The Common Priors Assumption

A COMMENT ON "BARGAINING AND THE NATURE OF WAR"

MARK FEY
Department of Political Science
University of Rochester
KRISTOPHER W. RAMSAY
Department of Politics
Princeton University

In a recent article in this journal, Smith and Stam (2004) call into question the usefulness and applicability of what is know as the *common priors assumption* in the modeling of countries' strategic behavior in international relations. While the authors of this comment acknowledge that it is possible to incorporate noncommon priors in models of politics in a mathematically consistent fashion, they do not agree with the article's claims regarding the limitations of the common priors approach, which motivate Smith and Stam's rejection of it.

Keywords: *war; bargaining; common priors*

In a recent article in this journal, Smith and Stam (2004) call into question the usefulness and applicability of what is know as the *common priors assumption* (hereafter CPA) in the modeling of countries' strategic behavior in international relations. While we acknowledge that it is possible to incorporate noncommon priors in models of politics in a mathematically consistent fashion, we do not agree with the article's claims regarding the limitations of the common priors approach, which motivate the authors' rejection of it. Moreover, we believe that if the objective of modeling international interactions using game-theoretic tools is to identify causal mechanisms—in this case, linking the knowledge and/or information of players to causes of war—it is much more fruitful to remain within the common priors framework. That is, the assumption that players share a common prior at some initial stage is not simply a convention that we are taught in a game theory course but an

AUTHORS' NOTE: We thank Scott Ashworth, Adam Meirowitz, Curt Signorino, Branislav Slantchev, Alastair Smith, and Allan Stam for comments and advice. Kris Ramsay gratefully acknowledges financial support from NSF grant SES-0413381.

JOURNAL OF CONFLICT RESOLUTION, Vol. 50 No. 4, August 2006 607-613
DOI: 10.1177/0022002706289182

epistemologically desirable modeling choice if our objective is explanation. In particular, for the study of information and/or knowledge and war, the CPA requires the theorist to explain both the divergence and subsequent convergence of beliefs about war in a logically consistent fashion. Absent the discipline imposed by the common priors assumption, we are left with explanations of war that, at their core, claim nations fight because they have heterogeneous priors. This is the informational equivalent to a superficial rational choice argument that nations fight wars because they like to.

While admitting their approach is a "radical" departure from the standard approach to information, Smith and Stam (2004, 784) claim that it "is a significant technical innovation and a more accurate account of the empirical world." The authors claim that the assumption of common priors implies that if two countries observe the same information, they must have the same beliefs. Since this does not happen in reality, the argument goes, we should reject the common priors assumption and instead view uncommon priors as modeling differing worldviews, through which the same information is interpreted differently.[1] However, we contend that the common priors framework is, in fact, flexible enough to account for the variety of situations that they argue require its abandonment. Our justification for this claim has two parts. First, we argue that Smith and Stam's "innovation" of modeling countries' differing worldviews (in which beliefs fail to converge in the presence of the same information) can be accommodated in the standard common priors framework. Second, we show that, contrary to the claim of Smith and Stam, it is possible within the common priors framework for new public information to lead to "learning" in which beliefs diverge, rather than converge.

THE SCOPE OF THE COMMON PRIORS ASSUMPTION

So what exactly is the common priors assumption? When we assume common priors, we are assuming that the beliefs of rational players in a game are generated by updating, via Bayes's rule, a single prior over a commonly known set of states (Lipman 2003). Clearly, the CPA places limits on what beliefs are possible among rational actors. The question remains, however, are these limits as severe as is implied by Smith and Stam (2004)? That is, can two rational actors observe the same event and reach different conclusions about their likely success in war? By our reading, there are really two questions here. First, does the CPA permit two rational countries to know that the same event has occurred and hold different beliefs about some second correlated event? Second, does the CPA allow two countries to learn the "same thing" is true and hold different beliefs about the decision-relevant event? In this section, we show via several examples that the answer to both questions is yes and conclude that, contrary to the claims of Smith and Stam, it is not necessary to reject the CPA to model these occurrences.

1. As the authors put it, "The heterogeneous-beliefs approach that we employ instead focuses on underlying differences in how actors think the world works" (Smith and Stam 2004, 784).

TECHNICAL PRELIMINARIES

In this note, we use a method of modeling information, sometimes called a knowledge model, that is quite general.[2] A knowledge model consists of three elements: a *state space*, denoted Ω, with generic element ω, that defines the possible realizations of decision-relevant variables for actors; a probability measure on that space, denoted $\pi(\omega)$, which is the common prior at the focus of our discussion; and a collection of information partitions $\mathbf{P}_i$, one for each of n players. For each $\omega \in \Omega$, $\mathbf{P}_i(\omega)$ is interpreted as the collection of states that individual i thinks are possible when the true state is ω. These partitions represent *all* relevant information for the actors, such as their educational background, any inputs they may receive from advisers, life experience, and so on. Our justification for assuming that the information of players can be represented by a partition is that it can be shown that a partition is the only information structure consistent with Bayesian rationality.[3]

PRIORS, POSTERIORS, AND WHAT IS KNOWN

Consider the example described by Smith and Stam (2004) of two countries (A and B) observing the effects of a new weapon in some third party's war. To incorporate this example into our knowledge model, define the following two events. Let Ω be given, and let G be the event (defined as a subset of states) that country A will win a war between A and B. Next, consider the third-party war. Since evidence from this war implies something about the outcome of a conflict between A and B, let the event E be correlated with G and define it as the set of states consistent with viewing the outcome of the third-party war. We may also suppose that $G \cap E \neq \emptyset$ and $E \not\subset G$, so that E does not imply that one of our two states will win for sure. That is, E does not resolve all uncertainty about the true state belonging to or being outside of G. To describe the information of countries, we begin by assuming for simplicity that the common prior is uniform. That is, $\pi(\omega) = 1/\kappa$, where k is $|\Omega|$.[4] Also, suppose that since each observing country has a different military history, different strategic planners, and different advisers, each country may assess the implications of the new weapon differently. As such, we can represent the differences in the countries' views by information partitions, one $\mathbf{P}_i$ for each. Since the countries' knowledge of the circumstances relevant for a war with the other differ, these partitions will differ. We can call each country's partition a worldview or theory. Like Smith and Stam (2004), without loss of generality, we can assume that these partitions are common knowledge. That is, each side knows how various aspects of the state of the world will influence the other's beliefs about the true state of the world, and each knows the other knows, and so forth.

2. For a more detailed discussion of models of knowledge applied to international relations, see Fey and Ramsay (2005).

3. For an accessible proof of this result, we refer the reader to Rubinstein (1998).

4. The results do not depend on this assumption, but it removes unnecessary clutter from the exposition.

Finally, suppose that E is known to each country at the true state of the world. Note, while both countries know E, it is not common knowledge that they know it. E is not common knowledge because A and B view the war from afar, and no single country can be *sure* that the other knows the war's details.

Given the difference in the two countries' private information, we may wonder, does the CPA require that knowing how the new weapons influenced the third-party war imply that the countries agree on the probability that A will win a bilateral war? Even though these countries are rational, the answer is no. Consider the following example.

Suppose that there are six possible states of the world, $\Omega = \{\omega_1, \omega_2, \omega_3, \omega_4, \omega_5, \omega_6\}$, that $G = \{\omega_1, \omega_2\}$, and $E = \{\omega_1, \omega_2, \omega_5, \omega_6\}$. Recall that, due to various histories, advisers, and theories of war, the "worldviews" of countries A and B differ, in a way that is determined by their information partitions. Let these two partitions be

$$\mathbf{P}_A = \{\{\omega_1\}, \{\omega_3, \omega_4, \omega_5\}, \{\omega_2, \omega_6\}\}, \tag{1}$$

$$\mathbf{P}_B = \{\{\omega_1, \omega_3\}, \{\omega_2, \omega_5, \omega_6\}, \{\omega_4\}\}. \tag{2}$$

Given these partitions, we can then also define a knowledge operator such that country i knows an event F when $\omega \in K_i(F)$, where

$$K_i(F) = \{\omega : P_i(\omega) \in F\}, \tag{3}$$

and as knowledge can be contingent on the state of the world, suppose that $\omega = \omega_2$.

Since $K_A(E) = \{\omega_1, \omega_2, \omega_6\}$ and $K_B = (E) = \{\omega_2, \omega_5, \omega_6\}$, both A and B know E at ω_2. Yet even though both A and B know E, at ω_2, the probability that A assigns to the true state being at a state in G is 1/2, and the probability that B assigns to being in G is 1/3. So, even though both countries know E and are rational, they need not hold the same assessment of A's likelihood of success in war. Why is this the case? For the exact reason Smith and Stam (2004) argue, the two countries have private information, or theories, that also influences their beliefs. As such, the simple fact that both countries *know* some event has occurred does not change the fact that their beliefs about A's success in war are different.

So, if mere knowledge of an event, correlated with the probability of A's success in war, does not necessarily lead countries to have the same assessment of the likelihood that A will win a war, does "learning" about the state of the world necessarily imply equal assessments? To assess this question, we must consider a dynamic model of knowledge.

PRIORS, POSTERIORS, AND WHAT IS LEARNED

We begin by noting that, as Geanakoplos (1992) points out, we can model limited knowledge (or private information) by analogy to a far-off observer who cannot quite distinguish some objects from another. That is, differences in information are

captured by the idea that different actors can tell different states of the world from one another. Therefore, if a player is "learning," she will have different partitions at different points in the learning process. We can also say that one actor has "worse" information than another (or herself at a different time) if her partition consists of a *coarsening*[5] of the other's. So, a player is learning if she can better distinguish between various possible states. In particular, learning is often about being able to eliminate some state of the world from being the true state. Such a logic underlies any signaling or screening model, where only certain types of players will take certain actions or where a random variable can only take on certain values under known conditions.

So let us now consider the effects of refining the information partitions of players. It is easy to show that there exist refinement processes that do not lead to immediate convergence. That is, if we consider the situation where two countries are learning about A's probability of winning a war with B, we need not conclude that the countries share the same posterior. Here, again, we proceed with an example.

Suppose that after observing the use of the new weapon described above, the players "learn" how to distinguish ω_6 from all other states. To show how this learning process can easily be translated into the setting found in Smith and Stam (2004), consider the following reformulation of their model. Let the state space be all the possible realizations of the sequence of battles during a war. Then suppose that in the true (realized) state, the war proceeds deterministically, with state A winning the first battle. Obviously, both countries know that if A wins the first battle, then any state of the world in which B wins the first battle could not be the realized state. Therefore, both countries have learned how to distinguish the states where B wins the first battle. In particular, they can eliminate those states from being possible and adjust their beliefs accordingly.

Clearly, this new information changes both countries' posterior probabilities that A will win a war between A and B. In fact, if we refine $\mathbf{P}_A$ and $\mathbf{P}_B$ as given above by allowing them to distinguish ω_6 from all other states, A then believes that it will win a war with probability 1, while B believes the probability of success for A is only 1/2. Note that the countries' posteriors are still not the same; that is, learning the "same thing," where the same thing is that the realized state is not ω_6, need not lead to common posterior beliefs.

Moreover, it is not only the case that learning may not lead to convergence, but it might be that learning the same "fact" induces countries' beliefs to *diverge*. Consider the following two information partitions for A and B:

$$\mathbf{P}'_A = \{\{\omega_1\}, \{\omega_3, \omega_4,\}, \{\omega_2, \omega_5, \omega_6\}\}, \tag{4}$$

$$\mathbf{P}'_B = \{\{\omega_1\}, \{\omega_2, \omega_3, \omega_4\}, \{\omega_5, \omega_6,\}\}. \tag{5}$$

5. One partition coarsens another if each coarser (less informative) partition has atoms that consist of unions of the atoms of the finer (more informative) one (Aumann 1998).

At $\omega = \omega_2$, Bayes's rule implies that both countries believe the probability that *A* will win a war between *A* and *B* is 1/3, albeit for different reasons. Next, refine both countries' partitions such that they have both learned how to differentiate ω_3 from all other states. Now, the countries' beliefs diverge; *A* still holds the belief that it will win a war with probability 1/3, but *B* now believes *A*'s probability of success is 1/2.[6] Therefore, we can conclude that assuming common priors does not prevent divergence in countries' beliefs, even if both countries learn the "same thing."

CONCLUSION

In sum, we have demonstrated that the common priors framework is consistent with the following points: (1) mutual knowledge of an event that is correlated with the likelihood of success in war is not sufficient to make countries hold the same beliefs; (2) it is not necessarily the case that when two countries "learn the same thing," by observing some state-dependent event, that they must hold the same posterior beliefs; and (3) with rational agents, it is possible for learning to produce divergence in posteriors where there was once convergence. These points highlight the fact that, in "Bargaining and the Nature of War," the common priors assumption is not the problem. That is, the common priors framework is flexible enough to allow for the various factors that the authors wish to incorporate in their model. However, this flexibility does not mean that the "agreeing to disagree" result would go away in a model of war. The real problem is that the common priors framework rules out the possibility of rational actors taking up "irrational" bets, like inefficient wars, when it is common knowledge that someone is going to lose and that there exists some Pareto-optimal alternative to fighting. While a full formal exposition of this fact is beyond the scope of this note, this result is quite general.[7]

Finally, we arrive at the substantive problem of interpreting the noncommon priors assumption. As Aumann (1998) has demonstrated, if we interpret differing worldviews or theories as the product of life experience, and we wish to maintain the claim that the actors in our model are rational (i.e., apply the rules of conditional probability), it can only be the case that the common priors assumption is appropriate. To make the point another way, suppose we were to abandon the common priors assumption. The question then becomes the following: what are heterogeneous priors if they are not differences in information? If they are not the product of education, culture, or life experience, what do they represent? As even Morris (1995) concedes, the common priors assumption enables us to focus on pure informational dynamics, but he then argues that not all modeling issues in economics are informational. While this may be fine for Morris, the problem is that the argument in "Bargaining and the Nature of War" is about information. It is about what countries

6. As may be obvious, if the two countries "learn enough," their beliefs about *A*'s probability of success will converge. For finite state spaces, this convergence will happen in a finite number of steps. See Sebenius and Geanakoplos (1983) and Geanakoplos and Polemarchakis (1982).

7. See Fey and Ramsay (2005) for a formal treatment of this claim.

know before they fight and why they believe what they do. By Aumann's result, one could even argue that Smith and Stam's (2004) argument *requires* the use of the CPA. But if they were to assume common priors, the question remains as to whether any fighting would occur in equilibrium.

So, in the end, we have three points to make. First, while we agree that there is nothing inherently wrong with assuming noncommon priors, the limitations of the common priors framework are not as severe as implied by Smith and Stam (2004). In particular, if we conceptualize players as having different theories as to how the world works resulting from private information (i.e., the way they were taught in graduate school, their cultural and social background, and their "life experiences"), it is relatively easy to show that even with common priors, the same event can lead rational agents to draw different conclusions. In fact, the only time when players' beliefs must agree is when their beliefs are about a common knowledge event, when actions are common knowledge, or when the players' posteriors are common knowledge (Geanakoplos 1992). That is, it is *not* true that different *worldviews* or *theories* require a departure from the CPA. Second, given that we often believe that differences in assessments of the probability of success in war result from differences in education, culture, and information and, as Aumman (1998) shows, internally consistent information arguments require the common priors assumption, we are left wondering how to interpret the noncommon priors assumption. In many ways, the argument that states fight wars because they have noncommon priors begs more questions than it answers. Finally, the CPA is not merely a convention we accept because we may be taught to invoke it in our analysis. Rather, it is an epistemologically desirable assumption that forces the theorist to provide a complete explanation of the way uncertainty may affect players' decision making in a strategic setting. As such, the abandonment of the CPA is something that should not be taken lightly.

REFERENCES

Aumann, Robert J. 1998. Common priors: A reply to Gul. *Econometrica* 66 (4): 929-38.

Fey, Mark, and Kristopher W. Ramsay. 2005. Mutual optimism and war. Mimeo, University of Rochester.

Geanakoplos, John. 1992. Common knowledge. *Journal of Economic Perspectives* 6 (4): 53-82.

Geanakoplos, John, and Heraklis M. Polemarchakis. 1982. We can't disagree forever. *Journal of Economic Theory* 28 (2): 192-200.

Lipman, Barton L. 2003. Finite order implications of common priors. Mimeo, Boston University.

Morris, Stephen. 1995. The common prior assumption in economic theory. *Economics and Philosophy* 11 (2): 227-53.

Rubinstein, Ariel. 1998. *Modeling bounded rationality*. Cambridge, MA: MIT 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): 424-6.

Smith, Alastair, and Allan C. Stam. 2004. Bargaining and the nature of war. *Journal of Conflict Resolution* 48 (6): 783-813.