

Mutual Optimism and War

Mark Fey University of Rochester
Kristopher W. Ramsay Princeton University

Working with the definition of mutual optimism as war due to inconsistent beliefs, we formalize the mutual optimism argument to test the theory's logical validity. We find that in the class of strategic situations where mutual optimism is a necessary condition for war—i.e., where war is known to be inefficient, war only occurs if both sides prefer it to a negotiated settlement, and on the eve of conflict war is self-evident—then there is no Bayesian-Nash equilibrium where wars are fought because of mutual optimism. The fundamental reason that mutual optimism cannot lead to war is that if both sides are willing to fight, each side should infer that they have either underestimated the strength of the opponent or overestimated their own strength. In either case, these inferences lead to a peaceful settlement of the dispute. We also show that this result extends to situations in which there is bounded rationality and/or noncommon priors.

Why do states fight costly wars when less costly negotiated settlements are possible? Must there not be some mutually agreeable alternative to war that can produce the same result without incurring the social loss? Could not decision makers agree to distribute the disputed territory or assets in a way consistent with their beliefs about the likely outcome of conflict, saving both sides significant death and destruction? In this paper, we address one specific rationalist answer to these questions. As Blainey (1988) intimates, the high hopes on the eve of war suggest a sad conclusion: wars only occur when both rivals believe they can achieve more through fighting than through peaceful means. How might this be so? Obviously, when two countries are involved in a war, if one side wins then the other loses. We might then conclude that at least one side, in particular the loser, would prefer some peaceful method of resolving the dispute if she were certain of the outcome. But war is an uncertain process. Given this uncertainty, the leaders of the two countries must each form expectations about the results of a conflict to guide their decision making. The

decision-making processes in these kinds of situations are frequently studied in international relations (Bueno de Mesquita, Morrow, and Zorick 1997; Kim and Bueno de Mesquita 1995; Niu, Ordeshook, and Rose 1989; Stein 1990; Wagner 1994). We know that when uncertainty prevails, leaders' expectations will be shaped by any special knowledge or information they might possess. If these expectations are inconsistent in that both antagonists think their side will be better off fighting a war, the argument goes, then neither side would be willing to participate in a peacefully negotiated settlement (Blainey 1988; Wittman 1979). It then seems that such *mutual optimism* could lead to a rational choice of war by both countries (Morrow 1985; Wagner 2000; Werner 1998; Wittman 2001). In this setting, the root cause of war is the inconsistent expectations that arise because of private information (Fearon 1995, 390).¹

A similar argument is made by scholars who study why wars end. Given a war, it is argued that countries continue to fight until their individual assessments of the likely outcome of combat converge. At that point, both

Mark Fey is associate professor of political science, 109E Harkness Hall, University of Rochester, Rochester, NY 14627 (mark.fey@rochester.edu). Kristopher W. Ramsay is assistant professor of politics, 033 Corwin Hall, Princeton University, Princeton, NJ 08544 (kramsay@Princeton.edu).

Earlier versions of this paper were presented at the University of Rochester, Columbia University, and NYU. We thank Scott Ashworth, Bruce Bueno de Mesquita, Songying Fang, Tanisha Fazal, Erik Gartzke, Shigeo Hirano, Adam Meirowitz, John Patty, Pablo Pinto, Robert Powell, Quinn Ramsay, Anne Sartori, Curt Signorino, Allan Stam, Randy Stone, Branislav Slantchev, and Robert Walker, as well as other seminar participants. We would also like to thank the editor of the *AJPS* and two anonymous reviewers for helpful comments. Any remaining errors are our own. Kris Ramsay acknowledges financial support from NSF grant SES-0413381.

¹We note that this explanation does not require that both sides think they are more likely to win than lose. That is, the countries' respective estimates of the probability of success need not be greater than one-half; they need only be inconsistent with the fact that the actual probability of winning for the two sides sum to 1.

American Journal of Political Science, Vol. 51, No. 4, October 2007, Pp. 738–754

©2007, Midwest Political Science Association

ISSN 0092-5853

sides can agree to a settlement they prefer to continued war. As Wittman says, inefficient wars start and continue because the probability that state 1 believes it may win need not equal one minus the probability that state 2 wins, "... as the probability estimates [of each side] are based on different sources of information (1979, 755)."²

In this paper, we reconsider the mutual optimism hypothesis by presenting a model integrating players' knowledge and strategy in a single framework and analyzing the equilibria of a class of games that capture the key features of the mutual optimism argument. We find that the simple logic of "war by mutual optimism" is misleading. That is, if war is known to be inefficient, and on the eve of combat war is self-evident, then war cannot occur because of mutual optimism. We show this result is robust by generalizing it to the case where leaders are not perfectly rational information processors.

Our goal here is significantly less ambitious than presenting an all-inclusive theory of war and peace. Rather, we aim to formalize the mutual optimism argument, show that this widely cited intuition has flaws, and provide a clear reason why the relationship between mutual optimism and war is not as simple as it first appears. As such, our analysis provides a framework for a rationalist theory of optimism and war. We do not, however, make the broader claim that incomplete information is an unimportant part of the war puzzle; instead we aim to reassess one particular mechanism, namely mutual optimism, that purports to link incomplete information to war.

In addition to this rational framework, we also consider more general assumptions regarding the way decision makers process information. In particular, we consider how our results may change if information is processed in a "boundedly rational" fashion. We are able to establish two results. First, in a setting where decision makers, for any number of reasons, either take information at face value or ignore information at particular states of the world, war still cannot be the result of mutual optimism. Second, we identify minimally sufficient conditions that, if violated, allow such mutually optimistic wars to occur. We also show that there is an equivalence between our results concerning information-processing errors and results from a model with fully rational decision makers that hold noncommon priors.

Before we continue, we pause to consider what it means to determine whether mutual optimism causes war. As there can be many conditions that lead to war, such as commitment problems (Fearon 1998; Leventoglu and Slantchev 2005), dissatisfied revisionist states (Powell 1999), or domestic political concerns (Downs and Rocke

1994) to name a few, we need to think carefully about the meaning of mutual optimism as an explanation of war. Obviously, in a situation in which any number of the above causes of war are present, along with mutual optimism, when war occurs we cannot be sure that mutual optimism was the cause. One way to resolve this problem, as an analyst, is to set a benchmark and consider a situation in which war occurs *only if* mutual optimism is present. This argument directs us to examine a class of models in which war occurs only if mutual optimism holds. Put another way, we need to think about strategic settings in which a lack of mutual optimism leads to peace. In order to ensure our model has this feature, we make two important assumptions. As just indicated, we make these assumptions not on the grounds that all crisis models should satisfy them, but rather that they are required to determine whether mutual optimism can logically explain war.

Our first assumption concerns how nations arrive at war. Here we assume that both parties must choose to *stand firm* for war to occur. This assumption recognizes that war is a mutual act and is often made in the coercive diplomacy literature (Bueno de Mesquita and Lalman 1992; Fearon 1994; Schultz 2001). Our second assumption is that there exists some settlement procedure that either side can choose when their opponent stands firm. This assumption just extends the logic of the first. At any given moment before war begins, a state could continue negotiations with the hopes of avoiding a fight. In order to understand how these two assumptions capture the logic of testing the mutual optimism explanation, recall that such a test requires a situation in which war is avoided when mutual optimism does not occur. So consider a crisis situation without mutual optimism. In such a situation, one of the two countries would prefer a negotiated settlement instead of war. Then if the two assumptions above hold, such a country could avoid war by choosing the settlement procedure. Therefore, war is avoided when mutual optimism is not present, as required.

In many ways these assumptions are quite general. For example, in our assumption on the existence of a settlement procedure, we do not wish to restrict ourselves to a single extensive form. We thus consider abstract settlements; moreover, we permit each side to possess private information about the likely end result of the settlement process, as well as the probability of success and the costs of war.³ By making these general assumptions

²For more recent work that builds on the ideas presented by Wittman, see Wagner (2000), Wittman (2001), Smith and Stam (2004), and Powell (2004).

³Although we permit the settlement outcome to depend on the private information of each side, we do not allow this outcome to depend on the actions chosen in the game. As we discuss in the fifth section, this assumption, as with the others, is made to insure that in our model, war is explained by mutual optimism, and not some other cause.

and considering the class of all games that satisfy them, we make our argument about mutual optimism and war “game free.” That is, our results apply to any game where peace prevails in the absence of mutual optimism. In particular, this ensures that our result is not an artifact of a particular game form.

At the risk of redundancy, we must emphasize that we make these assumptions *not* because they must hold in all conflict situations, but rather because we are interested in formalizing and exploring the mutual optimism explanation of war and, therefore, we must insure that we can distinguish this explanation from the many other possible rationalist explanations for war. Our assumptions serve to rule out these other explanations in order to isolate the viability of the mutual optimism explanation.

Our results are closely related to ones found in the economics literature on efficient exchange (Aumann 1976; Milgrom and Stokey 1982; Rubinstein and Wolinsky 1990; Sebenius and Geanakoplos 1983; Tirole 1982). Formulated as a “no trade theorem,” these results state that it cannot be common knowledge between two individuals that both will gain from a trade of a risky asset. In this formulation, the return on the risky asset is compared to a fixed status quo payoff. However, we consider the case in which both possible outcomes (in our case, war and peace) are uncertain, and agents possess private information about the relative value of each. In addition to this extension, our formalization of the mutual optimism argument and the corresponding theory of knowledge can help clarify these issues in the literature on international relations.

Within the literature on uncertainty and war, Niou, Ordershook, and Rose (1989) and Fearon (1995) briefly discuss the implications of rationality and common knowledge on the mutual optimism argument and link this discussion to the theoretical results in economics, but they focus on Aumann’s assumption about the common knowledge of the players’ posterior beliefs. They fail to discuss the relationship between what is known in equilibrium and what beliefs a player could hold. In addition, Patty and Weber (2006) discuss a similar problem in the context of the democratic peace.

Finally, although we restrict attention in this paper to models in which war is a mutual act because these models provide the clearest possible test of the mutual optimism explanation for war, many other models assume that war is not a mutual act. Such models of *unilateral* war have received considerable attention in the literature, and the results are well known (Filson and Werner 2002; Powell 2004; Slantchev 2003). In a companion paper (Fey and Ramsay 2007b) we explore how the mutual optimism explanation for war fares in this class of models. As is well

known, war can be an equilibrium outcome in games when war is a unilateral act. However, we show that if war occurs in equilibrium, then there must be states of the world in which war occurs *and* only one side is “optimistic” in that it thinks war is better than a negotiated settlement, given its private information. Thus, in a model of unilateral war, there will always be cases in which there is no mutual optimism but war still occurs. Under these alternative assumptions mutual optimism is neither necessary nor minimally sufficient and, as a theoretical concept, has little to tell us about the causes of war.

The rest of this paper proceeds as follows. In the next section we present a simple game that highlights the nature of our main results. The third section describes the class of games for which our results apply and gives our main result, that in such games mutual optimism cannot lead to war. Then we extend the result to players who process information imperfectly. We end with a discussion of the implications of our results for theories of war and peace.

Example: Betting on the Highest Die

We begin with an example that will provide some intuition for the general result that follows. Consider a game between two players, Alice and Bob, who have a choice between participating and not participating in a contest whose outcome is determined by chance. To start, each player is given a (fair) die to roll and the die are compared. If Alice’s die generates a higher number than Bob’s, Alice wins. If the number on Bob’s die is the larger he wins, and Bob and Alice tie otherwise. Suppose that each player maximizes expected utility and assigns a utility value of +1 to winning, –1 to losing, and 0 to a tie.

Before deciding whether to participate in the contest or not, each player receives some information about the roll of the dice. In particular, Alice observes the result of her die and Bob observes the result of his die. After receiving this information, the two players simultaneously announce whether or not they agree to play the game of chance. If they both agree, payoffs are awarded as above and, in addition, each player pays a small cost c .⁴ If one or both players do not agree to play, then both receive a payoff of 0.

To begin the analysis of this game, suppose a player naively considers only the private information he or she receives and does not make any additional inferences. That is, suppose Alice, for instance, observes the results of the first die and makes the assumption that the second die is

⁴In this example, we require $0 < c < 1/6$.

equally likely to be any of the six possible values. Then it is easy to show via a simple expected utility calculation that if Alice observes 4, 5, or 6 on the first die, she would have positive expected utility for the contest, and if she observes 1, 2, or 3, her expected utility would be negative. Such nonstrategic Alice would choose to play if she observed the first die is greater than or equal to 4 and would choose not to play for values less than 4. Clearly, the same would hold for Bob.

If we stop our analysis here, the game would be played when both players have seen a roll of their die that causes them to think they are likely to win. However, these decision rules neglect an important aspect of the game and are therefore incomplete. Specifically, a strategic player must evaluate the effects of their opponent's information on their opponent's decision and incorporate those inferences into their own calculations. Consider Alice. If the contest is played in equilibrium, then it must be the case that it is "rational" for Bob to play. Therefore, in judging whether or not it is best response for Alice to play, she should consider the implications of Bob's strategy when he participates in the contest. Alice can then infer, by the above analysis, that even an unsophisticated Bob will not play if he observes that the second die is 3 or less. So, given Bob's strategy, Alice knows that the second die is at least a 4. But in that case, if Alice has observed that the first die is a 4, there is no chance of winning, because conditional on Bob agreeing to play, the best she can do is tie. Of course, by the same logic, Bob will not play if he observes the second die is a 4, because if Alice agrees to play, he can infer that the first die is greater than or equal to 4, and at best he will tie. Based on this analysis, then, Alice would only choose to play if the first die is a 5 or a 6, and Bob will also reach the same conclusion.

Now, taking this line of reasoning one step forward, consider the case in which Alice, having made the above inference about Bob, observes the first die is a 5. Now knowing that Bob will only play if the second die is a 5 or 6, agreeing to play has negative expected utility! Likewise, Bob would choose not to play if he sees a 5. We are therefore left with the decision rules in which Alice only agrees to play if she sees a 6 and Bob only agrees if he sees a 6. But, as Alice and Bob can infer these decision rules, they know that if they both agree to play the result is sure to be a tie. However, because of the cost c , both players strictly prefer not playing to a tie, and therefore neither player will agree to play, regardless of how favorable their private information is. Thus, there cannot be a Bayesian-Nash equilibrium in which the game is played because of differing beliefs that result from private information.

Now consider what the intuition highlighted by this example reveals about mutual optimism between two ad-

versaries. Suppose instead of Alice and Bob, we have two countries, East and West. Also, suppose that instead of looking at the roll of a die, the countries receive private information about the "quality" of their troops, the reliability of their allies, or some other relevant factor that directly affects their likelihood of winning a war. For East, a strong report would be equivalent to Alice seeing a high number on her die. Similarly, for West, a positive report about their troop quality would be like Bob seeing a high number on his die. In such a world, if East gets a good report (a high roll for Alice) and so does the West (a high roll for Bob), each side's expected utility from fighting might exceed that of the status quo, or some efficient settlement. Yet, like Alice and Bob, when East and West reflect upon what it means to fight an opponent who is going to resist, they must come to the same conclusion as Alice and Bob, namely that fighting cannot be a profitable alternative to peace. To summarize, the sequence of inferences described above reflect the fact that rational players in each example understand that it cannot be the case that both players expect to win in equilibrium.

Mutual Optimism, Rationality, and War

A typical story for how war might result from mutual optimism is as follows. Suppose the leaders of two countries have information about their military forces and tactics that their opponent does not. Moreover, suppose that this information influences each leader's assessment of their country's likelihood of success in combat. If both leaders then believe that their side possesses the "stronger" force, both sides may think they will prevail militarily and thus both leaders may choose to fight rather than pursue a peaceful settlement. In such an environment, the leaders' mutual sense of optimism could create a situation where there are no *ex ante* bargains both sides prefer to war, even though war is known to be *ex post* inefficient. The paradox is then that although rational leaders know that both sides cannot benefit from war at the same time, they still start wars that they would have preferred to avoid. In this section, we construct a game-theoretic model and a corresponding framework for modeling knowledge that formalizes the idea of mutual optimism as a cause of war.

Model and Assumption

Suppose two countries are facing a potential conflict. The dispute can be settled by war or resolved without armed conflict. To represent war as a mutual decision we assume

that war only occurs if both countries decide to stand firm. We also assume that war is an inefficient method of settling disputes. In this construction, we focus on the private information that the countries may have about their ability to prevail in the event of war and their expectations regarding the bargaining process. Initially, countries are uncertain about particular facets of the crisis situation, such as the balance of forces, technological differences, military strategy, latent resources of each side, support from allies, etc. The set of all such possible situations is denoted by Ω . An element of Ω is therefore a complete description of one possible situation and thus we refer to $\omega \in \Omega$ as a (possible) state of the world. For simplicity, we assume Ω is a finite set. We may then ask, what do the two sides believe before they “know” anything about the true state of the world? If we assume that differences in people’s beliefs about the state of the world are the product of private information, such as their personal background, confidential intelligence information, any inputs they may receive from advisors, etc., then it is logical to suppose all players share a “common prior.”⁵ Let the probability distribution π on Ω be this shared prior.

The “information” in our model can be naturally classified into two categories. First, there is information that players agree will affect the two sides’ respective likelihoods of success in the same way. For example, geography can favor one side, such as the advantage that being an island gave the British in the face of German attacks during World War II; it would indeed be strange to claim that the fact that Great Britain is an island was private information. Such commonly known information is described as “common knowledge.” The second type of information is private and only known to one side. This information could describe such aspects as troop quality, military strategy, and the plans of other countries in the international system. For example, the capabilities and tactics of German units on the eve of World War II would be an obvious source of private information for Germany. This private information would then generate differences in the two sides’ respective assessments of the prewar likelihoods of success.

To formalize this information structure we will need a model of knowledge.⁶ Generally, knowledge within a game is characterized by the ability of a player to distinguish between elemental states ω in Ω , as a player may possibly distinguish among decision nodes in an exten-

sive form game. We are also interested in *events*, which are naturally defined as subsets of Ω . For example, if $\Omega = \{1, 2, 3, 4, 5\}$ then $E = \{1, 3\}$ is an event. So here, events are related to information sets and describe a set of states consistent with the history of the game.

To formalize what players know and when they know it, we use a *possibility correspondence* $P_i(\omega)$, that maps every state ω to a nonempty set of states $P_i(\omega) \subseteq \Omega$. For each $\omega \in \Omega$, $P_i(\omega)$ is interpreted as the collection of states that individual i thinks are possible when the true state is ω . Equivalently, a player’s knowledge can be formalized by a *knowledge correspondence* $K_i(E) = \{\omega: P_i(\omega) \subseteq E\}$, where $K_i(E)$ is the set of states where player i knows an event E has occurred, for sure. So i knows E at every state in $K_i(E)$.

As can be seen by the relationship between K_i and P_i , a player’s knowledge can be discussed in terms of P_i or K_i . For example, let $\Omega = \{1, 2, 3, 4, 5\}$ and let player i ’s knowledge be represented by $P_i(\omega)$ taking on the following values:

$$P_i(1) = P_i(3) = \{1, 3\} \quad \text{and}$$

$$P_i(2) = P_i(4) = P_i(5) = \{2, 4, 5\}.$$

P_i implies that if the state is 1, the player thinks that the true state is either 1 or 3. Similarly, if the state is 4, then she thinks that the true state is 2, 4, or 5. Now define the events $F = \{1, 2\}$ and $F' = \{1, 2, 3\}$. By the definition of K_i , player i “knows” an arbitrary event E at ω if $P_i(\omega) \subseteq E$. In this example, if $\omega = 1$ then i clearly does not know F because $P_i(1) = \{1, 3\} \not\subseteq \{1, 2\} = F$. However, she can deduce the conditional probability of F , given P_i and a prior π_i , using Bayes’ Rule. This conditional probability is her posterior belief that F has occurred given her private information, $P_i(1)$. On the other hand, at $\omega = 1$ player i does know F' since $P_i(1) = \{1, 3\} \subseteq \{1, 2, 3\}$.

Notice that in the preceding example, the two sets $\{1, 3\}$ and $\{2, 4, 5\}$ form a partition of Ω . It is typically the case that the structure of a player’s knowledge is represented by a collection of disjoint and exhaustive subsets of Ω , called a partition.

Definition 1. A *possibility correspondence* $P_i(\omega)$ for Ω is *partitional* if there is a partition of Ω such that for any $\omega \in \Omega$ the set $P_i(\omega)$ is the element of the partition that contains ω .

If P_i is a partition and if ω and ω' are two states in Ω , then when ω and ω' are in the same element of the partition, the decision maker cannot tell the difference between them. However, if ω and ω' are not in the same element of the partition, the decision maker can tell the two states apart.

⁵We examine issues of bounded rationality and noncommon priors below.

⁶Note that our model of knowledge is simply a generalization of the typical information set structure of extensive form games of incomplete information.

At this point we place restrictions on P_i to represent the types of properties, or axioms, we desire in the processing of information by players.

Definition 2. Let P_i be a possibility correspondence for individual i . We say

1. P_i is nondeluded if, for all $\omega \in \Omega$, $\omega \in P_i(\omega)$,
2. a player i knows that she knows [KTYK] if, for every $\omega' \in P_i(\omega)$, $P_i(\omega') \subseteq P_i(\omega)$,
3. a player i knows that she doesn't know if, for every $\omega \in \Omega$ and every $\omega' \in P_i(\omega)$, $P_i(\omega') \supseteq P_i(\omega)$.

These three properties are used to formalize the idea of rationality in knowledge. The first condition requires that a rational person always considers the true state of the world to be possible. The second condition requires that if any state an individual thinks is possible at the current state of the world were the true state, she would know at least her current knowledge, $P_i(\omega)$. That is, $P_i(\omega)$ cannot occur without the individual knowing that she knows it has occurred. Formally, this implies that if $K_i(E)$ is the event that i knows E , then $K_i(K_i(E)) \subseteq K_i(E)$. The final condition requires that players also know what they don't know, i.e., $\neg K_i(E) \subseteq K_i(\neg K_i(E))$. These three conditions ensure that $P_i(\omega)$ is consistent and that a player's possibility correspondence represents all that is knowable at each state.

To see an example of how rationality is related to the conditions of definition 2, let us consider a possibility correspondence that does not satisfy all three conditions. It is then easy to see why such a correspondence does not represent all that a rational player could know at a given state of the world, ω . Suppose that there are five states of the world, $\Omega = \{1, 2, 3, 4, 5\}$. Since a rational model should never have players place zero probability on the true state of the world, let us consider a correspondence $P_i(\omega)$ that satisfies nondeluded, but *not* know that you know and know that you don't know. In particular, let

$$P_i(\omega) = \{\omega - 1, \omega, \omega + 1\} \quad \text{if } \omega = \{2, 3, 4\},$$

$$P_i(\omega) = \{1, 2\} \quad \text{if } \omega = 1,$$

and

$$P_i(\omega) = \{4, 5\} \quad \text{if } \omega = 5.$$

Now suppose that $\omega = 2$, what could a rational player conclude? From $P_i(2)$ the player knows that the true state of the world is 1, 2, or 3. But player i also knows that the state is not 1, because if the state were 1 she would know that the state was not 3. So i can deduce that the true state is either 2 or 3. Moreover, if $\omega = 3$, i would know that the other possible states would be 2, 3, 4, but since she knows $\omega \neq 4$, ω cannot be 3. Therefore, i can deduce that

the true state is 2, and a rational player knows more than what is described by $P_i(\omega)$.

In fact, we can justify the use of partitional possibility correspondences in a model with rational actors because it is easy, if somewhat tedious, to show that P_i satisfies nondeluded, know what you know, and know what you don't know if and only if it is partitional. That is, a partitional possibility correspondence is the only internally consistent representation of a player's knowledge and, therefore, represents all that can be known by a rational decision maker at a given state of the world.⁷

Now to use this model of knowledge in equilibrium analysis, we also need to be able to say something about what players know about what others know, and what those others know that they know, etc. This is accomplished by considering events that are common knowledge. The concept of common knowledge was first explicitly described by Lewis (1969), and was later formalized by Aumann (1976) in terms of the *meet* of the information partitions for all the players at a state ω . Informally, an event F is *common knowledge* between players i and j , in state ω , if and only if ω is a member of every set in an infinite series such that, player i knows F , player j know F , player i knows player j knows F , player j knows player i knows player j knows F , etc.

Like common knowledge events, another important class of events is the *self-evident* event.

Definition 3. An event E is self-evident for a possibility correspondence P_i if and only if for all $\omega \in E$, $P_i(\omega) \subseteq E$.

In other words, an event E is self-evident if, for any state in E , a player knows E has occurred. Returning to our previous example where $\Omega = \{1, 2, 3, 4, 5\}$ and $P_i(\omega)$ induces a partition on Ω of $\{\{1, 3\}, \{2, 4, 5\}\}$, the event $\{1, 3\}$ is self-evident, but $\{1, 2, 3\}$ is not. The following useful fact is immediate. If P_i is nondeluded, then for a self-evident event E ,

$$E = \bigcup_{\forall \omega \in E} P_i(\omega).$$

Self-evident events are useful because of the following result, which states that an event being self-evident to *all* players is equivalent to it being common knowledge. The proof of the lemma follows from Proposition 3.5 of Rubinstein (1998).

Lemma 1. Suppose P_i is nondeluded for all i . An event F is common knowledge at a state ω if and only if there is an ω and a self-evident event E such that $\omega \in E \subseteq F$ for all P_i .

⁷Notice that information sets in extensive form games are always partitional.

We are often also interested in “public” events. A public event, unlike a private signal, is known to all players when it happens.⁸ Formally, we define a public event as follows:

Definition 4. *An event E is a public event if and only if, for all i , $K_i(E) = E$.*

Note that a public event is self-evident to all players; this equivalence is given in the following lemma.

Lemma 2. *If E is a public event, then for all i , E is self-evident.*

Proof: By definition of a public event $K_i(E) = \{\omega \in \Omega \mid P_i(\omega) \subseteq E\} = E$ for all i . Therefore, for all $\omega \in E$, $P_i(\omega) \subseteq K_i(E) = E$ and E is self-evident. ■

Now that we have specified a model of knowledge, we can talk about knowledge at a given state of the world and the decision to go to war.

Recalling that the state of the world is directly relevant to the question of which country will win a war, we define two functions, $p_1(\omega)$ and $p_2(\omega)$, that specify the probability that country 1 and 2 will win a war, given the true state of the world ω . Of course, $p_1(\omega) + p_2(\omega) = 1$ and $0 \leq p_i(\omega) \leq 1$ for all values $\omega \in \Omega$. Consider an arbitrary event E . If a country knows an event $E \subseteq \Omega$ has occurred, it can combine this information with the prior π via Bayes’s Rule to form a posterior belief about the value of p_i as follows:

$$E[p_i \mid E] = \frac{\sum_{\omega \in E} P_i(\omega)\pi(\omega)}{\sum_{\omega \in E} \pi(\omega)}. \quad (1)$$

From this expression, it is easy to verify that if $E[p_i \mid E'] \geq x$ and $E[p_i \mid E''] \geq x$ for disjoint sets of states E' and E'' , then $E[p_i \mid E' \cup E''] \geq x$.

It is equally likely that the negotiated settlement will depend on the underlying state of the world. We now define two additional functions, $r_1(\omega)$ and $r_2(\omega)$, that specify the bargaining outcome when the true state of the world is ω . Since bargaining is efficient, we assume that in each state $r_1(\omega) + r_2(\omega) = 1$. It is then immediate that countries’ beliefs regarding the outcome of the bargaining process will depend on their private information as well.

We represent the private information of country i by a possibility correspondence $P_i : \Omega \rightarrow 2^\Omega$, which we assume is partitional. Recall that $P_i(\omega)$ is the set of states that country i views as possible, given the true state ω . Given a true state ω , a country can combine its knowledge of $P_i(\omega)$ with the prior π and equation (1) to construct its

posterior belief about the probability it will win, $\hat{p}_i(\omega) = E[p_i \mid P_i(\omega)]$, and its expected payoff from bargaining, $\hat{r}_i(\omega) = E[r_i \mid P_i(\omega)]$. It is important to note that without additional assumptions or structure, it is certainly possible that $\hat{p}_1(\omega) \neq 1 - \hat{p}_2(\omega)$ or $\hat{r}_1(\omega) \neq 1 - \hat{r}_2(\omega)$ for some state (ω) . In this setting, mutual optimism occurs when $\hat{p}_1(\omega) > \hat{r}_1(\omega)$ and $\hat{p}_2(\omega) > \hat{r}_2(\omega)$.

In the rest of this section, we place the preceding informational assumptions in the context of a game and show that, in equilibrium, countries cannot have mutual optimism and therefore, war cannot occur as a result. Because the information structure is quite general and can capture many aspects of the strategic interaction, the description of the game is abstract.

Denote the set of actions for country i in some two-player strategic form game by A_i , with elements $\{a_i^1, \dots, a_i^k\}$. Depending on the choice of actions by both countries, the outcome of the game is either war or settlement. The expected payoff to war depends on the probability that a country will win, the utility of victory and defeat, and the inefficiencies present in fighting. We normalize the utility of countries to be 1 for victory in war and 0 for defeat, and we suppose there is a cost $c_i(\omega) > 0$ of fighting a war for country i . Thus the expected utility for country i of going to war is simply $\hat{p}_i(\omega) - \hat{c}_i(\omega)$, where $\hat{c}_i(\omega) = E[c_i \mid P_i(\omega)]$. On the other hand, the negotiation process provides an expected utility $\hat{r}_1(\omega)$ for country 1 and $\hat{r}_2(\omega)$ for country 2. It is important to note that we assume that neither the expected payoff to war or the expected outcome of negotiations depend on the choice of actions by the countries. That is, we exclude the possibility of gaining an advantage by surprise attack or making threats in order to gain bargaining leverage.⁹ However, we are completely general about how negotiations actually proceed and permit the outcome of the negotiation stage to depend on the private information of the two countries. In fact, the *revelation principle* tells us that any equilibrium of any choice of game for the negotiation stage can be mimicked by a mapping that takes states of the world into outcomes (Myerson 1979). This is because a game is defined as a mapping from players’ negotiation choices into outcomes, a strategy is defined as a mapping of player types into choices, and so without loss of generality we can define a new mapping, $r(\omega)$, that is a composition of the game form mapping and the strategy mapping. This new mapping takes states of the world into outcomes in the exact same way as an equilibrium to the underlying game. Our set up, therefore, lets us consider any distribution of the prize that countries may expect to get in the bargaining game that is an alternative to

⁸For a discussion of public events see Milgrom (1981).

⁹For more on this issue see the discussion in below.

war. Our bargaining protocol is consistent with previous work, like Powell (1999), that argues a peaceful settlement should be an equilibrium to a bargaining game and should represent, through that equilibrium, the underlying balance of power as captured by the state of the world. By construction, $r(\omega)$ covers both conditions.

To focus on the mutual optimism explanation of war, we assume that both countries must choose to “stand firm” for war to occur. Formally, this condition requires that for each country i , there is an action $\tilde{a}_i \in A_i$ such that, conditional on the opponent choosing to “stand firm,” the outcome is a settlement. In practice, if one country chooses to stand firm, the other country can stop a war by inducing the bargaining procedure instead. That is, war is an act of mutual consent and, by construction, also mutual optimism. As we have already discussed, if we drop this assumption and allow any single state to cause a war, the concept of war by mutual optimism loses meaning. Put simply, it is hard to understand what is *mutual* about mutual optimism if only one side’s expectations enter into the decision to fight. Moreover, Fey and Ramsay (2007a) show that in certain cases, if a single country can start a war, it is impossible to write down any game form that guarantees peaceful outcomes.

We now define strategies for each country. We reflect the fact that countries can condition their choice of action on their private information by defining a (pure strategy) strategy $s_i \in S_i$ as a function $s_i : \Omega \rightarrow A_i$ with the restriction that

$$P_i(\omega) = P_i(\omega') \Rightarrow s_i(\omega) = s_i(\omega').$$

This condition states that if a country cannot distinguish state ω from state ω' , then its action must be the same in both states.

Lastly, we discuss the event *war*. It follows from Lemma 2 that if the outcome of the game is a publicly observable war, then war is common knowledge whenever it occurs. Since strategies associate states with outcomes, we can now offer a rigorous definition of the statement that war is a public event given our information model. For a strategy profile (s_1, s_2) , let F be the set of states for which the outcome of the game is war. It follows that if war is publicly observable (for (s_1, s_2)), then the event F is a public event. If the event F is nonempty for a strategy profile (s_1, s_2) , we say that (s_1, s_2) is a strategy profile in which *war occurs*.

Results

Letting G denote any strategic form game of incomplete information that satisfies the preceding assumptions on

information structure, payoffs, and strategies. We are now prepared to state our main result.¹⁰

Theorem 1. *Suppose countries have a common prior, war is a public event, and P_i is partitional for $i = 1, 2$. Then there is no Bayesian-Nash equilibrium of G in which war occurs.*

Proof: Suppose not. That is, suppose that the strategy profile (s_1^*, s_2^*) is a Bayesian-Nash equilibrium in which war occurs. At state ω , war has an expected payoff to country i of $\hat{p}_i(\omega) - \hat{c}_i(\omega)$. By choosing action \tilde{a}_i , though, country i can ensure itself a payoff of $\hat{r}_i(\omega)$. Define the following two events:

$$O_1 = \{\omega \in \Omega \mid \hat{p}_1(\omega) \geq \hat{r}_1(\omega) + \hat{c}_1(\omega)\}$$

$$O_2 = \{\omega \in \Omega \mid \hat{p}_2(\omega) \geq \hat{r}_2(\omega) + \hat{c}_2(\omega)\}.$$

Here O_i is the event that country i will prefer war given that it knows the event $P_i(\omega)$ has occurred. At states outside O_i , country i will prefer to deviate to \tilde{a}_i . Thus, if (s_1^*, s_2^*) is a Bayesian-Nash equilibrium, the set of states for which the outcome of the game is war is $W \subseteq O_1 \cap O_2$. Moreover, since war occurs under this strategy profile, W is not empty; there exists an $\omega^* \in W$.

Since we assume war is a public event, the event W is a public event. By Lemma 2, W is self-evident to 1 and 2, and therefore

$$W = \bigcup_{\omega \in W} P_1(\omega) = \bigcup_{\omega \in W} P_2(\omega). \tag{2}$$

That is, as W is self-evident for nondeluded P_i , W is the union of $P_i(\omega)$ for all ω in W , and this is true for each player i . As the correspondence P_i is partitional, we can further write W as the union of disjoint sets $P_1(\omega)$, defined by some collection of states D^* with $D^* \subseteq W$. Since $D^* \subseteq W \subseteq O_1 \cap O_2$, we have $E[p_1 \mid P_1(\omega)] \geq E[r_1 \mid P_1(\omega)] + E[c_1 \mid P_1(\omega)]$ for every $\omega \in D^*$. That is, if each disjoint set $P_1(\omega)$ has conditional expectation $E[p_1 \mid P_1(\omega)]$ of at least $E[r_1 \mid P_1(\omega)] + E[c_1 \mid P_1(\omega)]$, then the conditional expectation over the union of these disjoint sets (i.e., $E[p_1 \mid W]$) is also at least $E[r_1 \mid W] + E[c_1 \mid W]$. That is, $E[p_1 \mid W] \geq E[r_1 \mid W] + E[c_1 \mid W]$. By a symmetric argument for player 2, $E[p_2 \mid W] \geq E[r_2 \mid W] + E[c_2 \mid W]$. Therefore

$$E[p_1 \mid W] + E[p_2 \mid W] \geq E[r_1 \mid W] + E[r_2 \mid W] + E[c_1 \mid W] + E[c_2 \mid W]. \tag{3}$$

But as $p_1(\omega) + p_2(\omega) = 1$, $r_1(\omega) + r_2(\omega) = 1$, and $c_i(\omega) > 0$ for all $\omega \in \Omega$, it follows from Bayes’s Rule that

¹⁰For simplicity, we state and prove our result using pure strategies, but the same result holds if we permit countries to use mixed strategies.

$E[p_1 | W] + E[p_2 | W] = 1$, $E[r_1 | W] + E[r_2 | W] = 1$, and $E[c_1 | W] + E[c_2 | W] > 0$. But then equation (3) yields a contradiction, which proves the result. ■

This theorem shows that there cannot be an equilibrium in which both sides think they are better off fighting, and as a result, go to war. The intuition underlying this theorem is as follows. Each country knows that the other is optimizing in equilibrium, knows when a war occurs and can deduce from the set of states, the prior probability, and the associated costs and benefits of each action for their opponent. Each player therefore knows that their opponent is willing to fight only in states where she is “likely to lose.” Knowing this, each player should condition her decision on this fact. As a result, the conjecture that the players’ strategies form an equilibrium where war is a public event would unravel just like in the dice game between Alice and Bob discussed above.¹¹ So the fact that wars are public events is inconsistent with inconsistent beliefs, even if leaders have private information about the likely outcome of conflict.

We conclude this section with some additional remarks regarding the theorem. First, our result does not require that bargaining be costless, but rather that fighting a war is more costly. This is because in any world where bargaining is costly, but less costly than war, we can always normalize the settlement outcome and think of the cost of war as the relative loss from fighting. Second, we note that Wittman’s (1979) classic mutual optimism argument is subsumed by a special case of Theorem 1. In particular, when $r_1(\omega) = r$ and $r_2(\omega) = 1 - r$ for all ω , we see that for any $r \in [0, 1]$ there cannot be war by mutual optimism. That is, it would be false to conclude that mutual optimism can create a situation where countries can find no agreement that would dissuade at least one country from wanting to fight in equilibrium. In fact, Theorem 1 implies any efficient division will do. Our result, however, is significantly stronger than this. Theorem 1 shows that there are no “optimistic” beliefs that allow a public war to be preferred to a peaceful settlement in equilibrium even when the probability of success in war *and* the value of a negotiated settlement are arbitrary functions that may depend on the state of the world. Moreover, as the mappings $r_1(\omega)$ and $r_2(\omega)$ can represent any equilibrium of any bargaining game one could imagine, our results do not depend on the extensive form of the negotiation phase reached after some country has opted for negotiations. Third, one may wonder how our assumption that the game is in strategic form influences the generality of

¹¹Note the $(r_1(\omega), r_2(\omega))$ negotiation subgame need not be reached in any equilibrium of the game for the result to hold.

our theorem. The answer stem from the fact that the only requirements on the possibility correspondences of the two sides are that they be partitional and that war is a public event. As a result, the players’ information may differ in a number of ways. One way they may differ is that one leader may know whether the other has chosen to stand firm or negotiate when making their decision. That is, our result applies equally to decisions made simultaneously and sequentially.¹² Finally, our assumption that war is a public event is used in the proof to show that the event W is a public event. For $\omega^* \in W$, by Lemma 1, this event is common knowledge at ω^* . So another way to state the conclusion of the theorem is that if war is common knowledge when it occurs, it cannot occur because of mutual optimism.

Mutual Optimism, Bounded Rationality, and War

While Theorem 1 is true for any mutual optimism game in which the decision makers rationally process information, one may wonder if the results depend on strictly rational learning. Can mutual optimism result in war if otherwise rational agents suffer from pathological misperception? In this section, we consider a class of games where, again, two countries are choosing whether to fight a war or resolve the dispute by some other means. Here, we show that even if players’ information processing suffers from cognitive biases, war still may not be possible in an equilibrium of a mutual optimism game. In particular, even if both players ignore “bad news” or are inattentive, then war cannot occur because of mutual optimism.

Processing Errors and Bounded Rationality

When it comes to information processing, a rational Bayesian may be able to deduce much more information from a “signal” than the signal carries at face value. For example, consider a world where there are two possible states, $\{a, b\}$. Suppose that when the true state is a it is brought to the player’s attention that the state is in fact a . However, when the true state is b nothing is brought to the player’s attention. In this a situation a rational Bayesian can always deduce the true state of the world. When the state is a , the player is informed of that fact and knows it is a . When the state is b , the player knows that if the state were a she would have been told, but since she was not, the state must be b . The rational Bayesian, like Sherlock Homes, learns from the dog that does not bark.

¹²For a more detailed discussion of this point see the next section.

There are, however, many cases in which we think that decision makers, particularly the leaders of countries, may not be processing information rationally. Consider the information processing errors found in the psychological international relations literature (Jervis 1976; Jervis, Lebow, and Stein 1985). For example, a decision maker who has many responsibilities may face a volume of information that induces flaws in their learning. In particular, such a decision maker may not update their beliefs when the state of the world is not explicitly brought to their attention. So while they may learn that the state is a when there is an explicit signal to indicate that is so, they may not deduce that the state must be b in the absence of a signal that the state is a . This error may occur because of a flaw in human psychology or it could be an information shortcut that allows decision makers to deal with a world far more complex than the two state example above.

Alternatively, due to what Jervis, Lebow, and Stein (1985, p. 4) call *motivated bias*, a player's knowledge may be partly a matter of choice. So given that some people have strong predispositions to believe certain things to be true, this may prevent them from recognizing new information inconsistent with their world view. That is, sometimes decision makers may consciously, or subconsciously, choose to ignore unpleasant information.

Next we consider a game with players whose information processing is flawed in ways consistent with the learning processes described above. A common component of these cognitive biases is that the player's information processing allows them to learn from new information in some states of the world, but not in others. To capture this idea formally, we define a new restriction on the players' possibility correspondences, P_i . In particular, while we still assume P_i is nondeluded, we now allow players to "ignore" or "throw out" information at a given state of the world that would be known to a fully rational Bayesian. To allow such pathologies, we must allow for the possibility that some information sets are nested within, or subsets of, other information sets. That way we can capture the idea that inattentive decision makers do not deduce that the state of the world is a by lack of a signal that it is a or that decision makers might choose to ignore the deeper implications of the information in front of them. To do this, we replace the *know that you don't know* and the *know that you know* conditions with the requirement that the players' possibility correspondences are *nested*.¹³

¹³For more on decision-theoretic approaches to bounded rationality in models of knowledge, see Geanakoplos (1989) or Rubinstein (1998) for game theoretic approaches. The concept of nestedness is taken from Geanakoplos (1989).

Definition 5. A player's possibility correspondence is nested if for all $\omega, \omega' \in \Omega$, either (1) $P_i(\omega) \cap P_i(\omega') = \emptyset$ or (2) $P_i(\omega) \subseteq P_i(\omega')$ or (3) $P_i(\omega') \subseteq P_i(\omega)$.

Possibility correspondences that satisfy nondeluded and nestedness represent a generalization of rational learning. That is, a decision maker with a nested possibility correspondence may process information in a rational way or she may ignore new information at a number of different states. Such a formalization is consistent with many forms of bias, because it is agnostic to the reason information is ignored. Players could fail to learn in some states because acquiring information is costly, because they are inattentive, or because they would rather not think about the implications of the information in front of them.

Results

As we have seen, a useful approach to model the limitations of a player's ability to process information is to consider a player's information partition that satisfies nondeluded and nested. Together, these conditions are weaker than the three conditions for a rational partition, yet are still sufficient to exclude optimistic war. In general, we now have:

Theorem 2. Suppose countries have a common prior, war is a public event, and P_i is nondeluded and nested for $i = 1, 2$. Then there is no Bayesian-Nash equilibrium of G in which war occurs.

Proof: Suppose not. That is suppose that (s_1^*, s_2^*) is a Nash equilibrium in which war occurs with a generalized information partition that satisfies nondeluded and nestedness.

As in Theorem 1, define the following two events:

$$O_1 = \{\omega \in \Omega \mid \hat{p}_1(\omega) \geq \hat{r}_1(\omega) + \hat{c}_1(\omega)\}$$

$$O_2 = \{\omega \in \Omega \mid \hat{p}_2(\omega) \geq \hat{r}_2(\omega) + \hat{c}_2(\omega)\}.$$

Recall, O_i is the event that country i will prefer war given that it knows the event $P_i(\omega)$ has occurred. At states outside O_i , country i will prefer to deviate to \bar{a}_i . Thus, if (s_1^*, s_2^*) is a Bayesian Nash equilibrium the set of states for which the outcome of the game is war is $W \subseteq O_1 \cap O_2$. Moreover, since war occurs under this strategy profile, W is not empty; there exists an $\omega^* \in W$.

Since we assume war is a public event, the event W is a public event. By Lemma 2, and the fact that a self-evident event is well defined for any nondeluded possibility correspondence, W is self-evident to 1 and 2, and therefore

$$W = \bigcup_{\omega \in W} P_1(\omega) = \bigcup_{\omega \in W} P_2(\omega). \quad (4)$$

Let $\bar{P}_i(\omega)$ be the set with the largest cardinality such that $\omega \in P_i(\omega')$ for some $\omega' \in W$. By nestedness $\bar{P}_i(\omega)$ is unique and $\bar{P}_i(\omega) \subseteq W$ for all $\omega \in W$. Because W is self-evident,

$$W = \bigcup_{\forall \omega \in W} \bar{P}_i(\omega)$$

and for all $\omega, \omega' \in W$, either $\bar{P}_i(\omega) = \bar{P}_i(\omega')$ or $\bar{P}_i(\omega) \cap \bar{P}_i(\omega') = \emptyset$. Therefore, country i 's nondeluded, and nested $P_i(\omega)$ induces a $\bar{P}_i(\omega)$ partition on W . We can then write W as the union of disjoint sets $\bar{P}_i(\omega)$, defined by some collection of states \hat{D}^* all contained in W , i.e., $\hat{D}^* \subseteq W$. The result then follows as in Theorem 1. ■

Theorem 2 show that for some plausible types of “boundedly rational” actors, mutual optimism cannot be the reason two decision makers go to war. In the appendix, we give two examples that show that the conditions in Theorem 2 cannot be relaxed. That is, for the impossibility of war by mutual optimism, nondeluded and nested are minimally sufficient or “tight” conditions. In particular, we show that if either condition in the theorem fails, then there exist examples of mutual optimism games in which both countries choose to fight in equilibrium because of their private information.

As mentioned above, considering information structures that relax the requirements of strict Bayesian rationality can help us understand just how general our mutual optimism result is. On the one hand, the analysis in this section shows that the mutual optimism result is not fragile. Clearly, some departure from rational Bayesian learning is acceptable and consistent with our results. In particular, if decision makers sometimes ignore unpleasant information or behave as if they have imperfect memory, then our result survives. On the other hand, the fact that Theorem 2 also contains the minimally sufficient conditions for our result to hold means that our analysis gives a criteria by which one can determine the minimum level of rationality needed to rule out war by mutual optimism.

Generalized Partitions and Noncommon Priors

It should be clear by now that the common priors assumption, while standard, plays a critical role in our analysis.¹⁴ Indeed, if we permit countries to hold different prior beliefs that are not explained by information, then it may be possible to find pairs of prior beliefs that result in war as an equilibrium outcome in our setting.

¹⁴For a discussion of some of the issues surrounding the common priors assumption see Gul (1998), Aumann (1998), Smith and Stam (2004), Fey and Ramsay (2006), Smith and Stam (2006).

In this section, prompted by Smith and Stam (2004) who explore a model where players do not have common priors and demonstrate that one possible consequence is war, we investigate what our results imply regarding the possibility of war by mutual optimism in an environment with noncommon priors. We show that there is an equivalence between common-prior models with boundedly rational players and models with noncommon priors and fully rational players. In fact, we show that our analysis implies a set of conditions on the “noncommonness” of priors required to result in war by mutual optimism.

Specifically, we give two propositions that identify the direct connection between games in which countries have a common prior and are boundedly rational in processing of information and games where countries have noncommon priors. The two propositions show that the characterization of minimally sufficient conditions on decision makers' information structures described in Theorem 2 have direct analogs in the noncommon priors framework. That is, our characterization of constraints on the possibility correspondences of decision makers imply bounds on the noncommon priors that can be held by rational Bayesians in a analogous noncommon priors model.¹⁵

To begin, we need to define a notion of strategic equivalence between any two games.

Definition 6. A game $G' = \langle \Omega', A, \mathbb{P}', u', \Pi' \rangle$ is strategically equivalent to a game $G = \langle \Omega, A, \mathbb{P}, u, \Pi \rangle$ if there is a onto function $\varphi : \Omega' \rightarrow \Omega$ such that for every state $\omega' \in \Omega'$,

$$E u'_i(a_i, a_{-i} | P'_i(\omega')) = E u_i(a_i, a_{-i} | P_i(\varphi(\omega'))). \quad (5)$$

This definition of strategic equivalence requires that the optimal actions, given beliefs, in one game be the same as in the other game once we properly “rename” the state space.

For our first proposition, we consider a game with non-partitional information and show that there exists an strategically equivalent analogue in the noncommon priors framework.

¹⁵Similar results are derived for decision-theoretic equivalence and equivalence of correlated equilibria of games with generalized partitions and correlated equilibria of games with partitional information structures and noncommon priors by Brandenburger, Dekel, and Geanakoplos (1992) and Geanakoplos (1989). Our result differs, however, because we use a different equivalence relation and because the analysis implies equivalence for the whole equilibrium correspondence. The idea and proof strategies, however, are inspired by Brandenburger, Dekel, and Geanakoplos (1992) and Geanakoplos (1989). The detailed technical conditions on priors implied by our results can be found in Shin (1989) and are equivalent to those restrictions for nonspeculation found in Morris (1991, 1994). Therefore, we do not explicitly characterize them here.

Proposition 1. For any finite game $G = \langle \Omega, \pi, A, u, \mathbb{P} \rangle$ with nonpartitional information structure and a common prior, there exists a game $G' = \langle \Omega', \Pi, A, u', \mathbb{P}' \rangle$ that has a common state space Ω' , noncommon priors $\Pi = \{\pi_1, \pi_2\}$, utility functions u' , a partitional information structure \mathbb{P}' , and is strategically equivalent to G .

Proof: Let G be a game with state space Ω , a common prior π , action space $A = \{A_1, A_2\}$, a utility function for each player $u_i : \Omega \times A_1 \times A_2 \rightarrow \mathbb{R}$, and a nonpartitional information structure.

Let $\mathbb{R}_i = \{R \subseteq \Omega : \exists \omega \in \Omega \text{ with } R = P_i(\omega)\}$ and let $\Omega' = \mathbb{R}_1 \times \mathbb{R}_2 \times \Omega$. An element of Ω' is given by $\omega' = (R_1, R_2, \omega)$ and let $P'_i(\omega') = P'_i(R_i, R_{-i}, \omega) = \{R_i\} \times \mathbb{R}_{-i} \times \Omega$. By construction, these P'_i partition Ω' for both players. Let $\varphi : \Omega' \rightarrow \Omega$ be an onto function such that $\varphi(R_1, R_2, \omega) = \omega$.

Next define the noncommon prior

$$\pi'_i(\omega') = \pi'_i(R_1, R_2, \omega) = \begin{cases} \frac{1}{|\mathbb{R}_i|} \frac{\pi(\omega)}{\sum_{\tilde{\omega} \in P_i(\omega)} \pi(\tilde{\omega})} & \text{if } P_1(\omega) = R_1 \text{ and } P_2(\omega) = R_2, \\ 0 & \text{otherwise.} \end{cases}$$

It follows that $\pi'_i(\omega' | P'_i(\omega')) = \pi'_i(R_1, R_2, \omega | \{R_i\} \times \mathbb{R}_{-i} \times \Omega)$ is given by

$$\pi'_i(\omega' | P'_i(\omega')) = \begin{cases} \frac{\pi(\omega)}{\sum_{\tilde{\omega} \in P_i(\omega)} \pi(\tilde{\omega})} & \text{if } P_1(\omega) = R_1 \text{ and } P_2(\omega) = R_2, \\ 0 & \text{otherwise.} \end{cases}$$

Finally, notice that if we let

$$u'_i(a_i, a_{-i}, \omega') = u_i(a_i, a_{-i}, \varphi(\omega')) = u_i(a_i, a_{-i}, \omega),$$

then

$$\begin{aligned} E u'_i(a_i, a_{-i} | P'_i(\omega')) &= \sum_{\tilde{\omega}' \in P'_i(\omega')} \pi'_i(\tilde{\omega}' | P'_i(\omega')) u'_i(a_i, a_{-i}, \tilde{\omega}') \\ &= \sum_{\{\tilde{\omega}' \in P'_i(\omega') | \pi'_i(\tilde{\omega}') > 0\}} \frac{\pi(\tilde{\omega})}{\sum_{\tilde{\omega} \in P_i(\omega)} \pi(\tilde{\omega})} u_i(a_i, a_{-i}, \varphi(\tilde{\omega}')) \\ &= \sum_{\tilde{\omega} \in P_i(\omega)} \pi(\tilde{\omega} | P_i(\omega)) u_i(a_i, a_{-i}, \tilde{\omega}) \\ &= E u_i(a_i, a_{-i} | P_i(\omega)) = E u_i(a_i, a_{-i} | P_i(\varphi(\omega'))). \end{aligned}$$

This new game has a common states space, a non-common prior, and a partitional information structure, and is strategically equivalent to our original game with common priors and a nonpartitional information structure. ■

Our second proposition is a converse of the first. That is, we consider a game with non-common priors and partitional information and show that there exists a strategically equivalent game with common priors and a generalized information structure.

Proposition 2. For any finite game $G = \langle \Omega, \Pi, A, u, \mathbb{P} \rangle$ with a partitional information structure and a noncommon prior ($\Pi = \{\pi_1, \pi_2\}$), there exists a game $\hat{G} = \langle \hat{\Omega}, \hat{\pi}, A, \hat{u}, \hat{\mathbb{P}} \rangle$ that has a common state space $\hat{\Omega}$, a common prior $\hat{\pi}$, utility functions \hat{u} , and a nonpartitional information structure ($\hat{\mathbb{P}}$), and is strategically equivalent to G .

Proof: Let G be a game with state space Ω , noncommon priors π_i , action space $A = \{A_1, A_2\}$, a utility function for each player $u_i : \Omega \times A_1 \times A_2 \rightarrow \mathbb{R}$, and a partitional information structure.

Let $\hat{\Omega} = \Omega \times \{1, 2\}$ and define a function $\varphi : \hat{\Omega} \rightarrow \Omega$ such that $\varphi(\omega, k) = \omega$. Let player i 's possibility correspondence be

$$P'_i(\omega, k) = \{(\tilde{\omega}, \tilde{k}) \in \hat{\Omega} : \tilde{\omega} \in P_i(\omega) \text{ and } \tilde{k} = i\}.$$

Let $\pi^*(\omega, i) = \pi_i(\omega | P_i(\omega))$ for every $(\omega, i) \in \hat{\Omega}$ and define the common prior $\hat{\pi}$ by

$$\hat{\pi}(\omega, i) = \frac{\pi^*(\omega, i)}{\sum_{(\tilde{\omega}, \tilde{k}) \in \hat{\Omega}} \pi^*(\tilde{\omega}, \tilde{k})}.$$

Finally, if we set $\hat{u}_i(a, a_{-i}, (\omega, i)) = u_i(a_i, a_{-i}, \varphi(\omega, i)) = u_i(a_i, a_{-i}, \omega)$, it is easy to verify that \hat{G} is strategically equivalent to G . ■

These two propositions show that the approach taken in this paper that explores boundedly rational information processing with common priors is strategically equivalent to an approach where information is processed rationally, but decision makers have noncommon priors. The real difference between these two approaches is then largely one of taste. In the noncommon priors framework, it is assumed that for reasons not related to private information or the way it is processed, decision makers “enter the game” with different beliefs. The generalized information structure approach, on the other hand, assumes that differences in beliefs are the consequence of private information and the way that information is processed. As a result, Theorem 2 gives a characterization of sufficient conditions for our optimism result to hold and a substantive interpretation of the limits in the rational decision process necessary for the result to fail. The non-common priors framework can provide—it turns out—analytically equivalent conditions for the failure of our result, but does not provide a theoretical mechanism for its origin. In light of this fact, we prefer the generalized

partition approach, but point out that, formally, the two perspectives are equivalent.

Discussion

Formalizing the mutual optimism hypothesis and using assumptions designed to test this hypothesis, we see that if war is a public event when it happens, it cannot occur between rational opponents because of mutual optimism. Moreover, when we consider some intuitive forms of “bounded rationality” in the way players learn, the result is robust. This result is somewhat surprising, given the existing work on the subject. In this section we discuss some specific aspects of our model as well as its implications and limitations for the study of international conflict.

Mutual optimism is one of several ways that private information can lead to war. Our approach in this paper is designed to serve as a test of mutual optimism as a valid explanation of war in a game-theoretic rationalist framework. As we have argued, in order to isolate mutual optimism as a cause of war, we must insure that when mutual optimism is not present, war does not occur. Otherwise, we cannot be sure if it is mutual optimism or some other aspect of private information that explains the conflict. Our game-theoretic setting is thus designed to create conditions in which there is a clear link between mutual optimism and war.

The reasoning behind our modeling choices also helps us to understand why mutual optimism is not a valid rationalist explanation for war, even though others have argued that it is. Intuitively, the mutual optimism explanation requires that, when the war starts, both sides believe they will prevail. Another way to say this is it must be that neither side prefers to settle rather than go to war. Clearly, if, in a model, an equilibrium (with war) exists in which one side wants to settle but is not able to, then this war was not caused by mutual optimism. Thus, war can be validated as a consequence of mutual optimism only in a model in which an equilibrium with war exists even though *either* side can deviate to negotiate some settlement. This is precisely what our main assumptions imply.¹⁶

Taken together, our assumptions are based on the idea that mutual optimism, as a cause of war, is a distinct mechanism and is, in fact, different from rationalist explanations for war that focus on the preferences players

have over actions that risk war, but have higher peaceful payoffs, and “sure things” that guarantee peace at a price. To illuminate more precisely what is and is not covered by our result on mutual optimism, it is useful to consider a simple bargaining model in which war is possible. Suppose there are two countries *A* and *B*. Country *A* is uncertain about the type of country *B*, which is either weak or tough. In the first stage, *A* makes a low demand or a high demand and in the second stage, *B* either accepts the demand or rejects it, in which case war results. Suppose a weak type of country *B* will lose a war with high probability and a tough type will win with high probability. In this model, it is straightforward to find parameters such that both types of *B* will accept the low demand but, in equilibrium, *A* makes a high demand that will be accepted by the weak type but rejected by the tough type. Therefore, war is a possible outcome of this equilibrium. Intuitively, the first mover in this “risk reward” equilibrium knows that with some probability she will face a tough type and be likely to lose, but that this risk is offset by the reward of achieving a valuable concession from a weak type. Thus, war arises because there is no way for *A* to react to the private information of *B*.

This sort of model is informative about the nature of optimal demands in the shadow of war. Yet, the occurrence of war in this model is related not just to the assumptions concerning the countries’ information, but also the nature of the bargaining process. In this kind of model, the fact that war occurs is a consequence of the commitment value of country *A*’s initial offer. It is not obvious that such a commitment to an initial demand is either reasonable or realistic. Moreover, the structure of the interaction implicitly assumes that all the bargaining power lies with country *A*. As shown by Leventoglu and Tarar (2006), however, if we relax the assumption that country *A* has all the bargaining power, then reasonable equilibria exist in which sufficiently patient players do not behave in the way suggested by our simple example.¹⁷

For our purposes, however, what is important in this simple example is that it is not that case that *both* sides believe they will prevail. In particular, at the terminal node in which war occurs on the equilibrium path, *A* can infer from the rejection of its demand that *B* is the tough type and will revise its belief about its probability of winning downwards. It is easy to show that after this updating, *A* would prefer to settle rather than go to war, but this option is not available at this point in the game. Thus, it is clear

¹⁶From a technical perspective, our assumptions serve to insure that a profitable deviation from a war equilibrium is present when there is no mutual optimism.

¹⁷This result, along with other work that drops the assumption that a decision maker is committed to their proposal after their bargaining partner has chosen to accept or reject it, suggests that the “risk-reward” result is driven to a large extent by (somewhat arbitrary) assumptions regarding the extensive form of the bargaining process.

that this example of war does not satisfy our definition of mutual optimism that, at the instant before the war starts, both sides believe they will prevail.

As we show in a companion paper (Fey and Ramsay 2007b), this is a specific instance of a general phenomena. Namely, in a large class of games including games in which war is a unilateral act rather than a mutual act, there will always be states of the world in which war occurs but one of the two sides does not believe it will prevail. In fact, if the correct model of war is one in which any single country can start a war, the presence or absence of mutual optimism is irrelevant and, therefore, not a coherent rationalist explanation of war.

Reflecting on this example, we can give an intuitive statement of our main result in the following way. If it is common knowledge that countries are going to fight, and these countries have a “hot line” available, then at least one side will always want to make a call and a proposal that will be accepted and avoid the war. That is, our result applies to a situation where countries can discuss war before making it, but after a proposal has been made and rejected. In situations where a firm offer is made that, if rejected, leads to certain war, equilibria with war can exist, but not because of mutual optimism; one side would prefer to settle but is locked into a war by the structure of the extensive form.¹⁸

Another important assumption in our model is that the settlement outcome does not depend on how the negotiation stage is reached. That is, while we assume that the underlying state of the world and the players’ information about that state are important for determining the outcome of the negotiated settlement, that outcome does not depend on which of the players pursued the peaceful settlement instead of a war. Once again, it is important to emphasize that we make this assumption in order to guarantee that the sole cause of war in our model is mutual optimism and not some alternative cause. The following simple example of a game that does not satisfy this assumption (but does satisfy all other assumptions of our model) will illustrate this point. In this example, there is complete information and therefore no possibility of mutual optimism, and yet war is the *unique* equilibrium outcome.¹⁹ Suppose each side is equally like to win a war, the cost of war is not too large ($c < 1/2$), and both sides simultaneously choose whether to stand firm or negotiate. If both sides choose to stand firm, war results with

payoff $1/2 - c$; otherwise a negotiated settlement results, the payoffs of which depend on the actions. If both sides choose negotiations, they both receive payoff $1/2$, and otherwise the settlement gives payoff 1 to the country who stands firm and 0 to the country who chooses negotiations. In this game, it is easy to see that war is the unique Nash equilibrium outcome as both sides have a strictly dominant strategy to stand firm.²⁰ It is just as easy to see that such a war has nothing to do with mutual optimism. In this example, war occurs because the structure of the settlement payoffs gives each side an incentive to take an aggressive posture in choosing actions, but these aggressive actions have nothing to do with any underlying uncertainty. While it is possible that, in the real world, strategic situations such as these may arise, and as a result wars may occur, the cause of such a war is not mutual optimism. Put another way, the mutual optimism explanation for war is a purely informational story and is not a story about “capitulation” or posturing or bargaining concessions. Our assumption about the settlement outcome serves to insure that these other causes of war are not present in our model so that we can focus on the explanatory power of the mutual optimism explanation.²¹

Finally, we should comment on two of the assumptions we make about our informational structure: that the information partitions are themselves common knowledge, and that the players share a common prior over the states of the world. The first assumption is easily justified by the fact that war is clearly a public event, and thus its occurrence is commonly known. The second assumption has been addressed above. As we have argued, the common prior assumption makes explicit how differences in beliefs are grounded in differences in information. For this reason, we prefer to start with common priors and consider generalized information structures. Nevertheless, Propositions 1 and 2 demonstrate that the generalized partitions approach and the noncommon priors approach are analytically equivalent and the choice of which direction to push when departing from traditional rationality assumptions is largely a matter of taste. What is settled by our results, however, is the set of constraints—interpreted in the generalized information structure or in the noncommon priors world—that qualify our result concerning mutual optimism. It should not be surprising that as we allow our decision makers to become “less rational” that our optimism result eventually breaks down.

¹⁸In this way, our arguments have the flavor of renegotiation-proofness (Bernheim, Peleg, and Whinston 1987; Laffont and Martimort 1997).

¹⁹Formally, in our model of knowledge, the case of complete information is simply a model with a single state of the world.

²⁰In fact, this game is a standard Prisoner’s Dilemma.

²¹We thank an anonymous reviewer for bringing this important point to our attention.

What is more surprising, however, is how far we can go down that road before the optimism result fails.

In sum, the results of our model are general, but not universal. Our results are not universal in the sense that we only address the use of mutual optimism as an explanation of war and have little to say about other mechanisms by which war occurs or other distinct rationalist explanations for war. However, our results are general in that the assumptions underlying the model are not heroic; the construction is consistent with, and even subsumes, arguments found in the informal and decision theoretic literature on mutual optimism and war; and our results hold for a number of different game forms and less than rational players.

Conclusion

We have established two main findings by formalizing the argument that countries fight wars due to the mutual optimism. First, if war is a public event among rational actors, war cannot occur because of mutual optimism. Second, relaxing the conditions in our model of knowledge shows that this result is robust to imperfect, or “boundedly rational,” learning. We have also shown that there is a formal equivalence between our bounded rationality approach to learning and the approach to modeling disagreement that assumes subjective (noncommon) priors. These claims are based on our demonstration that, for a general class of games where states contemplate war, there are no Bayesian-Nash equilibria with war as an outcome. This analysis also helps explain why models with private information appear to show that mutual optimism leads to war. There are other ways in which private information can cause war, other than mutual optimism. In order to eliminate these other causes, we impose our assumptions which serve to insure that in cases without mutual optimism, war cannot arise in equilibrium. In this way, our result that war cannot occur in equilibrium implies that mutual optimism is not a valid rationalist explanation for war.

The modeling approach taken in our analysis is consistent with a broader research program that looks to explain the occurrence of inefficient wars from a rational choice perspective. Clearly, such a framework already limits the set of theoretically consistent explanations for the causes of war. The theoretical landscape, however, is still dotted with many theories of war initiation that can rightly be called rationalist explanations for war. Our result shows that one prominent explanation, war by mutual optimism, is not a coherent and internally consistent theory of war within the rationalist framework.

Alternatively, our analysis leaves open many other rationalist explanations as possible avenues for further research. In particular, explanations for war that depend on the inability of countries to commit to agreements as there are shifts in the distribution of power, the imperfect indivisibility of disputed prizes, and incomplete information mechanisms that are not associated with mutual optimism, are all viable rationalist explanations and should continue to be a focus of theoretical and empirical research.

To the extent that our theoretical results have empirical content, they ask us to reconsider some narratives describing the onset of war, such as those presenting stylized facts surrounding the start of World War I, and to reinterpret the competing historical accounts in light of the internal inconsistency of the mutual optimism argument. Was the key mechanism causing these wars the onset of mutual optimism among groups of otherwise rational decision makers? Should we believe that decision-makers in Germany were optimistic concerning the outcome of a war with France and Russia? Or should we believe, as Moltke wrote, that the prospect of a conflict with Russia in a year or two left “. . . no alternative but to fight a preventative war so as to beat the enemy while we [Germany] could still emerge fairly well from the struggle” (Wohlforth 1987, 362) and thus understand the war as prompted by commitment problems, rather than mutual optimism? In the presence of multiple viable explanations for war, each with supporting historical evidence, we would suggest that those explanations which are consistent, given their overarching framework, should be given precedent over those that are incoherent. By this criteria our results suggest that, when it comes to mutual optimism, we should look elsewhere for our causes of war.

Appendix

In this section, we show by means of examples that the conditions in Theorem 2 cannot be relaxed. In particular, we will show that if any one of nondeluded, know that you know, or nestedness is violated, there are games in which war occurs due to mutual optimism.

The easiest case is a possibility correspondence that does not satisfy the nondeluded criterium. Any such correspondence would allow a decision maker to rule out the true state of the world. Suppose that at the some state of the world country 1 was sure to lose a war with country 2. If the deluded decision maker’s possibility correspondence ruled out that state, and instead put high probability on a state where country 1 would surely win, then clearly war by mutual optimism could occur.

The crucial role of the nestedness criteria are is not so obvious. First, consider a simple example of a game where the possibility correspondence satisfies nondeluded and know that you know, but not nestedness. In this example, let $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4\}$, $r(\omega) = r$ for all $\omega \in \Omega$, and let $c(\omega) = c$ in all states, with the restriction that $c > 0$. Assume $p(\omega)$ is a simple characteristic function that equals 1 if country 1 wins the war in state ω and 0 otherwise. In particular, suppose that country 1 wins the war in even numbered states and country 2 wins in odd numbered states. As before, normalize the value of the prize under dispute so that the utility of winning the war is $1 - c$ and losing is $-c$.

Now let country 1's possibility correspondence be $P_1(\omega_1) = \{\omega_1, \omega_2\}$, $P_1(\omega_2) = \{\omega_2\}$, $P_1(\omega_3) = \{\omega_2, \omega_3\}$, and $P_1(\omega_4) = \{\omega_4\}$ and let $P_2(\omega_1) = P_2(\omega_2) = P_2(\omega_3) = \{\omega_1, \omega_2, \omega_3\}$ and $P_2(\omega_4) = \{\omega_4\}$. Country 1's possibility correspondence then violates nestedness, but satisfies know that you know and nondeluded. Moreover, the event $E = \{\omega_1, \omega_2, \omega_3\}$ is a public event. Finally, assume that the countries have a common prior with $\pi(\omega_1) = \pi(\omega_3) = 1/4$, $\pi(\omega_2) = 3/8$ and $\pi(\omega_4) = 1/8$. Now suppose that country 1 and country 2 play strategies such that they both fight at $\omega_1, \omega_2, \omega_2$ and don't fight at ω_4 . Such a strategy implies that war is a public event and on the eve of conflict war is self-evident to both countries. Can such a strategy profile be a Bayesian-Nash equilibrium to this game with nonpartional information structures?

Equilibrium requires that, for country 1, the expected utility for war in each state where she chooses to fight must give her a higher expected payoff than a settlement. Given our conjectured strategy, and the beliefs implied by $P_1(\omega)$, the expected utility for fighting in state ω_2 for country 1 is greater than that of a settlement if and only if $1 - c \geq r$. Similarly, the expected utility of war for country 1 in ω_1 and ω_3 must also be greater than the value of a settlement in those states, implying $3/5 - c \geq r$. We can thus conclude that country 1 will play the conjectured strategy in equilibrium if $3/5 - c \geq r$.

For country 2 to play the above strategy, incentive compatibility requires that, given its information ($P_2(\omega)$), $r \geq 4/7 + c$. Taken together, we can then conclude that if r satisfies $4/7 + c \leq r \leq 3/5 - c$, then the configuration of preferences and information structure in this example imply war can occur by mutual optimism if the possibility correspondence satisfies nondeluded, known that you know, but not nestedness.

On the other hand, if country 1's possibility correspondence satisfied nestedness, there would be no equilibrium with war in this example. To see this, change the above example by allowing country 1's possibility corre-

spondence to be nested. Changing nothing else about the game, let $P_1(\omega_1) = \{\omega_1, \omega_2\}$, $P_1(\omega_2) = \{\omega_2\}$, $P_1(\omega_3) = \{\omega_1, \omega_2, \omega_3\}$, and $P_1(\omega_4) = \{\omega_4\}$ and consider player 1's incentive compatibility constraint for fighting a war when the state is ω_3 . Given this nested partition, incentive compatibility requires that $r \leq 3/7 - c$. From the above, we know that country 2 will fight when the true state is ω_1, ω_2 , or ω_3 if and only if $r \geq 4/7 + c$. Clearly these two conditions cannot hold at the same time and war cannot occur in equilibrium at state ω_3 . If country 1 were only willing to fight at ω_1 or ω_2 , a similar calculations shows that there is no feasible value of r such that country 2, given $P_2(\omega_1)$ and $P_2(\omega_2)$ respectively, would prefer war to a settlement. Similar arguments hold for other nested and nondeluded possibility correspondences for player 1 in this example, implying there would be no war in any equilibrium of this example if nestedness was satisfied.

References

- Aumann, Robert J. 1976. "Agreeing to Disagree." *The Annals of Statistics* 4(6): 1236–39.
- Aumann, Robert J. 1998. "Common Priors: A Reply to Gul." *Econometrica* 66(4): 929–38.
- Bernheim, B. Douglas, Bezalel Peleg, and Michael D. Whinston. 1987. "Coalition-proof Nash Equilibria I, Concepts." *Journal of Economic Theory* 42(1): 1–12.
- Blainey, Geoffrey. 1988. *The Causes of War*. New York: The Free Press.
- Brandenburger, Adam, Eddie Dekel, and John Geanakoplos. 1992. "Correlated Equilibrium with Generalized Information Structures." *Games and Economic Behavior* 4(2): 182–201.
- Bueno de Mesquita, Bruce, and David Lalman. 1992. *War and Reason: Domestic and International Imperative*. New Haven: Yale University Press.
- Bueno de Mesquita, Bruce, James D. Morrow, and Ethan R. Zorick. 1997. "Capabilities, Perception, and Escalation." *American Political Science Review* 90(1): 15–27.
- Downs, George W., and David M. Rocke. 1994. "Conflict, Agency, and Gambling for Resurrection: The Principle-Agent Problem Goes to War." *American Journal of Political Science* 38(2): 362–80.
- Fearon, James D. 1994. "Signaling versus the Balance of Power and Interests: An Empirical Test of a Crisis Bargaining Model." *Journal of Conflict Resolution* 38(2): 236–69.
- Fearon, James D. 1995. "Rationalist Explanations for War." *International Organization* 49(3): 379–414.
- Fearon, James D. 1998. "Bargaining, Enforcement, and International Cooperation." *International Organization* 52(2): 269–305.
- Fey, Mark, and Kristopher W. Ramsay. 2006. "The Common Priors Assumption: A comment on *Bargaining and the Nature of War*." *Journal of Conflict Resolution* 20(4): 607–13.

- Fey, Mark, and Kristopher W. Ramsay. 2007a. "Uncertainty and Incentives in Crisis Bargaining Games." Typescript. University of Rochester.
- Fey, Mark, and Kristopher W. Ramsay. 2007b. "Uncertainty, Optimism, and War." Typescript. University of Rochester.
- Filson, Darren, and Suzanne Werner. 2002. "A Bargaining Model of War and Peace." *American Journal of Political Science* 46(4): 819–38.
- Geanakoplos, John. 1989. "Game Theory without Partitions, and Applications to Speculation and Consensus." *Cowles Foundation Discussion Paper No. 914*.
- Gul, Faruk. 1998. "A Comment on Aumann's Bayesian View." *Econometrica* 66(4): 923–27.
- Jervis, Robert. 1976. *Perception and Misperception in International Relations*. Princeton: Princeton University Press.
- Jervis, Robert, Richard Ned Lebow, and Janice Gross Stein. 1985. *Psychology and Deterrence*. Baltimore: The John Hopkins University Press.
- Kim, Woosang, and Bruce Bueno de Mesquita. 1995. "How Perceptions Influence the Risk of War." *International Studies Quarterly* 39(1): 51–65.
- Laffont, Jean-Jacques, and David Martimort. 1997. "Collusion Under Asymmetric Information." *Econometrica* 65(4): 875–911.
- Leventoglu, Bahar, and Ahmer Tarar. 2006. "War and Incomplete Information." Typescript. Texas A&M University.
- Leventoglu, Bahar, and Branislav L. Slantchev. 2005. "The Armed Peace: A Punctuated Equilibrium Theory of War." Typescript. University of California–San Diego.
- Lewis, David K. 1969. *Convention: A Philosophical Study*. Cambridge: Harvard University Press.
- Milgrom, Paul. 1981. "An Axiomatic Characterization of Common Knowledge." *Econometrica* 49(1): 219–22.
- Milgrom, Paul, and Nancy Stokey. 1982. "Information, Trade and Common Knowledge." *Journal of Economic Theory* 26(1): 17–27.
- Morris, Stephen. 1991. "The Role of Beliefs in Economic Theory." *Ph.D. Dissertation*. Yale University.
- Morris, Stephen. 1994. "Trade with Heterogeneous Prior Beliefs and Asymmetric Information." *Econometrica* 62(6): 1327–47.
- Morrow, James D. 1985. "Capabilities, Uncertainty, and Resolve: A Limited Information Model of Crisis Bargaining." *American Journal of Political Science* 33(4): 941–72.
- Myerson, Roger B. 1979. "Incentive Compatibility and the Bargaining Problem." *Econometrica* 47(1): 61–73.
- Niou, Emerson M. S., Peter C. Ordeshook, and Gregory F. Rose. 1989. *The Balance of Power: Stability in International Systems*. New York: Cambridge University Press.
- Patty, John W., and Roberto A. Weber. 2006. "Agreeing to Fight: An Explanation of the Democratic Peace." *Politics Philosophy and Economics* 5(3): 227–47.
- Powell, Robert. 1999. *In the Shadow of Power: States and Strategies in International Politics*. Princeton: Princeton University Press.
- Powell, Robert. 2004. "Bargaining and Learning While Fighting." *American Journal of Political Science* 48(2): 344–61.
- Rubinstein, Ariel. 1998. *Modeling Bounded Rationality*. Cambridge: MIT University Press.
- Rubinstein, Ariel, and Asher Wolinsky. 1990. "On the Logic of 'Agreeing to Disagree' Type Results." *Journal of Economic Theory* 51(1): 1184–93.
- Schultz, Kenneth. 2001. *Democracy and Coercive Diplomacy*. New York: Cambridge University Press.
- Sebenius, James K., and John Geanakoplos. 1983. "Don't Bet on It: Contingent Agreements with Asymmetric Information." *Journal of the American Statistical Association* 78(382): 424–26.
- Shin, Hyun. 1989. "Non-Partitional Information on Dynamic State Spaces and the Possibility of Speculation." Typescript. University of Michigan.
- Slantchev, Branislav L. 2003. "The Principle of Convergence in Wartime Negotiation." *American Journal of Political Science* 97(4): 621–32.
- Smith, Alastair, and Allan C. Stam. 2004. "Bargaining and the Nature of War." *Journal of Conflict Resolution* 48(6): 783–813.
- Smith, Alastair, and Allan C. Stam. 2006. "Divergent Beliefs in 'Bargaining and the Nature of War'." *Journal of Conflict Resolution* 50(4): 614–18.
- Stein, Arthur A. 1990. *Why Nations Cooperate; Circumstances and Choice in International Relations*. Ithaca: Cornell University Press.
- Tirole, Jean. 1982. "On the Possibility of Speculation under Rational Expectations." *Econometrica* 50(5): 1163–82.
- Wagner, R. Harrison. 1994. "Peace, War and the Balance of Power." *American Political Science Review* 88(3): 593–607.
- Wagner, R. Harrison. 2000. "Bargaining and War." *American Journal of Political Science* 44(3): 469–84.
- Werner, Suzanne. 1998. "Negotiating the Terms of Settlement: War aims and Bargaining Leverage." *Journal of Conflict Resolution* 42(3): 321–43.
- Wittman, Donald. 1979. "How Wars End." *Journal of Conflict Resolution* 23(4): 743–63.
- Wittman, Donald. 2001. "War or Peace." Typescript. University of California–Santa Cruz.
- Wohlforth, William C. 1987. "The Perception of Power: Russia in the Pre-1914 Balance." *World Politics* 39(3): 353–81.