Frank C. Zagare*

Two Against One: Deterrence in the Triad

https://doi.org/10.1515/peps-2023-0050
Received September 12, 2023; accepted November 30, 2023; published online January 2, 2024

Abstract: This essay explores a deterrence relationship between one defender and two challengers that are interconnected game-theoretically but are otherwise acting independently. In the double deterrence game model, the primary challenger makes the first move, defender, the second, and the secondary challenger makes its move only after the defender and the primary challenger have made their choices. In other words, the threat of the secondary challenger is both latent and contingent. There are four types of defenders, and three types of the two challengers. These preference assumptions define 36 distinct strategic environments, or games. These games are analyzed under both complete and incomplete information. The results are highly sensitive to the specific combination of player types. Deterrence is most likely to succeed when a defender's threat is highly credible, and to fail when it is low. The secondary challenger is most likely to have an impact on play only at intermediate levels of the defender's *and* the primary challenger's credibility. Assuming a rising challenger, or a declining defender, or both, a danger zone will exist just prior to, or immediately after, the primary challenger has achieved parity with the defender, that is, when a balance of capabilities exists. The principal aim of a defender of the status quo facing two challengers acting independently should be on preventing a challenge by its primary opponent. This conclusion is robust no matter what configuration of player types one assumes.

Keywords: deterrence; perfect deterrence theory; game theory; power transition theory; three body problem

JEL Classification: C70; F51; F52

During the cold war conventional wisdom held that an all-out war between the United States and the Soviet Union was all-but precluded. The key to this strategic nirvana was a carefully calibrated *balance* of strategic weapons and the high costs associated with nuclear conflict (Intriligator and Brito 1984). The policy that was credited with bringing this state of affairs about was labeled Mutual Assured Destruction, or MAD. Each side threatened to obliterate the other if it were attacked.

***Corresponding author: Frank C. Zagare**, Political Science, University of Buffalo, 504 Park Hall, Buffalo, 14260, NY, USA, E-mail: fczagare@buffalo.edu

Based on this logic some strategic thinkers argued for the selective proliferation of nuclear weapons, to Iran for example, in order to stabilize the relationship between two otherwise hostile states (Waltz 2012). Others argued that Ukraine would be misguided if it surrendered its nuclear arsenal, as it eventually did in 1994 (Mearsheimer 1983).

However, the theory underlying this policy, sometimes called Classical (or misappropriately, rational) Deterrence Theory, does not pass the test of strict logic (Zagare 1996). It assumes, simultaneously, that the players are both rational and irrational: they are rational when they are being deterred and are irrational when they are deterring. Bernard Brodie (1959: 293), considered by many to be the seminal deterrence theorist, put it this way: "For the sake of deterrence before hostilities, the enemy must expect us to be vindictive and irrational if he attacks us." The noted game theorist, Thomas Schelling, who was the recipient of the 2005 Nobel Prize in economics, also argued that nuclear deterrence only worked if an aggressor was convinced that its opponent would retaliate – irrationally. As he wrote so succinctly: "another paradox of deterrence is that it does not always help to be, or to be believed to be, fully rational, cool headed, and in control of one's country." In other words, it was rational to be irrational (Schelling 1966: 37).

Logically inconsistent theories are *prima facie* seductive, yet they are also fatally flawed. They invite theorists with a point of view to draw almost any conclusion, including its exact opposite, depending on the analyst's policy preferences. So, it is unsurprising that other classical deterrence theorists, working from the same set of assumptions, oppose disseminating nuclear weapons to Iran or to any other state actor including Ukraine.

To overcome the logical inconsistency of Classical Deterrence Theory, an alternative theory called Perfect Deterrence Theory was constructed. Perfect Deterrence Theory assumes that the players are rational (or purposeful) at all times (Zagare and Kilgour 2000).[1] In this alternative specification there are certain conditions under which wars cannot be avoided. For example, according to this theory it is *possible* that Russia may still have invaded Ukraine even if the Ukrainians had not given up their nuclear capability once the Soviet Union broke up. Low level conflicts are very difficult to deter, as are situations where one state seeks to deter an attack on an ally. Worse still is the propensity of these conflicts to escalate. Sunk costs may play a role here. But so may uncertainty about the extent of resistance, if any. Risk taking leaders are the most dangerous.

1 The theory's name reflects its reliance on Selten's (1975) definition of rational strategic behavior. Strict adherence to his *perfectness* criterion assures logical consistency and sets it apart from Classical Deterrence Theory which presumes the possibility of both rational and conditionally irrational decisions by the same player in the same game. No claim is made that the theory itself is perfect.

For obvious reasons, both Classical Deterrence Theory and Perfect Deterrence Theory have focused on dyadic relationships. It is clear, however, that that focus now needs adjustment. In the current environment there are currently two dissatisfied major nuclear powers that would prefer, *ceteris paribus*, an adjustment of the rules that support the international political and economic system as it operates today. At the global level, then, a key question is how a defender of the status quo might deal with two potential challengers.

The answer to this question depends, in part, on the relationship of the two challengers. There are three possibilities:

1. Both dissatisfied actors act independently in completely separate disputes. If this is the case, there is no need for a theory shift (Lakatos 1970). Current theory still applies. Assuming that the conditions needed to deter the most problematic case are met, these same requirements would suffice to deter the less problematic case as well. The assumption of complete independence implies that it is highly unlikely that the two disputes would break out simultaneously.
2. Both actors operate as one in a tacit alliance. Again, current theory can handle this case, as it did during the Cold War.
3. The secondary opponent becomes a player *if and only if* the primary opponent upsets the status quo. It is clear that it is this special situation that requires further scrutiny.

The purpose of this essay is to extend the logic of Perfect Deterrence Theory to the third case which, for expository purposes, will be referred to as the three-body problem. To my knowledge, this case has not, as yet, been analyzed game-theoretically.[2] Nonetheless, it is not a novel problem as discussed below.

It is important to note that, unlike Classical Deterrence Theory, Perfect Deterrence Theory is not simply a divergent theory of nuclear war avoidance. Rather, it is a universal theory of conflict initiation, escalation, and resolution, applicable to both nuclear and nonnuclear interactions. In other words, it applies to both conventional challenges to the status quo that have the potential to escalate, and to direct deterrence confrontations between and among nuclear powers.

The collection of game-theoretic models that constitute the theory assumes that the players in the game prefer to win or, if necessary, to lose at the lowest level of conflict. Most other preference relationships are taken as strategic variables. For example, some players might prefer conflict to losing. Players with such a preference

2 The literature on multi-party conflicts is large. For a representative sample see Altfeld and Bueno de Mesquita (1979), Schweller (1998), Krainin and Wiseman (2016), Powell (2017), and Iwanami (2023). For three-person game models that were developed within the umbrella of Perfect Deterrence Theory, see Quackenbush (2006) and Zagare and Kilgour (2003).

are assumed to have *credible* threats, that is, threats that are rational to execute. Other players, with the opposite preference, have threats that lack credibility. The players also may or may not prefer the status quo to conflict. A player whose opponent prefers the status quo to conflict is said to have a capable threat, that is, a threat that hurts (Schelling 1966: 7). Threats that do not hurt are considered incapable. Thus, there may be four types of threats: threats that are both credible and capable, threats that are neither, threats that are credible but not capable, and threats that are capable but lack credibility.[3]

In the analysis that follows, a simple three-person game model will be described and analyzed under a variety of (plausible) assumptions about the credibility of the players' threats. These assumptions will be made explicit and justified. As a first cut the game will be explored under the condition of complete information. Information is complete when the players are fully informed about each other's preferences (or goals). After that, the assumption is relaxed. A focused version of the game will be analyzed under incomplete information, that is, when the players are unsure about whether the other players prefer to execute their threats when challenged.

1 Game Form

Three-body deterrence games may arise in a variety of empirical circumstances. For example, at various times, in the 19th century, Germany faced potential threats from both Austria and Russia, not to mention France. Not surprisingly, an important goal of Bismarck's diplomacy was to neutralize Austria in order to minimize the probability of a conflict with Russia. He was successful. In 1879 he negotiated an alliance with Austria that ended only when Germany was defeated after the First World War. Similarly, in 1939 after the invasion of Poland, the United States feared that it might be drawn into a war with both Germany and Japan. Ultimately, deterrence failed twice, first in the Pacific and, shortly thereafter, in Europe, as it did in 1912 when the Balkan League took advantage of Turkey's ongoing war with Italy (Overy 2021: 10).

Regardless of the conditions that give rise to them, three-body deterrence situations, by definition, share a number of common characteristics. The Double Deterrence Game depicted in Figure 1 is a simple, yet plausible, model that reflects the most salient of these commonalities. For example, there is one player, called Defender, and two potential opponents, a primary opponent, here called Challenger

3 In Perfect Deterrence Theory capable threats constitute a necessary condition for deterrence success. In other words, deterrence of a dissatisfied actor is not possible when a defender of the status quo lacks a threat that is capable. Thus, for the purposes of this essay, all deterrent threats are assumed to be capable. By contrast, threat credibility will be treated as a variable.

1, and a secondary opponent, here called Challenger 2.[4] The choices of the two challengers are interconnected game-theoretically, but they are assumed to make autonomous strategic choices. In other words, the choice of the second challenger will depend not only on the Defender's choice, but on the choice of Challenger 1. To put this in a slightly different way, the assumption will be that Challenger 1 presents the immediate threat to Defender, while Challenger 2's threat is both contingent and latent.

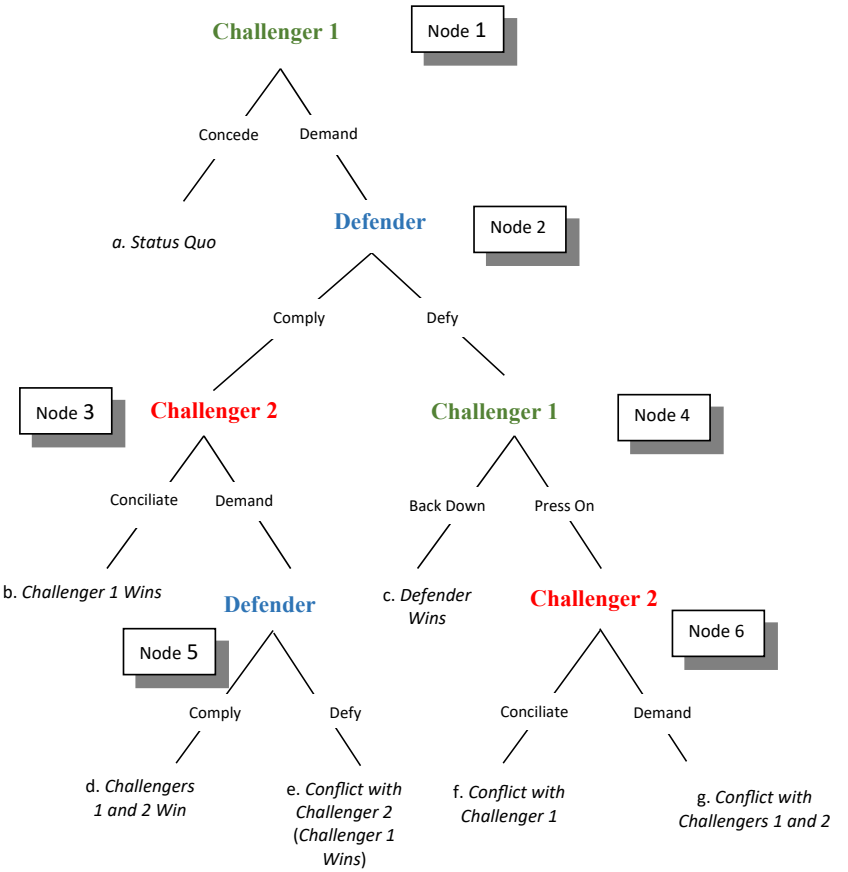


Figure 1: The Double Deterrence Game.

4 There is nothing in the logic of the Double Deterrence Game that requires that the primary challenger to be stronger than the secondary challenger. Challenger 1 should be thought of as the immediate challenger and Challenger 2 as the *possibly* opportunistic challenger.

The Double Deterrence Game is a three-person noncooperative game model that fully captures what Snyder and Diesing (1977) refer to as the "precipitant-challenge-confrontation scheme." As Figure 1 shows, in this game Challenger 1 begins play at Node 1 by deciding whether to *Demand* a change in the *Status Quo*, or to *Concede* the issue by doing nothing. Of course, if challenger does nothing, the game ends before it begins, and the *Status Quo* holds. But if Challenger 1 makes a demand, Defender (at Node 2) is faced with a difficult decision: whether to *Comply* with the demand or *Defy* Challenger 1. What happens after Defender's choice will depend on the choice of one of the two challengers.

Consider first the implications of Defender's choice to resist Challenger 1's demand at Node 2. In this case, Challenger 1 could *Back Down* at Node 4. If and when this occurs, the game ends and the outcome is *Defender Wins*. But if Challenger 1 decides to *Press On*, a conflict at some unspecified level occurs. The type of conflict that breaks out, however, depends on Challenger 2's choice. If (at Node 6) Challenger 2 decides to remain neutral and *Conciliate* Defender, Defender's conflict is solely with Challenger 1. On the other hand, if Challenger 2 makes its own demand, Defender's conflict is with both challengers, its nightmare scenario.

Consider now the implications of Defender's choice to *Comply* with Challenger 1's demand at Node 2. Once Defender capitulates, *Challenger 1 Wins* regardless of Challenger 2's choice at Node 3. But Challenger 2's choice has implications for Defender. If at Node 3 Challenger 2 issues its own demand, Defender has another difficult choice to make. If at Node 5 Defender chooses to *Comply*, *Challengers 1 and 2 Win*. But if Defender chooses to *Defy* the demand, the outcome is a bilateral *Conflict with Challenger 2* (but a win for Challenger 1 nonetheless).

To be sure, the Double Deterrence Game of Figure 1 is extremely austere. Specifically, the consequences of Defender's choice at Node 5, and of Challenger 2's choice at Node 6, have been truncated. These decision nodes could easily be extended to include a subsequent choice by Challenger 2 in the first instance, or a subsequent choice by Defender in the second, or both. Nonetheless the presumption is that these simplifications are useful. A "no-fat" modeling approach (Rasmusen 1989) allows a firm focus on the critical role played by deterrent threats in a game in which there are two challengers. In a future study, if warranted, these simplifications could be modified or relaxed.

2 Outcomes and Preferences

There are seven outcomes in the Double Deterrence Game. Of the seven, two involve a win for one or both Challengers (and a loss for Defender). Another is associated

with a win for Defender. There are also three distinct types of conflicts. Of course, when deterrence succeeds, the *Status Quo* will hold.

A game is defined by both the rules that govern play, that is, by the number of players, their choices, the set of possible outcomes, and the preferences of the players over the set of possible outcomes. With three players, each with at two moves, and seven outcomes, the Double Deterrence Game can have many variants. To gain tractability, not to mention theoretical relevance, some additional simplifying assumptions are needed. We begin with Defender's preferences.[5]

2.1 Defender's Preferences and Types

A number of preference relationships are manifest. For the purpose of this study, the assumption will be that Defender most prefers the *Status Quo*, and next most prefers the outcome *Defender Wins*. Of the two outcomes associated with a loss for Defender, the assumption is that Defender prefers the outcome *Challenger 1 Wins* to the outcome *Challengers 1 and 2 Win*. Finally, it is assumed that Defender prefers a bilateral conflict with Challenger 1 to a concurrent conflict with both Challengers 1 and 2. Defender's preference among the remaining outcomes determines its type. There are four plausible types of Defenders: *Staunch*, *Semi-Staunch*, *Submissive*, and *Semi-Submissive*.

Staunch and *Semi-Staunch* Defenders strictly prefer a *Conflict with Challenger 1* and a *Conflict with Challengers 1 and 2* to the outcome *Challenger 1 Wins*. *Staunch* Defenders also prefer a unilateral *Conflict with Challenger 2* to *Challengers 1 and 2 Win*. By contrast, *Semi-Staunch* Defenders prefer to conciliate both challengers rather than resist Challenger 2's demands. The columns of Table 1a list the postulated

Table 1a: Preferences of Staunch and Semi-Staunch Defenders.

Rank	Staunch	Semi-Staunch
1	a. *Status Quo*	a. *Status Quo*
2	c. *Defender Wins*	c. *Defender Wins*
3	f. *Conflict with Challenger 1*	f. *Conflict with Challenger 1*
4	g. *Conflict with Challengers 1 and 2*	g. *Conflict with Challengers 1 and 2*
5	b. *Challenger 1 Wins*	b. *Challenger 1 Wins*
6	e. *Conflict with Challenger 2*	d. *Challengers 1 and 2 Win*
7	d. *Challengers 1 and 2 Win*	e. *Conflict with Challenger 2*

5 For a formal summary of the restrictions placed on the players' preferences, see the Appendix.

Table 1b: Preferences of Submissive and Semi-Submissive Defenders.

Rank	Submissive	Semi-Submissive
1	a. *Status Quo*	a. *Status Quo*
2	c. *Defender Wins*	c. *Defender Wins*
3	b. *Challenger 1 Wins*	b. *Challenger 1 Wins*
4	d. *Challengers 1 and 2 Win*	e. *Conflict with Challenger 2*
5	e. *Conflict with Challenger 2*	d. *Challengers 1 and 2 Win*
6	f. *Conflict with Challenger 1*	f. *Conflict with Challenger 1*
7	g. *Conflict with Challengers 1 and 2*	g. *Conflict with Challengers 1 and 2*

preferences of a *Staunch* and *Semi-Staunch* Defender, respectively, from best (i.e. "1") to worst (i.e. "7"). For example, both types of Defenders prefer the *Status Quo* to *Defender Wins*, and *Defender Wins* to *Conflict with Challenger 1*, and so on. Similarly, the columns of Table 1b list the postulated preferences of a *Submissive* and *Semi-Submissive* Defender, respectively, again from best to worst.

As Table 1b suggests, both types of *Submissive* Defenders strictly prefer the outcome *Challenger 1 Wins* to all three conflict outcomes and the outcome *Challengers 1 and 2 Win*, least prefers a *Conflict with Challengers 1 and 2*, and next-least prefers a *Conflict with Challenger 1*. Defender's preference between *Challengers 1 and 2 Win* and *Conflict with Challenger 2* determines it sub-type. *Submissive* Defenders prefer *Challengers 1 and 2 Win*. *Semi-Submissive* Challengers prefer not to accede to Challenger 2's demands at Node 5, that is, they prefer a *Conflict with Challenger 2*.

2.2 Challenger 1's Preferences and Types

Consider now Challenger 1's preferences. Any challenger, by definition, prefers to upset the *Status Quo*. In Perfect Deterrence Theory the assumption is that the players prefer to win, or if it comes to it, lose, at the lowest conflict level. It follows therefore that Challenger 1 prefers the outcome *Challenger 1 Wins* to the *Status Quo*. The further assumption will be that if and when Defender has fully complied with Challenger 1's demand at Node 2, Challenger 1's involvement in the game will be terminated. In other words, Challenger 1 is assumed to be indifferent with respect to what Challenger 2 does at Node 3 and what, if anything, Defender does at Node 5.[6] To put this in a slightly different way, the assumption is that Challenger 1 prefers *each*

6 The assumption is strategically unimportant. It is made simply for convenience.

of the outcomes that result from Defenders choice to *Comply* at Node 2 to *any* of the three outcomes that are possible when Defender is defiant at Node 2.

Challenger 1, however, may or may not be indifferent with respect to what Challenger 2 does at Node 6. If in fact Challenger 1 is indifferent, there are only two possibilities. Challenger 1 might prefer to *Back Down*, in which case the outcome is *Defender Wins*, or to *Press On* and induce a conflict that may or may not involve Challenger 2. An indifferent Challenger 1 that prefers to *Back Down* is called *Hesitant*. *Determined* Challengers, by contrast, prefer to *Press On*. On the other hand, Challenger 1 may be invested in what Challenger 2 does at Node 6. A *Cautious* Challenger 1 will *Press On* if and only if it anticipates that Challenger 2 will also confront Defender.

To summarize briefly, Challenger 1 may be one of three types. A *Determined* Challenger 1 prefers to *Press On* at Node 2, a *Hesitant* Challenger 1 prefers to *Back Down* at Node 2. And a *Cautious* Challenger 1 will *Press On* if and only if it expects Challenger 2 to issue its own demand at Node 6. Tables 2a and 2b contains the particulars.

Table 2a: Preferences of a Hesitant and Determined Challenger 1.

Rank	Hesitant	Rank	Determined
1	b. *Challenger 1 Wins*	1	b. *Challenger 1 Wins*
1	d. *Challengers 1 and 2 Win*	1	d. *Challengers 1 and 2 Win*
1	e. *Conflict with Challenger 2*	1	e. *Conflict with Challenger 2*
2	a. *Status Quo*	2	a. *Status Quo*
3	c. *Defender Wins*	3	g. *Conflict with Challengers 1 and 2*
4	g. *Conflict with Challengers 1 and 2*	3	f. *Conflict with Challenger 1*
4	f. *Conflict with Challenger 1*	4	c. *Defender Wins*

Table 2b: Preferences of a Cautious Challenger 1.

Rank	Cautious
1	b. *Challenger 1 Wins*
1	d. *Challengers 1 and 2 Win*
1	e. *Conflict with Challenger 2*
2	a. *Status Quo*
3	g. *Conflict with Challengers 1 and 2*
4	c. *Defender Wins*
5	f. *Conflict with Challenger 1*

2.3 Challenger 2's Preferences and Types

Unlike Challenger 1, Challenger 2's choices are strictly reactive. At Node 3, it must decide to initiate a conflict with the Defender by itself. By contrast, at Node 6 it must decide whether or not to join an existing conflict. In the latter case, its choice is strictly determined by its type. *Opportunistic* Challenger 2s will issue its own demand bringing about a *Conflict with Challengers 1 and 2*, By contrast, A *Reluctant* Challenger 2 will stand aside. In this case, the outcome will be a *Conflict with Challenger 1*.

As will be seen, Challenger 2's choice at Node 3 depends on both its and Defender's type. Clearly, of the three outcomes in the subgame that begins at Node 3, Challenger 2 most prefers that Defender accede to whatever is demanded, that is, the outcome *Challengers 1 and 2 Win*. But to achieve this outcome, Challenger 2 must risk a conflict. A Challenger 2 that prefers to avoid the risk is called *Restrained*. A *Restrained* Challenger 2 strictly prefers the outcome *Challenger 1 Wins* to *Conflict with Challenger 2*. Risk-taking Challenger 2s, with the opposite preference, are called *Persistent*.

All things being equal, a *Persistent* Challenger 2 is also likely to be *Opportunistic*. But a *Restrained* Challenger 2 may or may not be *Reluctant*. Thus, in this analysis we consider three types of Challenger 2s: *Persistent* and *Opportunistic*, *Restrained and Opportunistic*, and *Restrained* and *Reluctant*. The preferences that define these sub-types are summarized in Table 3.[7]

Table 3: Sub-types of Challenger 2.

Rank	Persistent	Restrained	Reluctant	Opportunistic
1	d. *Challengers 1 and 2 Win*	d. *Challengers 1 and 2 Win*	d. *Challengers 1 and 2 Win*	d. *Challengers 1 and 2 Win*
2	e. *Conflict with Challenger 2*	b. *Challenger 1 Wins*	f. *Conflict with Challenger 1*	g. *Conflict with Challengers 1 and 2*
3	b. *Challenger 1 Wins*	e. *Conflict with Challenger 2*	g. *Conflict with Challengers 1 and 2*	f. *Conflict with Challenger 1*

7 Challenger 2 never chooses between outcomes b, d and e and outcomes f and g. Therefore, no assumption is made about the relative rankings of the outcomes in these two distinct sets of outcomes. The ordinal utilities given here, then, should be understood to reflect Challenger 2's preference solely within each set and not across the two sets. Challenger 2 also never chooses between outcomes a and c and any other outcome. For convenience, the assumption will be that Challenger 2 is indifferent between both of these outcomes and the least preferred alternative in each of the two sets of outcomes in Table 3. This assumption is of no strategic consequence.

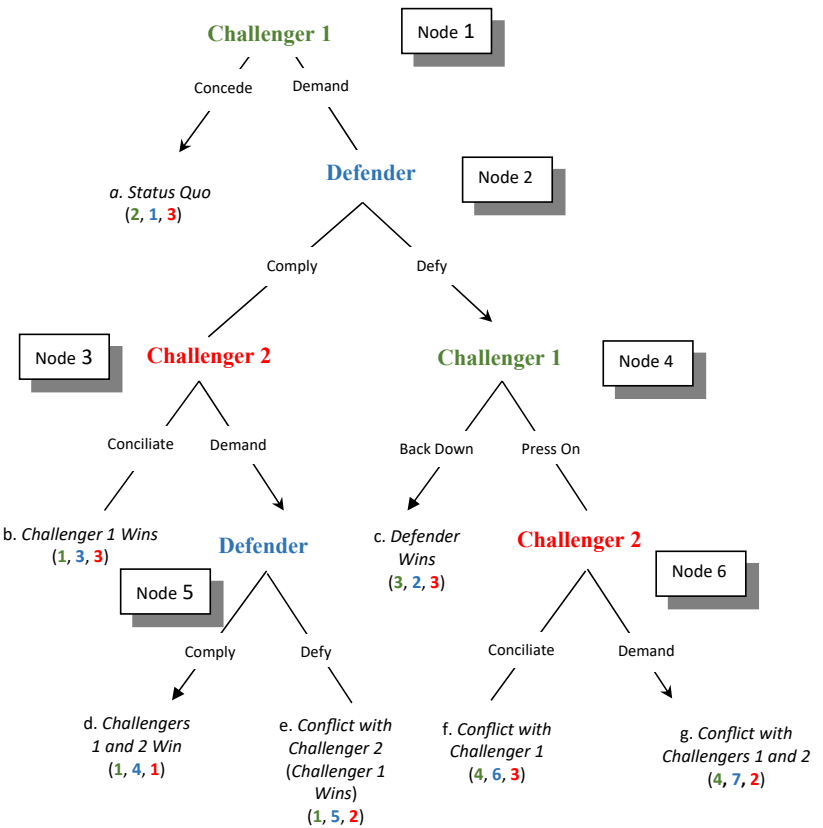
There are four distinct types of Defenders: *Staunch*, *Semi-Staunch*, *Submissive*, and *Semi-Submissive*; three types of Challenger 1s: *Hesitant*, *Determined*, and *Cautious*; and three types of Challenger 2s: *Persistent* and *Opportunistic*, *Restrained and Opportunistic*, and *Restrained* and *Reluctant*. Thus, there are $4 \times 3 \times 3 = 36$ different combinations of player types. Clearly, the addition of a second challenger to the game adds a great amount of complexity to what would otherwise be a very simple strategic relationship between one defender and one challenger. Were Challenger 2's choices at Nodes 3 and 6 eliminated, the *Double Deterrence Game* would reduce to a game-form that Zagare and Kilgour (2000: ch. 5) call the *Unilateral Deterrence Game*. In this truncated game there are only two types of Defenders, only two types of Challengers, and therefore, but $2 \times 2 = 4$ games. In three of these games the *Status Quo* holds, and deterrence succeeds. Significantly, deterrence fails in the fourth case: Challenger makes a demand, and Defender capitulates.

3 The Double Deterrence Game with Complete Information

In this section I consider the Double Deterrence Game with complete information. The standard measure of rational play in a dynamic (or extensive-form) game with complete information is *subgame perfect equilibrium* (Selten 1975). A subgame perfect equilibrium requires that the players plan to choose rationally at every node of a game tree whether they expect to reach a particular node or not. The concept of a subgame perfect equilibrium is a refinement of Nash's well-known equilibrium concept, which is the accepted measure of rational behavior in a static (or strategic-form) game. Nash equilibria, however, may be supported by irrational threats, that is, by threats that are not believable or credible. Selten's perfectness criterion eliminates that possibility.

In an extensive-form game of complete and perfect information in which the players know not only their place on the game tree at all times but are also fully informed about each other's preferences, a simple procedure known as *backwards induction* can be used to identify which outcomes are subgame perfect. As its name suggests, backwards induction involves working backwards up the game tree to determine, first, what a rational player would do at the last (or terminal) node (or nodes) of the game tree, what the player with the previous move would do given that the player with the last move is assumed to be rational, and so on until the first (or initial) node of the tree is reached. Outcomes that survive the backwards induction process are, by definition, subgame perfect.

To illustrate, consider for now the variant of the Double Deterrence Game depicted in Figure 2. In this representation the assumption is that Defender is *Submissive*. Recall that a *Submissive* Defender prefers the outcome *Challenger 1 Wins* to all three conflict outcomes. In other words, whenever Defender is *Submissive*, both of its deterrent threats lack credibility. The game of Figure 2 also assumes that Challenger 2 is the *Persistent/Opportunistic* type. From Defender's point of view, the most problematic secondary challengers are of this type. Finally, this version of the Double Deterrence Game assumes that Challenger 1 is *Hesitant*. Like a *Submissive*



Key: (x, y, z) = payoff to Challenger 1, payoff to Defender, payoff to Challenger 2
1 = best, 2 = next best; *etc*.
→ = Rational choice

Figure 2: The Double Deterrence Game when Defender is Submissive, Challenger 1 is Hesitant, and Challenger 2 is Persistent and Opportunistic.

Defender, a *Hesitant* Challenger 1's deterrent threat lacks credibility. These assumptions are reflected in the three-tuple of ordinal utilities beneath each of the outcomes in Figure 2. The outcomes, for Challenger 1 (marked in green). Defender (marked in blue), and Challenger 2 (marked in red) are ranked from best to worst. For example, the *Status Quo* is Challenger 1's second best outcome (i.e. 2), Defender's best (i.e. 1), and one of Challenger 2's four worst outcomes (i.e. 3).[8]

We begin by considering the calculus of Challenger 1 at the first node of the tree. At Node 1 Challenger 1 can either do nothing (i.e., *Concede*) which brings about its second ranked outcome, the *Status Quo*, or insist on an adjustment of it. Of course, the consequences of demanding a change in the *Status Quo* are uncertain since they depend, at least in part, on Defender's likely response at Node 2. Defender choice, in turn, depends on what Challenger 2 will do if it complies, and on the decisions of both challengers if it defies Challenger 1 at Node 2.

If Defender complies, Challenger 1 will win regardless of what Challenger 2 does at Node 3. But, given complete information, Defender will know that compliance at Node 2 will eventually lead to its fourth best outcome, *Challengers 1 and 2 Win*. The reason is manifest: as the arrows indicate, at Node 5, a *Submissive* Defender will rationally comply with Challenger 2's demand, in order to avoid its fifth best outcome, *Conflict with Challenger 2 (Challenger 1 Wins*). And given that Challenger 2 is *Persistent*, it will force that choice on Defender by rationally issuing its own demand at Node 3. By contrast, defiance at Node 2 will induce Defender's second-best outcome, *Defender Wins*. As the arrows indicate, a *Hesitant* Challenger 1 will back down at Node 4 since, by assumption, it prefers the outcome *Defender Wins* to either outcome that would come about if it pressed on.

Under the conditions that define the game in Figure 2, then, Defender's choice at Node 2 and Challenger 1's choice at Node 1, is straightforward. At Node 2, Defender rationally defies Challenger 1 in order to induce its second-best outcome, *Defender Wins*, and to avoid its fourth best-outcome, *Challengers 1 and 2 Win*. Similarly, at Node 1, Challenger 1 does nothing (i.e. *Concedes*), to avoid its second-worst outcome, *Defender Wins*. Instead, deterrence succeeds when Challenger 1 rationally settles for its second-best outcome, the *Status Quo*.

As Table 4 shows, whenever Challenger 1 is *Hesitant* and lacks a credible threat to *Press On* at Node 4, deterrence always succeeds, and the *Status Quo* always holds, even if Defender is *Submissive,* that is, when all of its threats lack credibility. A credible threat to retaliate, then, does not constitute a necessary condition for deterrence success in a game of complete information. For example, Defender's threat to defy Challenger 1 in the game depicted in Figure 2 lacks credibility (i.e. it prefers to submit to Challenger 1's demand *should* Challenger 1 *Press On* at Node 4).

8 See Footnote 7.

Table 4: Subgame perfect equilibria in the double deterrence game.

		Defender											
		Staunch			*Semi-Staunch*			*Submissive*			*Semi-Submissive*		
		Challenger 1											
		Hesitant	*Determined*	*Cautious*	*Hesitant*	*Determined*	*Cautious*	*Hesitant*	*Determined*	*Cautious*	*Hesitant*	*Determined*	*Cautious*
Challenger 2	*Persistent/ Opportunistic*	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Challenger 1 and 2 Win	Challenger 1 and 2 Win	Status Quo	Conflict (Defender Challenger 2)	Conflict (Defender Challenger 2)
	Restrained/ Opportunistic	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Challenger 1 and 2 Win	Challenger 1 and 2 Win	Status Quo	Challenger 1 Wins	Challenger 1 Wins
	Restrained/ Reluctant	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Status Quo	Challenger 1 and 2 Win	Status Quo	Status Quo	Challenger 1 Wins	Status Quo

But since Challenger 1's threat to *Press On* also lacks credibility, it cannot deter Defender from resisting its demand at Node 2.

It is important to note that deterrence will also succeed whenever Defender is *Staunch* or *Semi-Staunch*, regardless of either Challenger's type. Finally, if Challenger 1 is *Cautious* and Challenger 2 is *Restrained and Reluctant*, Challenger 1 will refrain from contesting the *Status Quo* and deterrence will be successful.

Under most conditions, then, the *Status Quo* holds in the Double Deterrence Game. But when is it likely to fail? *Determined* Challenger 1s will never be deterred if Defender is either *Submissive* or *Semi-Submissive*. And, unless Challenger 2 is both *Restrained and Reluctant* it will also fail even when Challenger 1 is *Cautious*.

Of course, Defender's type will determine the specific resolution of the game in the event of a deterrence failure. Both Challengers will win if Defender is *Submissive*; Challenger 1 alone will win if Defender is *Semi-Submissive* and Challenger 2 is *Restrained*. Challenger 1 will still win if Defender is *Semi-Submissive*, but an all-out conflict with Challenger 2 will result when Challenger 2 is *Persistent*.

Significantly, however, a rational conflict between a defender of the status quo and its primary challenger is precluded. There are three keys to deterrence success: the *Status Quo* will hold (1) if Defender's threat to retaliate is credible, or (2) if the primary challenger's threat to press on lacks credibility, or (3) if the primary challenger's willingness to press on depends on support from a reluctant secondary challenger. Of course, a Defender of the *Status Quo* should always work to maintain a credible retaliatory threat. But there may be situations where the stakes are such that Defender's threat to resist a challenge is hardly believable. It follows, therefore, that another focus of policy for the Defender should aim to eliminate, or at least reduce, the primary challenger's willingness to risk a confrontation. Finally, a third policy objective could aim to weaken the secondary challenger's propensity to exploit a crisis between Challenger 1 and a Defender. But this should not be the main thrust of Defender's policy given the restricted set of circumstances in which the secondary challenger's participation in a crisis actually matters. Of the 36 games with complete information, there are only 4 games in which the secondary challenger's type makes a significant difference (see Table 4).

4 The Double Deterrence Game with Incomplete Information

In the previous section we saw that the secondary challenger played a relatively minor role in the dénouement of the Double Deterrence Game with complete information, and concluded that in a real-world situation that satisfies the model's assumptions, Defender's focus should be on its relationship with its primary

challenger. Notably, three outcomes never occur when information is complete: *Defender Wins* and the two conflict outcomes that involve Challenger 1. Of course, we know, empirically, that sometimes *faits accompli* fail and that all-out wars occur. To understand the conditions under which these and other outcomes may rationally occur, we next examine a simplified version of the Double Deterrence Game with incomplete information (see Figure 3). Since Challenger 2's choice at Node 3 is immaterial to Challenger 1, and its choice at Node 6 only matters when Challenger 1 is *Cautious*, its role will be suppressed, at least for now. Thus, only two types of Challenger 1s, *Hesitant* and *Determined*, and two types of Defender, *Staunch* and *Submissive*, are considered.[9] In the analysis of the game of Figure 3, the assumption

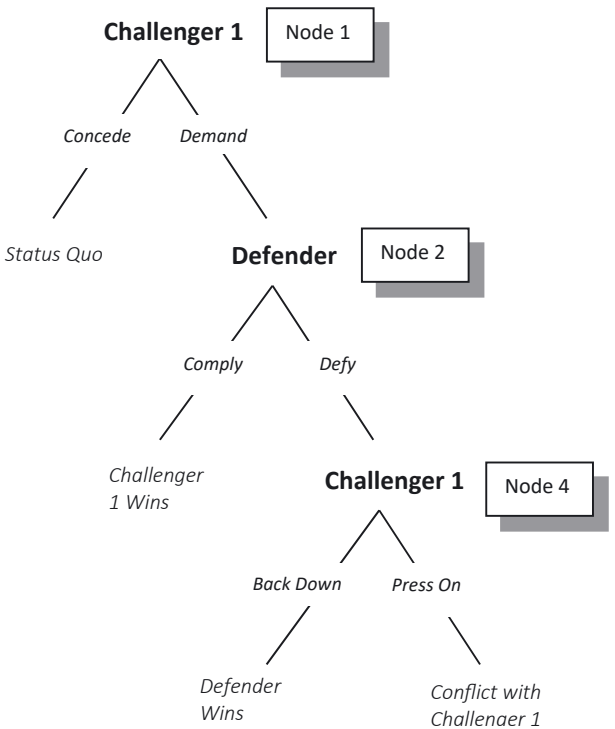


Figure 3: Simplified version of the Double Deterrence Game.

9 As noted previously, Zagare and Kilgour (2000: ch. 5) call this truncated game the *Unilateral Deterrence Game*. Quackenbush (2010, 2011) finds strong empirical support for the equilibrium predictions of this game. The Challenger 1 – Challenger 2 subgame starting at Node 4 is analyzed below (see Figure 5).

will be that each player knows its own type but has only probabilistic knowledge of its opponent's type, and that these probabilities are common knowledge. Specifically, Defender is believed to be *Staunch* with probability p_{Def} and *Hesitant* with probability $(1 - p_{Def})$. Similarly, Challenger 1 is assumed to be *Determined* with probability p_{Ch}, and *Hesitant* with probability $(1 - p_{Ch})$.

The standard measure of rational choice in a game with incomplete information is *perfect Bayesian equilibrium*. A perfect Bayesian equilibrium specifies an action choice for every type of every player at every decision node (or information set) belonging to the player; it must also indicate how each player updates its beliefs about the other player's type in the light of new information obtained as the game is played out.[10]

As it turns out, there are five major types of perfect Bayesian equilibria in the simplified version of the Double Deterrence Game with incomplete information (see Table 5).[11] Two are deterrence equilibria in which the *Status Quo* is never contested. The good news is that one of them always exists. The *Certain Deterrence Equilibrium* exists whenever the credibility of Defender's threat to resist a demand at Node 2 is sufficiently high, that is, when p_{Def} is large and Defender is most likely *Staunch*, regardless of Challenger 1's type (see Figure 4). A *Steadfast Deterrence Equilibrium* exists only when the credibility of Defender's threat falls below the threshold (c_t) that sustains the Certain Deterrence Equilibrium. When it exists, the Certain Deterrence Equilibrium exists uniquely, as one might expect from an examination of the game under complete information. Unfortunately, the Steadfast Deterrence Equilibrium will always co-exist with one of the three remaining Perfect Bayesian equilibria under which an all-out conflict remains a distinct possibility.[12] That is the bad news.

There are two rational strategic possibilities when the credibility of Defender's threat is low (i.e. when $p_{Def} < c_s$), that is, when Defender is likely *Submissive*: a *Bluff Equilibrium* and an *Attack Equilibrium*. The Bluff Equilibrium exists when Challenger

10 In an extensive-form game of incomplete information, the initial (or *a priori*) beliefs of the players are taken as givens. The assumption is that the players update their beliefs rationally (i.e., according to Bayes' rule) given the actions it observes during the play of the game. The definition of a perfect Bayesian equilibrium, however, places no restriction on the players' updated (or *a posteriori*) beliefs "off the equilibrium path," that is, on beliefs at nodes that are never reached under rational play. It is sometimes the case that a perfect Bayesian equilibrium is supported by *a posteriori* beliefs that are inconsistent with a player's *a priori* beliefs. Perfect Bayesian equilibria that are based on internally inconsistent beliefs are implausible. In consequence, they are not considered as rational strategic possibilities for the purposes of this essay.

11 For a proof and additional technical details, see Zagare and Kilgour (2000, Appendix 5).

12 At either deterrence equilibrium, the *Status Quo* is the only possible outcome. Although the *Status Quo* can sometimes result when another equilibrium is in play, all remaining equilibria carry with them the possibility of other outcomes, depending on the players' types, beliefs, and choices.

Table 5: Action Choices of Perfect Bayesian Equilibria for the Simplified Version of the Double Deterrence Game with Incomplete Information.

Equilibrium	Strategy profiles				*Existence conditions*
	Challenger		Defender		
	x_{Det}	x_{Hes}	y_{Sth}	y_{Sub}	
Certain deterrence	0	0	1	Unrestricted	$p_{\text{Def}} \geq c_{\text{t}}$
Steadfast deterrence	0	0	1	*u*	$p_{\text{Def}} < c_{\text{t}}$
Separating	1	0	1	0	$c_{\text{s}} \leq p_{\text{Def}} \leq c_{\text{t}}$
Bluff	1	*v*	1	*u*	$p_{\text{Def}} < c_{\text{s}}$ and $p_{\text{Ch}} \leq d_{\text{n}}$
Attack	1	1	1	0	$p_{\text{Def}} < c_{\text{s}}$ and $p_{\text{Ch}} \geq d_{\text{n}}$

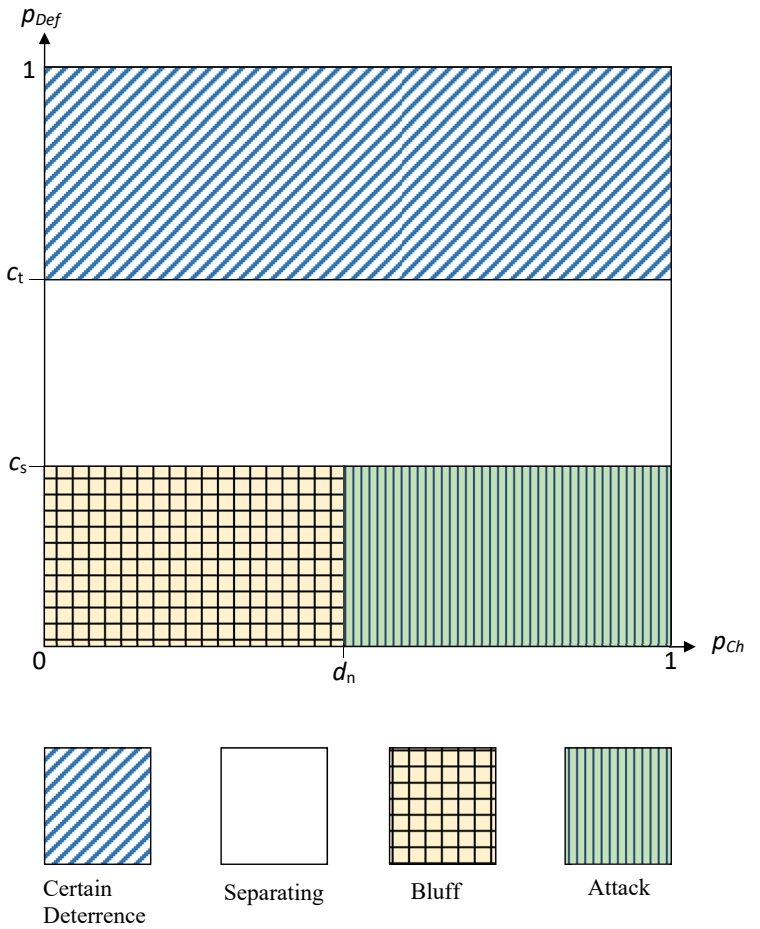
Key: x_{Det} = the probability that a *Determined* Challenger will issue a demand at Node 1, x_{Hes} = the probability that a *Hesitant* Challenger will issue a demand at Node 1, y_{Sth} = the probability that a *Staunch* Defender will choose to defy Challenger 1 at Node 2, y_{Sub} = the probability that a *Submissive* Defender will choose to defy Challenger 1 at Node 2, "*u*" = fixed value between 0 and 1, "*v*" = fixed value between 0 and 1, c_t, c_s, and d_n = constant threshold values that distinguish the existence regions of the Perfect Bayesian Equilibria. Source: (Zagare and Marc Kilgour 2000, Ch. 5).

1 is most likely *Hesitant* (i.e. when $p_{Ch} \leq d_{\text{n}}$). When Challenger 1 is more likely to be *Determined*, (i.e. when $p_{Ch} \geq d_{\text{n}}$), the Attack Equilibrium governs play.

The *Status Quo* is the most likely outcome when play takes place under the conditions that support the existence of the Bluff Equilibrium. After all, Challenger 1 is most likely *Hesitant*. Nonetheless, since Defender is also most likely the *Submissive* type, even a *Hesitant* Challenger 1 *sometimes* issues a demand. If and when this happens, *sometimes* Defender will defy Challenger 1, who most likely will back down and the outcome will be *Defender Wins*. But if Defender complies with the demand, *Challenger 1 Wins*. *Conflict* is also a possibility under a Bluff Equilibrium, but this possibility is remote, at best.

The *Status Quo* is the only outcome that is precluded under the conditions associated with the existence of the Attack Equilibrium. Of the remaining outcomes, *Challenger 1 Wins* is, by far, the most likely. Since Challenger 1 is more likely than not to be *Determined*, it is certain to issue a demand at Node 1. But since Defender is likely to be *Submissive*, even a *Hesitant* Challenger 1 will contest the *Status Quo* with certainty. All of which is to say that deterrence is not even a remote possibility when an Attack Equilibrium is in play. Of course, if Defender is *Staunch*, it will always resist the demand and the outcome will be *Conflict* unless Challenger 1 is actually *Hesitant*, in which case the outcome will be *Defender Wins*. Nonetheless, the possibility of either of these two rational strategic possibility is relatively low.

The final rational strategic possibility is the *Separating Equilibrium*. Under a Separating Equilibrium Defender's credibility is not high enough to support a Certain



Key:

p_{Def} = the probability Defender is Staunch
p_{Ch} = the probability that Challenger is *Determined*
c_t, c_s, and d_n = constant threshold values that distinguish the existence regions of the Perfect Bayesian Equilibria

Figure 4: Location of perfect Bayesian equilibria for the simplified version of the Double Deterrence Game with incomplete information subgame.

Deterrence Equilibrium, but not so low that either a Bluff or an Attack Equilibrium can exist (i.e. when $c_s \leq p_{Def} \leq c_t$). At a Separating Equilibrium, the players' preferences are fully revealed by their strategy choices: a *Determined* Challenger 1 always issues a demand, and a *Hesitant* Challenger never does. Likewise, if challenged, a *Staunch* Defender always defies and a *Submissive* Defender always capitulates. The

Status Quo may remain stable, therefore, when separating strategies are selected, but only when Challenger 1 is *Hesitant*. When Challenger 1 is *Determined*, it gains an advantage if Defender is *Submissive* (i.e. *Challenger 1 Wins*) but precipitates *Conflict* if Defender is *Staunch*. Thus, three of the four possible outcomes of the game can arise under a Separating Equilibrium. As explained below, it is under the conditions that support the existence of a Separating Equilibrium that Challenger 2 may have a significant impact on whether deterrence succeeds or fails.

To summarize briefly: Under the Certain Deterrence Equilibrium, the *Status Quo* is the only possible outcome. The *Status Quo* is also a possible outcome when a Steadfast Deterrence Equilibrium exists. But since this equilibrium form will always co-exist with one of the three remaining Perfect Bayesian Equilibria, there will always be other possibilities. These possibilities, in turn, depend on which of the three Perfect Bayesian Equilibria also exist. Under the Separating Equilibrium, only the outcome *Defender Wins* is precluded. Under the Bluff Equilibrium, the *Status Quo* will most likely hold, whether play takes part under it or under the Steadfast Deterrence Equilibrium. And if and when the Attack Equilibrium comes into play, deterrence always fails, and Challenger 1 will most likely win.

5 Subgame Analysis

To this point we have examined the double deterrence game under complete information, and a focused version of it under incomplete information. In this section, I examine the Challenger 1 – Challenger 2 subgame[13] that begins at Node 4 (see Figure 5). Any perfect Bayesian equilibrium in the truncated version of the Double Deterrence Game must be consistent with rational play in this and all of its other subgames, relevant or not. However, the Challenger 1 – Challenger 2 subgame, and the proper subgame that begins at Node 6 which it subsumes, are the only subgames of the Double Deterrence game that have a bearing on Challenger 1's Node 1 decision.[14]

In the Challenger 1 – Challenger 2 subgame, Challenger 2's choice at Node 6 is strictly determined by its type. A *Reluctant* Challenger 2 will always conciliate Defender and the outcome will be a bilateral conflict between Defender and Challenger 1. But if Challenger 2 is *Opportunistic*, it will issue its own demand and join Challenger 1 in a entirely separate conflict with Defender. We know that, at Node 4, a *Determined* Challenger 1 will always press on and the outcome will depend on

13 A subgame is that part of an extensive-form game that can be considered a game unto itself.

14 For an analysis of the Challenger 2 – Defender subgame that begins at Node 3, see Zagare (2019: 18–14).

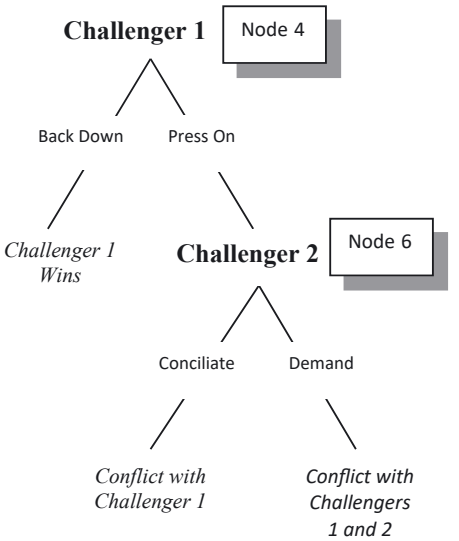


Figure 5: The Challenger 1 – Challenger 2 subgame.

Challenger 2's type, as just explained; and if Challenger 1 is *Hesitant*, it will always back down at Node 4 and the outcome will be *Defender Wins*. All of which is to say that the dynamics of this subgame are clear when Challenger 1 is either *Determined* or *Hesitant*.

When Challenger 1 is *Cautious*, however, things are different. Not surprisingly, the key here is the probability that Challenger 2 is *Reluctant*. When that probability is high enough, Challenger 1 will mimic the behavior of a *Hesitant* Challenger 1; when it is low, Challenger 1 will act as if it were *Determined*. All of which is to say that, given uncertainty and depending on Challenger 1's estimate of Challenger 2's type, actual play in the Challenger 1 – Challenger 2 subgame when Challenger 1 is *Cautious* will be all-but indistinguishable from play when it is either *Determined* or *Hesitant*. How it plays out will have no effect on Defender's choice at Node 2 in the larger game. *Staunch* and *Semi-Staunch* Defenders will defy Challenger 1 at Node 2 no matter how the Challenger 1 – Challenger 2 subgame plays out, while *Submissive* and *Semi-Submissive* Defenders will always comply with Challenger 1's demand should it make one. But it may have an impact on Challenger 1's Node 1 choice and whether or not deterrence succeeds.

In the Double Deterrence Game with incomplete information, *Determined* Challenger 1s are assumed to have highly credible threats, while *Hesitant* Challenger 1s have threats that are barely credible. It is natural, then, to interpret a *Cautious* Challenger 1's credibility as lying between these two extremes, that is

at middling levels of credibility. Assuming this to be the case, there are two situations where Challenger 2's type, and Challenger 1's estimate of Challenger 2's type (credibility) are salient. The first is in the area that separates the Bluff Equilibrium from the Attack Equilibrium, that is when $p_{\text{Ch}} \approx d_{\text{n}}$; the second is under the conditions that support the existence of the Separating Equilibrium, that is when $c_{\text{s}} \leq p_{\text{Def}} \leq c_{\text{t}}$ (see Figure 4).

As previously mentioned, the Bluff and Attack Equilibria exist only when Defender's credibility is very low, that is when $p_{\text{Def}} \leq c_{\text{s}}$. When the credibility of Challenger 1's threat is also low, when $p_{\text{Ch}} \leq d_{\text{n}}$, the Bluff Equilibrium will exist, but when it is high, that is when $p_{\text{Ch}} \geq d_{\text{n}}$, play will take place under the constraints of the Attack Equilibrium. If Challenger 1 is *Cautious* under these conditions, a big if, it is unlikely that Challenger 2's type will have a significant impact on Challenger 1's Node 1 decision. Recall that under the Bluff Equilibrium, even a *Hesitant* Challenger 1 may rationally contest the *Status Quo*, and under the Attack Equilibrium, a *Hesitant* Challenger 1 will do so with certainty. Thus, in the area around the threshold value (d_{n}) that separates these two equilibrium forms, Challenger 1's propensity to make a demand at Node 1 will be slightly higher under Bluff conditions and slightly lower under Attack conditions. The cumulative impact of Challenger 2's type, therefore, is likely to be a one marginal, at best.

Challenger 1 is much more likely to be *Cautious* when a Separating Equilibrium exists than when either a Bluff Equilibrium or an Attack equilibrium exists, that is, when Defender's credibility is at a middling level, too low to support the existence of a Certain Deterrence Equilibrium, but high enough to avoid play under either the Bluff or the Attack Equilibrium. Under a Separating Equilibrium *Determined* Challengers always issue a demand at Node 1, *Hesitant* Challengers never do. *Cautious* Challengers can go either way. It is under these conditions that Challenger 2's type is most likely to play a major role in determining Challenger 1's Node 1 choice.

It is important to note, however, that a prior condition for this to be the case is that Challenger 1 is the *Cautious* type, suggesting once again that the relationship between a defender and its primary challenger will be the key to deterrence success, or failure, in a deterrence game with two potential challengers acting autonomously. The danger zone, then, is at intermediate levels of Defender's *and* Challenger 1's credibility, that is when $c_{\text{s}} \leq p_{\text{Def}} \leq c_{\text{t}}$ and $p_{\text{Ch}} \approx d_{\text{n}}$ respectively. Assuming a rising Challenger 1, or a declining Defender, or both, this zone will exist just prior to, or immediately after, Challenger 1 has achieved parity with the Defender, that is, when a balance of capabilities exists (Organski and Kugler 1980).

6 Summary and Conclusion

This essay explores a deterrence relationship between one defender and two challengers that are interconnected game-theoretically but are otherwise acting independently. In the Double Deterrence game model, the primary challenger makes the first move, defender, the second, and the secondary challenger makes its move only after the defender and the primary challenger have made their choices. In other words, the threat of the secondary challenger is both latent and contingent. There are four types of defenders, and three types of the two challengers. These preference assumptions define 36 distinct strategic environments, or games. These games are analyzed under both complete and incomplete information.

The results are highly sensitive to the specific combination of player types and the credibility of their deterrent threats. Deterrence is most likely to succeed when a defender's threat is highly credible, and to fail when it is low. The secondary challenger is most likely to have an impact on play only at intermediate levels of the defender's *and* the primary challenger's credibility. Assuming a rising challenger, or a declining defender, or both, a danger zone will exist just prior to, or immediately after, the primary challenger has achieved parity with the defender, that is, when a balance of capabilities exists. The principal aim of a defender of the status quo facing two challengers acting independently (a major qualification), then, should be on preventing a challenge by its primary opponent. This conclusion is robust no matter what configuration of player types one assumes. While this conclusion is limited to the special case under consideration – the three-body problem – it in no way is inconsistent with the major findings and the policy recommendations of Perfect Deterrence Theory which, for convenience, are summarized and contrasted with those of Classical Deterrence Theory in Table 6.

Table 6: Classical deterrence theory and perfect deterrence theory: policy prescriptions.

	Classical deterrence theory	**Perfect deterrence theory**
Policies:		
Overkill capability	Supports	Opposes
Minimum deterrence	Opposes	Supports
"Significant" arms reductions	Opposes	Supports
Proliferation	Supports	Opposes
Negotiating stances	Coercive, based on increasing war costs and inflexible bargaining tactics	Conditionally cooperative, based on reciprocity

Appendix

Tables A1 and A2

Table A1: List of possible outcomes.

Outcome
a. *Status Quo*
b. *Challenger 1 Wins*
c. *Defender Wins*
d. *Challengers 1 and 2 Win*
e. *Conflict with Challenger 2* (*Challenger 1 Wins*)
f. *Conflict with Challenger 1*
g. *Conflict with Challengers 1 and 2*

Table A2: Restrictions on preferences.

Types of players	Restrictions on preferences
All types of defender:	$a \succ c$; $b \succ (d, e)$; $f \succ g$
Staunch	$(f, g) \succ (b, d, e)$; $e \succ d$
Semi-Staunch	$(f, g) \succ (b, d, e)$; $d \succ e$
Submissive	$(b, d, e) \succ (f, g)$; $d \succ e$
Semi-Submissive	$(b, d, e) \succ (f, g)$; $e \succ d$
All types of Challenger 1:	$b \sim d \sim e \succ a \succ (c, f, g)$
Hesitant	$c \succ f \sim g$
Determined	$f \sim g \succ c$
Cautious	$g \succ c \succ f$
All types of Challenger 2:	$d \succ (b, e, f, g)$
Persistent	$e \succ b$
Restrained	$b \succ e$
Reluctant	$f \succ g$
Opportunistic	$g \succ f$

References

Altfeld, Michael F., and Bruce Bueno de Mesquita. 1979. “Choosing Sides in Wars.” *International Studies Quarterly* 23: 87–112.

Brodie, Bernard. 1959. “The Anatomy of Deterrence.” *World Politics* 11: 173–9.

Intriligator, Michael D., and Dagobert L. Brito. 1984. “Can Arms Races Lead to the Outbreak of War?” *Journal of Conflict Resolution* 28: 63–84.

Iwanami, Yukari. 2023. “Asymmetric Burden-sharing and the Restraining and Deterrence Effects of Alliances.” *Journal of Peace Research*. https://doi.org/10.1177/00223433231158146.
Krainin, Colin, and Thomas Wiseman. 2016. “War and Stability in Dynamic International Systems.” *The Journal of Politics* 78: 1139–52.
Lakatos, Imre. 1970. “Falsification and the Methodology of Scientific Research Programs.” In *Criticism and the Growth of Knowledge*, edited by Imre Lakatos, and Alan Musgrave. Cambridge: Cambridge University Press.
Mearsheimer, John J. 1983. “The Case for a Ukrainian Nuclear Deterrent.” *Foreign Affairs* 72: 50–66.
Organski, A.F.K., and Jacek Kugler. 1980. *The War Ledger*. Chicago: University of Chicago Press.
Overy, Richard. 2021. *Blood and Ruins: The Last Imperial War, 1931 – 1945*. New York: Penguin Books.
Powell, Robert. 2017. “Taking Sides in Wars of Attrition.” *American Political Science Review* 111: 219–36.
Quackenbush, Stephen L. 2006. “Not Only Whether but Whom: Three-Party Extended Deterrence.” *Journal of Conflict Resolution* 50: 562–83.
Quackenbush, Stephen L. 2010. “General Deterrence and International Conflict: Testing Perfect Deterrence Theory.” *International Interactions* 36: 60–85.
Quackenbush, Stephen L. 2011. *Understanding General Deterrence: Theory and Application*. New York: Palgrave Macmillan.
Rasmusen, Eric. 1989. *Games and Information: An Introduction to Game Theory*, 2nd ed. Cambridge: Blackwell.
Schelling, Thomas C. 1966. *Arms and Influence*. New Haven: Yale University Press.
Schweller, Randall. 1998. *Deadly Imbalances: Tripolarity and Hitler’s Strategy of World Conquest*. New York: Columbia University Press.
Selten, Reinhard. 1975. “A Re-examination of the Perfectness Concept for Equilibrium Points in Extensive Games.” *International Journal of Game Theory* 4: 25–55.
Snyder, Glenn H., and Paul Diesing. 1977. *Conflict Among Nations: Bargaining, Decision Making and System Structure in International Crises*. Princeton: Princeton University Press.
Waltz, Kenneth N. 2012. “Why Iran Should Get the Bomb: Nuclear Balancing Would Mean Stability.” *Foreign Affairs* 91: 2–5.
Zagare, Frank C. 1996. “Classical Deterrence Theory: A Critical Assessment.” *International Interactions* 21: 365–87.
Zagare, Frank C. 2019. *Game Theory, Diplomatic History and Security Studies*. Oxford: Oxford University Press.
Zagare, Frank C., and D. Marc Kilgour. 2000. *Perfect Deterrence*. Cambridge: Cambridge University Press.
Zagare, Frank C., and D. Marc Kilgour. 2003. “Alignment Patterns, Crisis Bargaining, and Extended Deterrence: A Game-Theoretic Analysis.” *International Studies Quarterly* 47: 587–615.