B = 0)$. They extract the least from an offer $q(q_A = 0, q_B = t)$.

Modeling war as a factor creates corner many different corner conditions because S's minimum demand can be less than 0. I avoid these conditions by assuming that $w < \theta_L$ and

$2w < \theta_H p$. Relaxing these assumptions does not change my basic result.

C.2 Analysis

The purpose of the analysis is to show that including diplomacy alters the offer that O makes in a way that raises the risk of war. I progress to that conclusion in four steps. First, I define four types of offers based on what O can achieve by making these offers. I argue that these four offers are the only offers that O can make on the equilibrium path given any set of posterior beliefs (Lemma C.1).[25] Second, I define equilibrium behavior that follows in a babbling equilibrium and identify the associated probabilities of war that O accepts for his equilibrium offer for all parameters of the game (Lemma C.2).[26] Third, I define equilibrium behavior that follows from a credible message similar to the message I outlined in the manuscript for all parameters of the game. I also define the associated risk of war that O accepts in equilibrium (Lemma C.3).

Finally, I contrast the results from the babbling equilibrium and the credible cheap talk equilibrium to understand how credible cheap-talk influences the risk of war that O accepts form his on path offer (Proposition C.4). I reach two conclusions. If O observed a credible message and made the same type of offer that he would have made in the babbling equilibrium, then the credible message led to an offer that (weakly) reduced the risk of war. But if O observed a credible message and made a different type of offer that he would have made in the babbling equilibrium, then the credible message led to an offer that increased the risk of war.

C.2.1 Four plausible offers O makes on the path

I now focus on the 4 plausible offers that O could make in equilibrium. Each of these offers factors in two different components: the total length of the offer (t), and the mapping

[25] This approach is different from what I took in the baseline model. In the baseline model I defined the minimum demand given each offer type. In this extension, it is easier to look directly at differences between types because there are only a few different types.

[26] The babbling message is uninformative. Therefore, the results are equivalent to the case where S is prohibited from sending a message.

of the offer onto each issues: $t \rightarrow t_1, t_2$, such that $t = t_1 + t_2, q(q_1 = t_1, q_2 = t_2)$. Below I explain O's logic for each offer and describe O's expected utility from making each offer. Once I've defined these different offers, I argue that O's equilibrium offer will always be one of these 4 offers.

First, O can make the smallest offer that satisfies all types with certainty. I notate this offer $\bar{q}$ of total length $\bar{t}$. This offer must meet the high-resolved type's minimum demand: $\bar{t}\theta_H = p\theta_H - w$. This solves for: $\bar{t} = p - \frac{w}{\theta_H}$. Since this high-resolved type values both issues equally, she would be willing to accept any offer $q(q_1, q_2)$ of a fixed length $\bar{t}$.[27] But this is not always true for the mixed types. For example, the mixed type $\theta_1 = \theta_L, \theta_2 = \theta_H$ values an offer $q : q_1 = \bar{t}, q_2 = 0$ relative to her war pay-off: $\frac{\theta_L}{2}(p - \frac{w}{\theta_H}) > \frac{p(\theta_L+\theta_H)}{2} - w$. Solving this inequality, this type is only willing to accept concentrated offers against her least valuable issue if: $0 > \frac{\theta_L w}{\theta_H} + \theta_H p$; impossible.

It follows that the smallest offer that all types will accept for all values of the game's parameters is the balanced offer that meets the high-resolved type's minimum demand: $\bar{q} : (q_1 = q_2 = \bar{t}/2)$.[28] This offer leaves O with an expected utility:

$$EU^O(\bar{q}) : 1 - p + \frac{w}{\theta_H} = 1 - \bar{t} \tag{15}$$

Notice this expected utility does not depend on O's beliefs about S's type. The reason is that all types accept this offer. O's utility is simply what he keeps given that R will accept the offer (which R always will).

Second O can make the smallest possible offer that at least one type will accept. I notate this offer $\tilde{q}$ of length $\tilde{t}$. This offer is the concentrated offer that leaves just one mixed type indifferent with war. This offer must satisfy the following inequality: $\tilde{t}\theta_H > \frac{p(\theta_L+\theta_H)}{2} - w$.

[27]Clearly O is indifferent between different offers of the same length if and only if there is an equal probability that S will accept them.

[28]My results are identical if I study an unbalanced offer that all types accepted of length $\bar{t}$. It is easier to focus on a single offer that I know all types will accept for all parameters of the game.

This inequality assumes that O makes an offer $\tilde{q}(q_A = \tilde{t}, q_B = 0)$ where A is the issue that a mixed type values high and B is the issue that the mixed type values low. Solving for $\tilde{t} = \frac{p(\theta_L+\theta_H)}{2\theta_H} - \frac{w}{\theta_H}$.

O's expected utility from playing an offer $\tilde{q}$ is

$$EU^O(\tilde{q}|\gamma(.), \tilde{t} \to q_A) = \left[pr(\theta_1 \neq \theta_2) \times pr(\tilde{t} \to q_A)\right]\left(1 - \tilde{t}\right) + \left[1 - pr(\theta_1 \neq \theta_2) \times pr(\tilde{t} \to q_A)\right](1 - W) \quad (16)$$

This expression captures the uncertainty and risk that goes with making a smaller, concentrated offer. Notice that O's expected utility from making this offer is conditional on two factors: (1) O's beliefs about S's type $\gamma(.)$; (2) O's beliefs about whether his concentrated offer correctly targets S's highest value issue: $\tilde{t} \to q_A$. The probability reported in the first term defines O's expectation that S is a mixed type ($pr(\theta_1 \neq \theta_2)$ and that O has correctly guessed the issue that the mixed types values high ($pr(\tilde{t} \to q_A)$) . This term is multiplied by O's utility in the case that S accepts an offer $\tilde{q}$ (O gets $1 - \tilde{q}$). The second term is the probability that S is any other type, multiplied by O's war pay-off $(1 - W)$ given these types reject $\tilde{q}$.

Third O can make the best offer that leaves the weakly resolved types exactly indifferent with war: $\theta_L \hat{t} = p\theta_L - w$.[29] I notate this offer $\hat{q}$ of length $\hat{t}$. This offer is of length: $\hat{t} = p - \frac{w}{\theta_L}$. Since the weakly-resolved type values both issues equally, she would be willing to accept any offer $q(q_1, q_2)$ of a fixed length $\hat{t}$. But if O concentrates this offer on a single issue he can also satisfy one of the mixed types: $\theta_H \hat{t} > \frac{p(\theta_L+\theta_H)}{2} - w \equiv \theta_H p - \frac{\theta_H w}{\theta_L} > \frac{p(\theta_L+\theta_H)}{2} - w$. This solves for $\theta_L > \frac{2w}{p}$, true by assumption.

As a result, O can make an offer of length $\hat{t}$ concentrated against a single issue that satisfies the weakly-resolved types and one mixed type. O's expected utility from playing an

[29]By best, I mean the version of this offer that maximizes O's utility.

offer $\hat{q}$ is:

$$EU^O(\hat{q}|\gamma(.), \hat{t} \rightarrow q_A) = \left[pr(\theta_1 \neq \theta_2) \times pr(\hat{t} \rightarrow q_A) + pr(\theta_1 = \theta_2 = \theta_L)\right]\left(1 - \hat{t}\right)$$
$$+ \left[1 - pr(\theta_1 \neq \theta_2) \times pr(\hat{t} \rightarrow q_A) - pr(\theta_1 = \theta_2 = \theta_L)\right](1 - W) \quad (17)$$

The probability in the first term defines O's expectation that S is either a mixed type and O has correctly guessed that mixed type's high value issue, or that O is the weakly resolved type. This probability is multiplied by O's expected utility from S accepting an offer $\hat{q}$ of length $\hat{t}$. The second term is O's expectation that war will occur.

Fourth, O can make an offer that satisfies both mixed types. I notate this offer $\underline{q}$ of length $\underline{q}$. This offer is a balanced offer that leaves both mixed types indifferent with their minimum demand: $\frac{t(\theta_L+\theta_H)}{2} > \frac{p(\theta_L+\theta_H)}{2} - w$. This solves for $\underline{t} = p - \frac{2w}{\theta_H - \theta_L}$.[30]

O's expected utility from playing an offer $\underline{q}$ is:

$$EU^O(\underline{q}|\gamma(.)) = [1 - pr(\theta_1 = \theta_2 = \theta_H)]\,(1 - \underline{t}) + [pr(\theta_1 = \theta_2 = \theta_H)]\,(1 - W) \quad (18)$$

Since this is a balanced offer that satisfies all mixed types, O's utility is not conditional on where O chooses to make an offer ($\hat{t} \rightarrow q_A$). O expects S to accept this offer no matter what S's favorite issue is. The only type that will not accept the offer is the highly resolved type $(1 - pr(\theta_1 = \theta_2 = \theta_H))$. In that case, O expects his war payoff $(1 - W)$.

Lemma C.1 *For any set of feasible posterior beliefs $\gamma^\dagger(.)$, and given that S rejects offer if and only if they are strictly less than her war payoff, there is no offering strategy that leaves O with larger expected utility than $\tilde{q}, \hat{q}, \bar{q}, \underline{q}$.*

By feasible posterior beliefs, I mean O accepts some non-negative probability that S is a type in the type space. I exclude the possibility that O believes S is some type not in the type space. The proof is obvious. Clearly, O will never set $t > \bar{t}$ because all types accept $\bar{t}$

[30]It is easy to see that $\underline{t} > \hat{t}$ and so also $\underline{t}$ also satisfies the weakly resolved type.

and O loses in the size of his offers. O will never set $t < \hat{t}$ because that would guarantee war and $1 - \bar{t} > 1 - W$. Further, O gets no additional benefit from deviating from these four offers because the probability that S accepts is discrete not marginal in the length of t.

Lemma C.1 implies that O's best reply given any set of beliefs about O's type is the one that maximizes his expected utility given the four utility equations above. All that's left to do is fill in O's values for his beliefs. This point is critical in how I structure the analysis that follows: No matter what O's posterior beliefs about S's type are, O will make one of four offers. As a result, I can emphasize how different messages alter O's expectation about S's type, but hold constant O's expected value given each realization of the game (whether or not S accepts or rejects a particular offer of a fixed length).

In summary, O's best offer is either $\hat{q}, \tilde{q}, \bar{q}, \underline{q}$. In choosing between these offers, O weighs 2 factors. First, O considers what he will get if S accepts or does not accept. This varies in the size of what O gives up in the different condition: $W > \bar{t} > \underline{t} > \tilde{t} > \hat{t}$. Second O considers the probability that S will accept, and the adverse consequences that follow (war) if S does not accept (with payoff $1 - W$). As we shall see, given the posterior beliefs that we consider, the probability that S accepts is the inverse of the offers' lengths. This creates a trade-off for O between making larger offers that more types accept, and smaller offers with a larger risk of war.

The extent that O faces this trade-off depends on O's posterior beliefs: $\lambda^{\dagger}(m(), 2 \times \lambda)$. To the extent that S's message is credible, S can influence O's beliefs which alters O's best choice.

C.2.2 Babbling equilibrium

I now report the equilibrium results that follow from a babbling message where S's message has no effect on O's beliefs. Define a messaging strategy $m(bab)$ as one where S sends a message about her type drawn from the complete message space independent of her actual type. As with all babbling equilibrium, this message cannot be informative and so $\gamma^{\dagger}() = \gamma(2 \times \lambda)$. As a result, the babbling equilibrium is equivalent to the case where S is prohibited from sending messages.

First, I describe strategic behavior and beliefs under the babbling equilibrium. In particular, I highlight one difference between this result and the baseline model about the chance of war:

Lemma C.2 *In the babbling equilibrium S sends message $m(bab)$. O processes that message and forms posterior beliefs $\gamma^{\dagger}() = \gamma()$. If*

$$\lambda^2 > \frac{\bar{t} - \underline{t}}{W - \underline{t}} \tag{19}$$

$$\lambda(1-\lambda) < \frac{W - \bar{t}}{W - \tilde{t}} \tag{20}$$

$$\lambda > \frac{\bar{t} - \hat{t}}{W - \hat{t}} \tag{21}$$

are satisfied, O offers $\bar{q}$ and accepts no risk of war. If

$$\lambda^2 < \frac{\bar{t} - \underline{t}}{W - \underline{t}} \tag{22}$$

$$\lambda < \frac{W - \underline{t}}{\underline{t} - \tilde{t}} \tag{23}$$

$$\lambda > \frac{\underline{t} - \hat{t}}{W - \underline{t}} \tag{24}$$

are satisfied, O offers $\underline{q}$ and accepts a λ^2 probability of war. If

$$\lambda < \frac{\bar{t} - \hat{t}}{W - \hat{t}} \tag{25}$$

$$\lambda < \frac{W - \hat{t}}{2 - W - \tilde{t}} \tag{26}$$

$$\lambda < \frac{\underline{t} - \hat{t}}{W - \underline{t}} \tag{27}$$

are satisfied, O offers $\hat{q}$ and accepts a λ probability of war. If

$$\lambda(1-\lambda) > \frac{W - \bar{t}}{W - \tilde{t}} \tag{28}$$

$$\lambda > \frac{W - \hat{t}}{2 - W - \tilde{t}} \tag{29}$$

$$\lambda > \frac{W - \underline{t}}{\underline{t} - \tilde{t}} \tag{30}$$

are satisfied, O offers $\tilde{q}$ and accepts a $1 - (1-\lambda)\lambda$ probability of war.

S rejects these offers if they are less than her minimum demand and accepts otherwise.

The equilibrium result follows from finding O's highest expected utilities from the four offering strategies discussed above given O's beliefs following $m(bab)$. In the babbling equilibrium, I can plug in the prior probabilities into O's expected utility equations.[31]

$$EU^O(\bar{q}) : 1 - \bar{t} \quad (31)$$

$$EU^O(\tilde{q}|\gamma^\dagger(m(bab)), pr(\tilde{t} \to q_A) = 1/2) = (1-\lambda)\lambda\left(1-\tilde{t}\right) + (1-(1-\lambda)\lambda)(1-W) \quad (32)$$

$$EU^O(\hat{q}|\gamma^\dagger(m(bab)), pr(\tilde{t} \to q_A) = 1/2) = (1-\lambda)\left(1-\hat{t}\right) + \lambda(1-W) \quad (33)$$

$$EU^O(\underline{q}|\gamma^\dagger(m(bab))) = (1-\lambda^2)(1-\underline{t}) + \lambda^2(1-W) \quad (34)$$

The conditions reported in Lemma C.2 follow from comparing these four expected utility equations. Each condition finds the parameters where one of these conditions dominates all the rest. Notably, I have written all denominators and numerators on the RHS of all these conditions as positive values. We've already shown that O will not make any other offer. Further, we've shown that S will accept any offer if it beats her war payoff and reject otherwise. Finally, there are no off-path messages. Thus, I need not consider any deviations in S's message. This completes the proof.

As in the baseline model, the value O can extract by a concentrated offer is reduced because O does not know which issue S values the most:($pr(\tilde{t} \to q_A) = 1/2$). But this only effects the two offers in which O concentrates his offer on a single issue. When O balances his offer, this is not a relevant factor.

A difference between this model and the one reported in the manuscript is that O makes offers that accept a positive probability of war even with no additional information about S's type. The risk that O accepts depends on O's relative value between the different options.

[31] For help deriving these probabilities, return to Figure C.1 where the probabilities of each type are laid out.

I chose to report the results in terms of $\tilde{t}, \hat{t}, \bar{t}, \underline{t}, W$ and not the parameters θ_H, θ_L, p, w. The reason is that I am only interested in whether or not S accepts a smaller probability of war following a credible message. As we shall see in a moment, it is simpler to do so the way that I've presented the information.[32]

C.2.3 Credible Cheap-talk equilibrium

I'll now study equilibrium behavior that follows from a credible cheap-talk equilibrium that most closely resemble the cheap-talk message in the manuscript. I notate this messaging strategy as strategy $m(A)|\theta_1, \theta_2$. Under this strategy, all types of S choose a signal that is consistent with her preference ordering (if they have one) but reveals no information about the relative value of their motives. If S is the mixed type $\theta_1 > \theta_2$, she sends a message $m(A = 1) \implies \theta_1 = \theta_H$. If S is mixed type $\theta_2 > \theta_1$ she sends message $m(A = 2)$, which mixes over all remaining messages. All other types mix over all the feasible messages proportionately.[33]

Proposition C.3 *There is a credible cheap-talk equilibrium in which S sends message $m(A)$. On the path, S observes her type and sends message $m(A)|\theta_1, \theta_2$. Without loss of generality suppose O observes a message $m(A = 1)$, O's posterior beliefs $\gamma^\dagger \neq \gamma$ such that $pr(\theta_2 > \theta_1) = 0$, $pr(\theta_1 = \theta_2 = \theta_H) = \frac{\lambda^2}{1-(1-\lambda)\lambda}$, $pr(\theta_1 = \theta_2 = \theta_L) = \frac{(1-\lambda)^2}{1-(1-\lambda)\lambda}$, $pr(\theta_1 > \theta_2) = \frac{(1-\lambda)\lambda}{1-(1-\lambda)\lambda}$. If*

$$EU^O(\hat{q}|\gamma^\dagger(m(A))) : \frac{1-\lambda}{\lambda^2} < \frac{W - \bar{t}}{\bar{t} - \hat{t}} \tag{35}$$

$$EU^O(\tilde{q}|\gamma^\dagger(m(A))) : (1-\lambda)\lambda > \frac{W - \bar{t}}{2W - \bar{t} - \tilde{t}} \tag{36}$$

[32] With algebra, I can confirm that all four conditions are possible for some values of the parameters. Contact author for more details about the conditions under which each set of results holds. I have omitted these details because they are irrelevant to my main point.

[33] By proportionately, I mean that they mix in a way such that no message provides information about the relative intensity of their preferences.

O offers $\bar{q}$ and accepts no risk of war. If

$$EU^O(\bar{q}) : \frac{\lambda^2}{1-\lambda} > \frac{\bar{t}-\hat{t}}{W-\bar{t}} \tag{37}$$

$$EU^O(\tilde{q}|\gamma^\dagger(m(A))) : \lambda < \frac{W-\hat{t}}{2-\bar{t}-\tilde{t}} \tag{38}$$

O offers $\hat{q}$, which carries a $1 - \frac{1-\lambda}{1-(1-\lambda)\lambda}$ risk of war. If

$$EU^O(\bar{q}) : (1-\lambda)\lambda < \frac{W-\bar{t}}{2W-\bar{t}-\tilde{t}} \tag{39}$$

$$EU^O(\hat{q}|\gamma^\dagger(m(A))) : \lambda > \frac{W-\hat{t}}{2-\bar{t}-\tilde{t}} \tag{40}$$

O offers $\tilde{q}$, which carries a $1 - \frac{(1-\lambda)\lambda}{1-(1-\lambda)\lambda}$ risk of war. O never offers $\underline{q}$ on the path. S rejects any offer if it is less than her minimum demand and accepts otherwise.

S's message $m(A=1)$ rules out the mixed type $\theta_1 < \theta_2$. Yet provides no additional information. Assuming that S sends an honest message, then O's expected utilities from each offer are as follows:

$$EU^O(\bar{q}) : 1-\bar{t} \tag{41}$$

$$EU^O(\tilde{q}|\gamma^\dagger(m(A)), pr(\tilde{t} \to q_A) = 1) = \frac{(1-\lambda)\lambda}{1-(1-\lambda)\lambda}\left(1-\tilde{t}\right) + \left(1 - \frac{(1-\lambda)\lambda}{1-(1-\lambda)\lambda}\right)(1-W) \tag{42}$$

$$EU^O(\hat{q}|\gamma^\dagger(m(A)), pr(\hat{t} \to q_A) = 1) = \frac{1-\lambda}{1-(1-\lambda)\lambda}\left(1-\hat{t}\right) + \left(1 - \frac{1-\lambda}{1-(1-\lambda)\lambda}\right)(1-W) \tag{43}$$

$$EU^O(\underline{q}|\gamma^\dagger(m(A))) = \frac{1-\lambda}{1-(1-\lambda)\lambda}(1-\underline{t}) + \left(1 - \frac{1-\lambda}{1-(1-\lambda)\lambda}\right)(1-W) \tag{44}$$

O's expected utilities different from the babbling equilibrium because O has different expectations that S will accept.[34] These different probabilities follow from the Bayes' Rule application given a credible signal $m(A=1)$.

O's expectations that S will accept change in two ways. First, O's posterior beliefs rule out the possibility that S is one specific mixed type. There are only three feasible types

[34]Since all S accept $\bar{q}$, the result is identical.

remaining. The highly resolved type (prior probability λ^2), the weakly resolved type (prior probability $(1-\lambda)^2$), and one specific mixed type (prior probability $\lambda(1-\lambda)$). It follows that all probabilities are re-weighted by the remaining expectation: $\lambda^2 + (1-\lambda)\lambda + (1-\lambda)^2 = 1-(1-\lambda)\lambda$. This appears on the denominator in each posterior probability.

Second, O's posterior beliefs imply that O can target concentrated offers based on O's new high-confidence beliefs that he knows which issue S (weakly) values the most. How each probability updates depends on whether or not the offer was concentrated on a single issue (and only satisfied one mixed type).

Thus, the two concentrated offers $\tilde{q}, \hat{q}$ now satisfy all feasible mixed types. Leading to a higher overall probability that S will accept. Critically, the balanced offer $\underline{q}$ already satisfied both mixed types. As a result, O's posterior beliefs factor in that the offer already covered both of them.

The conditions for O to make each type of offer reported in Proposition C.3 follows from a direct comparison between these expected utilities. Since all offers that S rejects lead to war, the probability that O's equilibrium offer is insufficient is equivalent to the risk of war that O accepts on the equilibrium path.

We've already shown that O cannot profit from deviating from these four basic offers. Thus, these offers must be O's equilibrium best replies if S's message is incentive compatible.

Turning to S's incentives to send an honest message. Only mixed types can deviate from the equilibrium message and there only one deviation. The type $\theta_1 > \theta_2$ can send a message that implies $\theta_2 > \theta_1$. For the same reason reported in the manuscript, no mixed type can profit from such a deviation because such a message will not increase the size of the offer (both messages return offers of equal length t). Further, O will concentrate the offer on an issue that is no more valuable, and possibly less valuable to S.

Finally, I consider why $m(A)$ implies O never plays $\underline{q}$ on the path. Since $m(A)$ allows O to rule out one mixed type with certainty, O's posterior expectation that S will accept $\hat{q}$ is identical to O's expectation that S will accept $\underline{q}$. Since $\hat{t} < \underline{t}$ it follows that

$EU^O(\hat{q}|\gamma^\dagger(m(A))) > EU^O(\underline{q}|\gamma^\dagger(m(A)))$ for any parameters of the game. Thus, following a message $m(A)$ O never plays $\underline{q}$. This completes the proof.

C.2.4 When Cheap-talk raises the probability of war

I'll now contrast the results from the babbling equilibrium and the credible cheap-talk equilibrium. Let $x \in X$ be a list of values for all the parameters of the game, and X be the complete set of all lists of all parameters of the game. This is every possible condition of the game that can arise. Let $\rho(x|m(bab)) \in \left\{\tilde{q}, \hat{q}, \bar{q}, \underline{q}\right\}$ be O's equilibrium offer in the babbling equilibrium for parameters of the game x, which carries a probability of war $\omega(x|m(bab))$. Let $\rho(x|m(A)) \in \{\tilde{q}, \hat{q}, \bar{q}\}$ be O's equilibrium offer in the credible cheap talk equilibrium described in proposition C.3, which carries a probability of war $\omega(x|m(A))$. When O's offer is of the same type in both games I write: $\rho(x|m(bab)) = \rho(x|m(A))$. When O's credible message induces a different offering strategy I write $\rho(x|m(bab)) \neq \rho(x|m(A))$.

The conditions under which credible cheap-talk (weakly) increases the risk of war are those in which $\omega(x|m(bab)) \leq \omega(x|m(A))$. The conditions under which credible cheap-talk (weakly) decreases the risk of war are those in which $\omega(x|m(bab)) \geq \omega(x|m(A))$

Proposition C.4 *Suppose (1) a list of fixed parameters x such that O's best reply to an equilibrium babbling message is the same type of offer as it would have been in the equilibrium with a credible cheap-talk message (i.e $\rho(x|m(bab)) = \rho(x|m(A))$), then diplomacy decreases the risk of war O accepts in equilibrium: $\omega(x|m(bab)) \geq \omega(x|m(A))$. Suppose (2) a list of fixed parameters x such that O's best reply to an equilibrium babbling message is a different type of offer than it would have been in the equilibrium with a credible cheap-talk message (i.e $\rho(x|m(bab)) \neq \rho(x|m(A))$), then diplomacy increases the risk of war O accepts in equilibrium: $\omega(x|m(bab)) \leq \omega(x|m(A))$.*

Proposition C.4(1) states that any time a credible message $m(A)$ does not entice O to alter his strategy from what he would have played absent a diplomatic message, then diplomacy decreases the risk of conflict (assuming that O's offer would carry a positive risk of war following a message from $m(bab)$). This result, is largely consistent with how prior studies have characterized the role of cheap-talk and carries over an intuition from the model in the manuscript. There are cases where O knows that S may prefer one issue over the other, but does not know which issue. Nevertheless, O's incentives are such that he makes a concentrated offer even though O is unsure which issue S prefers. This concentrated offer carries some risk that the offer will fail because O has incorrectly guesses S's favorite issue.[35] In cases where O will use information about S's preferred issue to change where O concentrates his offer, but not alter the size of the offer he makes, then more information must reduce the risk of war.

The result follows immediately from comparing the risks of war that O accepts in the babbling equilibrium and the credible cheap-talk equilibrium given the same offer $\tilde{q}, \hat{q}$.[36] That is: $\omega(x|m(bab), \hat{q}) < \omega(x|m(A), \hat{q})$ and $\omega(x|m(bab), \tilde{q}) < \omega(x|m(A), \tilde{q})$. This result is intuitive. Each of these offers only satisfied one mixed type. Following a babbling message, there were two possible mixed types and always some risk that O would get the offer wrong because O targets the wrong issue. This risk is gone in the credible cheap talk equilibrium. It follows that O's offer is safer, and carries a smaller risk of conflict.

Proposition C.4(2) states that in every case that a credible message $m(A)$ entices O to alter his strategy from what he would have played absent a diplomatic message, O's equilibrium offer following $m(A)$ carries a larger risk of conflict.

The result follows from three facts. First, the conditions that define O's preferences for different offers imply that S's message $m(A)$ never entices O to change his offer from one kind of concentrated offer to another. On the path, O only makes one of two concentrated offers: $\tilde{q}, \hat{q}$. Notice that O's preference for choosing one of these concentrated offers over the other is identical in both games. That is, the inequality that defines $EU^O(\hat{q}|\gamma^\dagger(m(A))) >$

[35] In this analysis there is also a risk that O's offer fails because O mis-understands S's value for issues relative to the cost of war. In the manuscript that this type of risk was not necessary to produce this kind of result.We'll show in a moment that this other kind of risk is not driving the result here either.

[36] There is no risk of war with offer $\bar{q}$ and O never offers $\underline{q}$ so I need not consider it.

$EU^O(\tilde{q}|\gamma^\dagger(m(A)))$ is the same as the inequality that defines $EU^O(\hat{q}|\gamma^\dagger(m(bab))) > EU^O(\tilde{q}|\gamma^\dagger(m(bab)))$: $\lambda < \frac{W-\hat{t}}{2-\hat{t}-\tilde{t}}$. Given the condition is identical in both equilibrium, for the same set of parameters x O's preference for one over the other must be constant.[37]

Second, S's message $m(A)$ can entice O to shift from a balanced offer to a concentrated offer but not the other way around. We already saw that following a message $m(A)$ O's strategy $\hat{q}$ strictly dominates $\underline{q}$. As a result, O cannot play $\underline{q}$ on the path following a credible message $m(A)$.

I now turn to the conditions where O's equilibrium offer is $\bar{q}$ on the path. Comparing these conditions across the equilibrium defined in Lemmas C.3 and C.2, $\rho(x|m(A)) = \bar{q} \subset \rho(x|m(bab)) = \bar{q}$. Thus, for any set of parameters x, that O plays $\bar{q}$ on the path following a message $m(A)$, O would also play $\bar{q}$ following an equilibrium message $m(bab)$. But the opposite is not true. As a result, it must be that there are conditions where O shifts his offer from $\bar{q}$ to one of the concentrated offers but not the other way around.

Finally, the risk of war O accepts for a concentrated offer following a message $m(A)$ is larger than the risk O would have accepted from any balanced offer given a message $m(bab)$.[38] Obviously, O accepts no risk of war following an offer $\bar{q}$. Thus, if O shifts from this offer to a concentrated offer he must accept a greater risk of war.

I now check the risk that O accepts from offering $\underline{q}$ in the babbling equilibrium—$\omega(x|m(bab), \underline{q}) = \lambda^2$— against the the risk O accepts from offering the concentrated offer $\hat{q}$ following a message $m(A)$—$\omega(x|m(A), \hat{q}) = 1 - \frac{1-\lambda}{1-(1-\lambda)\lambda}$. Comparing these two values, O must always accept a larger risk of war following $m(A)$ if: $1 - \frac{1-\lambda}{1-(1-\lambda)\lambda} > \lambda^2 \implies 1 > \lambda$, which is always true.[39]

In summary, I've shown that every balanced offer that O makes on the path in the

[37] Similarly, O never switches between different balanced offers.

[38] This is not obviously true, because as proposition C.3(1) argued, S's credible message implies that O accepts a smaller risks from concentrated offers.

[39] The risk of war is clearly more in the case that O plays an offer $\tilde{q}$.

babbling equilibrium carries a smaller risk of war than any concentrated offer that O plays on the path following a message $m(A)$. I've also shown that there are certain parameters x for which O would make a balanced offer in the babbling equilibrium, but a concentrated offer in the equilibrium that follows from a message $m(A)$. However, there are no conditions in which O's offer reverts from a concentrated offer to a balanced offer. Finally, I've shown that there are no conditions in which O switches between different concentrated offers given S's different equilibrium messages. It follows that O must accept a risk of war that is at least as large, and sometimes larger in an equilibrium in which S plays a message $m(A)$ compared to $m(bab)$ and this leads O to alter his offer such that $\rho(x|m(bab)) \neq \rho(x|m(A))$.