Stabilizing unstable outcomes in prediction games

Brams, Steven and Kilgour, Marc

New York University, Wilfrid Laurier University

March 2017

Stabilizing Unstable Outcomes in Prediction Games

Steven J. Brams
Department of Politics
New York University
New York, NY 10012
UNITED STATES
steven.brams@nyu.edu

D. Marc Kilgour
Department of Mathematics
Wilfrid Laurier University
Waterloo, Ontario N2L 3C5
CANADA
mkilgour@wlu.ca

Abstract

Assume in a 2-person game that one player, Predictor (P), does not have a dominant strategy but can predict with probability $p > 1/2$ the strategy choice of an opponent, Predictee (Q). Q chooses a strategy that maximizes her expected payoff, given that she knows p—but not P's prediction—and that P will act according to his prediction. In all 2×2 strict ordinal games in which there is a unique Pareto-inferior Nash equilibrium (Class I) or no pure-strategy equilibrium (Class II), and which also has a Pareto-optimal non-Nash "cooperative outcome," P can induce this outcome if p is sufficiently high. This scenario helps to explain the observed outcomes of a Class I game modeling the 1962 Cuban missile crisis between the United States and the Soviet Union, and a Class II game modeling the 2015 conflict between Iran and Israel over Iran's possible development of nuclear weapons.

1. Introduction

In this paper, we analyze a set of 2-person games, which we call "prediction games," in which a one player, Predictor (P) (assumed male), does not have a dominant strategy but can predict with probability $p > 1/2$ the strategy choice of the other player, Predictee (Q) (assumed female), who may or may not have a dominant strategy. In our model, Q knows p and that P will act according to his prediction, and she chooses a strategy to maximize her expected payoff.

For simplicity, we restrict most but not all our analysis to 2×2 strict ordinal games, in which each player has two strategies and there are four possible outcomes. The players strictly rank the four outcomes from best to worst, so each player's payoffs, which are expressed in cardinal utilities that reflect their rankings, are different.

Of the 57 2×2 distinct "conflict games," wherein there is no mutually best outcome, we focus on the 32 (56%) in which at most one player has a dominant strategy.[1] In these games, P is the player or one of the players who lacks a dominant strategy. Not having such a strategy, P is motivated to try to predict Q's choice in order to determine his own optimal choice.

We show that in exactly eight of these 32 games (25%), which fall into two classes, good predictions by P can obviate two significant problems, thereby enabling the players to implement a Pareto-optimal outcome in these games that is not a pure-strategy Nash equilibrium but is at least the next-best outcome for both players:

---

[1] For a complete listing of the 57 games, see Brams, 1994, 2011). Bruns (2015) offers a different listing and classification of these games. Later, when we distinguish between the roles of P and Q, we will see that, except for six symmetric games, the 51 asymmetric games bifurcate into two games each, depending on which player is P and which is Q.

[2] In Newcomb's problem, there may be a conflict between maximizing one's expected payoff and choosing a dominant strategy, which is analyzed in Brams (1975) when Newcomb's problem is treated as a game. The Class II prediction games discussed in section 4 pose this problem, wherein we assume Q's choice is

- Class I (5 games): There is no Nash equilibrium in pure strategies; furthermore, the equilibrium outcome in mixed strategies is Pareto-inferior to the Pareto-optimal non-Nash outcome.
- Class II (3 games): There is a unique pure-strategy Nash equilibrium outcome, but it is Pareto-inferior to the Pareto-optimal non-Nash outcome.

For each class, we show that good predictions by P, when incorporated into the players' rational-choice calculations, stabilize the Pareto-optimal non-Nash outcome, which we call a "cooperative outcome" because it is at least the next-best outcome for both players. Thus, in these problematic games, there is a way of achieving outcomes preferable for both players to the Nash-equilibrium outcome (pure or mixed).

We analyze two international conflicts, the first modeled by a Class I game and the second by a Class II game. The first is between the United States and the Soviet Union during the 1962 Cuban missile crisis, the second is between Iran and Israel over Iran's possible development of nuclear weapons in 2015.

In both conflicts, a cooperative outcome halted further escalation, and possibly war, between the antagonists. In the specific prediction game modeling each conflict, one country's (P's) prediction about what its opponent (Q) would do seems to have played a major, and if not decisive, role in attenuating the conflict and inducing both countries to reach a settlement.

An important lesson from our analysis is that even in the two classes of troublesome games in which the unique cooperative outcome cannot be supported in equilibrium, there is hope for achieving it if at least one player (P) can predict with a high enough probability the behavior of an opponent (Q). Knowing this, Q will choose a cooperative strategy, which P will have predicted with this probability and Q, in turn, will have used in deciding to cooperate. Making available to policy makers better information about the choices of adversaries, and their likelihood, will tend to foster interlocking cooperative decisions.

2. The 2 × 2 Model

Assume Q has two strategies, s_1 and s_2, and P has two strategies, t_1 and t_2, leading to the four possible outcomes in the 2 × 2 game shown in Figure 1. Q's payoffs, a_{ij}, and P's payoffs, b_{ij}, at each outcome are assumed to be cardinal utilities (see Figure 1).

Figure 1. 2 × 2 Payoff Matrix

$$
\begin{array}{cc|cc}
 & & \multicolumn{2}{c}{\boldsymbol{P}} \\
 & & t_1 & t_2 \\
\hline
 & s_1 & (a_{11}, b_{11}) & (a_{12}, b_{12}) \\
\boldsymbol{Q} & & & \\
 & s_2 & (a_{21}, b_{21}) & (a_{22}, b_{22})
\end{array}
$$

Assume that P does not have a dominant strategy. Then the strategies can be arranged so that P prefers outcomes on the main diagonal (from upper left to lower right) to those that he could switch to on the off-diagonal (from upper right to lower left). In other words,

$$b_{11} > b_{12} \text{ and } b_{22} > b_{21}.$$

Thus, P's preferred strategy, t_1 or t_2, depends on Q's strategy choice (P prefers t_1 if Q chooses s_1, t_2 if Q chooses s_2).

Now suppose that P, based on information or intelligence that he possesses, can correctly predict Q's strategy choice, s_1 or s_2—in advance of Q's making it—with probability p. Knowing p, assume that Q chooses her strategy, s_1 or s_2, to maximize her expected payoff.[2]

Thus, if P predicts Q will choose s_1, he will choose t_1; if P predicts Q will choose s_2, he will choose t_2. Consequently, if Q knows that P will predict her choice with perfect accuracy (i.e., $p = 1$), then Q does best by choosing her strategy, s_1 or s_2, consistent with her, but not P's, preference between the two diagonal outcomes, (a_{11}, b_{11}) and (a_{22}, b_{22}). If P and Q prefer different diagonal outcomes, then P's ability to make perfect predictions means that he will correctly predict whether Q will choose s_1 or s_2.[3] Knowing this, Q will choose her strategy that yields the greater of b_{11} and b_{22}.

More generally, assume that P predicts Q's choice correctly with probability $p >$ ½ (and incorrectly with probability $1–p <$ ½). Q knows not only her own payoffs and p but also that P will choose his strategy in accordance with his prediction. Therefore, if Q chooses s_1, P will

- correctly predict her choice with probability p and choose t_1, giving her a payoff of a_{11}; or
- incorrectly predict her choice with probability $1–p$ and choose t_2, giving her a

[2] In Newcomb's problem, there may be a conflict between maximizing one's expected payoff and choosing a dominant strategy, which is analyzed in Brams (1975) when Newcomb's problem is treated as a game. The Class II prediction games discussed in section 4 pose this problem, wherein we assume Q's choice is determined by her expected-payoff calculation.

[3] Perfect predictions imply that P has a perfectly informed spy, or is omniscient, which can lead to a "paradox of omniscience" (Brams, 1994, 2011), whereby P—in terms of the players' comparative rankings of outcomes—does worse than Q. This paradox occurs in six of the 2 × 2 strict ordinal games, all with two pure-strategy equilibria (including Chicken, which we discuss later). None of these games coincides with those analyzed here, wherein perfect predictions induce cooperative outcomes that P never ranks lower than Q.

payoff of a_{12}.

Thus, Q's expected payoff from choosing s_1 is $EQ(s_1) = a_{11}p + a_{12}(1–p)$. Similarly, if Q chooses s_2, Q's expected payoff from choosing s_2 is $EQ(s_2) = a_{22}p + a_{21}(1–p)$.

It follows that Q prefers s_1 to s_2 (or is indifferent) when $EQ(s_1) \geq EQ(s_1)$:

$$a_{11}p + a_{12}(1–p) \geq a_{22}p + a_{21}(1–p)$$

which is equivalent to

$$p[(a_{11}–a_{22}) + (a_{21}–a_{12})] \geq (a_{21}–a_{12}).$$

Define differences $\Delta_1 = a_{21} - a_{12}$ and $\Delta_2 = a_{11} - a_{22}$. Then we have that Q prefers s_1 to s_2 (or is indifferent), which we write as $s_1 \succcurlyeq s_2$, if and only if

$$p(\Delta_1 + \Delta_2) \geq \Delta_1 . \tag{1}$$

A parallel calculation shows that $s_2 \geq s_1$ if and only if the direction of inequality (1) is reversed.

Because, as we assumed earlier, the players' payoffs are all different, Δ_1 and Δ_2 are never zero. The four cases, given in Table 1 below, show how these values determine whether Q prefers s_1 or s_2 (or is indifferent). In cases 2 and 3, Q's best choice is simple: Always select s_2 in case 2 and s_1 in case 3, whatever the value of p. In cases 1 and 4, "p is large enough" and "p is small enough," depend on whether p is larger than or smaller than the threshold value,

$$p_t = \frac{\Delta_1}{\Delta_1 + \Delta_2} .$$

Note that p_t is well-defined in cases 1 and 4, because $\Delta_1 + \Delta_2$ cannot equal zero since both summands the same sign. Also, $0 < p_t < 1$.

Table 1. Four Cases for Determining *Q*'s Optimal Choice of a Strategy

Case	Δ_1	Δ_2	**Is $s_1 \succcurlyeq s_2$?**
1	> 0 (positive)	> 0 (positive)	$s_1 \succcurlyeq s_2$ if p is *large* enough
2	> 0 (positive)	< 0 (negative)	It is *never* true that $s_1 \succcurlyeq s_2$
3	< 0 (negative)	> 0 (positive)	It is *always* true that $s_1 \succcurlyeq s_2$
4	< 0 (negative)	< 0 (negative)	$s_1 \succcurlyeq s_2$ if p is *small* enough

To clarify the play of a 2×2 prediction game between *P* and *Q*, we present the following timeline:

1. *P* makes a prediction of whether *Q* will choose s_1 or s_2.

2. If *P* predicts that *Q* will choose s_1, he chooses t_1; if *P* predicts that *Q* will choose s_2, he chooses t_2.

3. *Q* chooses s_1 or s_2 optimally, knowing that *P*'s prediction of her strategy is correct with probability p and incorrect with probability $1-p$:

- in case 1, *Q* will choose s_1 if $p > p_t$, be indifferent between s_1 and s_2 if $p = p_t$, and choose s_2 if $p < p_t$.
- in case 4, *Q* will choose s_1 if $p < p_t$, be indifferent between s_1 and s_2 if $p = p_t$, and choose s_2 if $p > p_t$.
- in cases 2 and 3, *Q*'s optimal choice does not depend on p; she will always choose s_2 in case 2 and s_1 in case 3.

We emphasize that P's choice in step 2, and Q's choice is step 3, are made independently of each other. While Q knows p, she does not know P's actual choice so bases her choice on an expected-payoff calculation.

We do not attempt to explain the basis of P's prediction, but presumably it depends on intelligence and other information that P is able to obtain about Q's probable behavior. In particular, P's prediction does not depend on Q's payoffs in the 2×2 game, which may be known only to Q.

In sum, we assume that Q has no information on P's likely behavior, but what she does know is the accuracy of P's prediction about her behavior (i.e., p) and that P will act according to his prediction. Knowing her own payoffs, Q is able to calculate p_t and ascertain, in cases 1 and 4, whether she is better off choosing s_1 or s_2 (or is indifferent)—based on her expected payoff—whereas in cases 2 and 3, as noted earlier, her optimal choice is independent of p.

We can, in fact, generalize the foregoing reasoning to any 2-person game (not just 2×2 games) to show how Q can select one pure-strategy Nash equilibrium out of many when predictions by P are perfect (i.e., $p = 1$). In any such game with two or more pure-strategy Nash equilibria, Q should evaluate each of her strategies in the following way: If no Nash equilibria coincide with strategy s, ignore this strategy; if only one Nash equilibrium coincides with strategy s, then evaluate s as Q's payoff at the equilibrium; if more than one Nash equilibrium coincides with strategy s, then evaluate s as Q's payoff at the equilibrium that P most prefers. Then Q should select her equilibrium strategy with the highest evaluation. Of course, because of continuity, this procedure will be optimal for values of p close enough to 1.

3. Distinguishing the Predictor and the Predictee in 2 × 2 Games

Because P does not have a dominant strategy, as we assumed in section 2, he would like to predict Q's strategy choice as accurately as possible. The more accurate P's prediction is, the more likely he will obtain the outcome he prefers between the two realized by his best responses to each of Q's strategies, s_1 or s_2.

Because there are infinitely many 2 × 2 games, it is useful to classify them according to each player's ranking of the four cardinal outcomes, where

4 = best; 3 = next best; 2 = next worst; 1 = worst.

Any 2 × 2 game in which neither player is indifferent between any two outcomes can be transformed into a strict ordinal game in this way.

If games are identical except for the naming of strategies (i.e., when the two players, A and B, are not distinguished*)*, there are 144 distinct 2 × 2 strict ordinal games. Of these, 12 are symmetric: They present the players with the same strategic choices.

In the remaining 132 games, the strategic choice problems of the row and column players are fundamentally different. More specifically, there are 12 + (½ × 132) = 78 distinct strict ordinal 2 × 2 games that do not distinguish between players A and B: None of these games can be transformed into any other by any combination of the operations of interchanging the players, interchanging the rows, or interchanging the columns (Rapoport, Gordon, and Guyer, 1976). Of the 78 games, 21 have a mutually best (4,4) outcome; in the remaining 57 games, which we call *conflict games*, there is no such outcome.

In Figure 2*a* we give one of these games, which is game 46 in Brams (1994, 2011). It can be transformed into the game in Figure 2*b*, which we call game 46*t* ("*t*" for transformed) by reversing the positions of *A* and *B* (making *A* the column player and *B* the row player, which can be achieved by switching the two off-diagonal outcomes and interchanging the players' payoffs in all four of the resulting entries. Thereby the choices of *A* and *B* in Figure 2a become the choices of *B* and *A* in Figure 2b.

Figure 2. Prediction Game 46 (2*a*) and Its Transformed Counterpart (2*b*)

Figure 2*a*

	B	
A	(3,4)	(2,1)
	(4,2)	(1,3)

Figure 2*b*

	A	
B	(4,3)	(2.4)
	(1,2)	(3,1)

In both games, we assume that the column player is *P* and the row player is *Q*. Although these games are strategically equivalent with respect to the aforementioned transformation, they give very different results when we distinguish the Predictor and the Predictee.

To illustrate this difference, assume the ordinal ranks are cardinal utilities.[4] In the Figure 2*a* game, $\Delta_1 = 4 - 2 = 2$ and $\Delta_2 = 3 - 1 = 2$, which fits case 1 and yields $p_t = ½$, so any $p > ½$ would induce *A* as *Q* to choose s_1. But in the Figure 2*b* game, $\Delta_1 = 1 - 2 = -1$ and $\Delta_2 = 4 - 3 = 1$, which fits case 3 and would induce *B* (now *Q*) always to choose s_1 (p_t is defined and used only in cases 1 and 4 in Table 1).

[4] We treat the ranks as utilities only for illustrative purposes. When the players' utilities are different from the ranks but consistent with them (i.e., the players' utilities can ordered from best to worst in the same order as the ranks), the value of p_t will be different, as we illustrate below.

In the Figure 2*a* game, *Q*'s (*A*'s) optimal choice depends on whether p is greater or less than the threshold probability p_t, whereas in the Figure 2*b* game *Q*'s (*B*'s) optimal choice does not depend on p—she will always choose s_2. In the Figure 2*a* game, if *P*'s predictions are better than random (i.e., $p > ½$), then *Q* will choose s_1.

If *Q*'s cardinal values in the Figure 2*a* game associated with ranks (4, 3, 2, 1) were (9, 3, 2, 1), then $\Delta_1 = 9 – 2 = 7$ and $\Delta_2 = 3 – 1 = 2$, yielding $p_t = 7/9 = 0.78$. Thus, only if p exceeds 78% would *Q* be induced to choose s_1 in the Figure 2*a* game. But this is not the case in the Figure 2*b* game, even when the utilities are (9, 8, 7, 1), because the game of Figure 2*b* is a case 1 game.

When we distinguish the roles of *P* and *Q*, we render the games in Figure 2*a* and 2*b* different in our prediction model. We always arrange for *P* to be the column player, as stated earlier, which accounts for the difference between the Figure 2*a* and Figure 2*b* games.

In a symmetric 2 × 2 ordinal game, each player faces the same strategic choices. In these games, the strategies can be ordered so that the players rank the two main-diagonal outcomes the same, and the rankings of the two off-diagonal outcomes become mirror images of each other: If the ranks of one are (a, b), with $a \neq b$, then the ranks of the other are (b, a). (Prisoners' Dilemma and Chicken are examples of symmetric but not prediction games, which we refer to later.) Because *P* and *Q* assume the same role in a symmetric game— whoever is the row or column player—these games do not give rise to a different prediction games when the row and column players are interchanged.

Of the 144 distinct 2 × 2 ordinal games when *P* and *Q* are distinguished, 36 have a mutually best (4,4) outcome, which we call *no-conflict* games. Thus, there are 108 *conflict games*, in which there is no mutually best outcome, when *P* and *Q* are distinguished. Of the 108 conflict games, 6 are symmetric and 102 are not, which implies that there are 6 + (½ × 102) = 57 different conflict games—a figure we obtained in a different way earlier—prior to the assignment of roles to *P* and *Q*.

Of the 57 conflict games, we focus in section 4 on the eight games (14%) in which there is one Pareto-inferior or no pure-strategy Nash equilibrium and a Pareto-optimal non-Nash outcome that is at least next-best for both players.[5] In all these games, we show that if *P*'s prediction probability *p* is sufficiently high, the players are induced to behave so as to attain the Pareto-optimal non-Nash outcome with this probability, which benefits both players compared with the Nash equilibrium outcome.

4. When Good Predictions Induce a Non-Nash Cooperative Outcome

We begin with the five Class I games, shown in Figure 3, in which there is no Nash equilibrium in pure strategies. Games 46 and 47 are as given in Brams (1994, 2011), whereas games 29*t*, 30*t*, and 31*t* are the transformed counterparts of the games with the same numbers. However, the roles of *P* and *Q* are interchanged and the players' strategies are reversed in some cases.[6] These five games are the 2 × 2 strict ordinal games that satisfy the following conditions:

- There is no mutually best (4,4) outcome.
- *P* has no dominant strategy.
- There are no Nash equilibria (in pure strategies).

---

[5] There is one Pareto-inferior Nash equilibrium in Prisoners' Dilemma (game 32), but it is not a prediction game because *P* (as well as *Q*) has a dominant strategy. Recall that in a prediction game, *P* bases his strategy choice on his prediction of the strategy that *Q* will choose precisely because the *P* does *not* have a dominant strategy.

[6] This reversal may be required to make the Pareto-optimal non-Nash outcome the upper right entry of the payoff matrix in the transformed game, as illustrated in Figure 3 for Class I prediction games.

- The (s_1, t_1) outcome is Pareto-optimal
- The (s_1, t_1) outcome is cooperative, offering both players at least a next-best outcome (i.e., $a_{11} \geq 3$ and $b_{11} \geq 3$).
- $\Delta_1 > 0$ and $\Delta_2 > 0$.

Class I games fall into Case 1 in Table 1: A sufficiently high prediction probability stabilizes the cooperative (s_1, t_1) outcome, even though it is not a Nash equilibrium (there are none) and even though neither s_1 nor t_1 is a dominant strategy.

Figure 3. Five Class I Games (No Pure-Strategy Nash Equilibrium)

Game 46 (p_t = ½)

(3,4)	(2,1)
(4,2)	(1,3)

Game 47 (p_t = ½)

(3,3)	(2,1)
(4,2)	(1,4)

Game 29*t* (p_t = ½)

(3,4)	(2,3)
(4,1)	(1,2)

Game 30*t* (p_t = ½)

(3,3)	(2,2)
(4,1)	(1.4)

Game 31*t* (p_t = ½)

(3.4)	(2,2)
(4,1)	(1,3)

If *Q*'s ranks are cardinal utilities, then *Q* will choose s_1 if *p* exceeds the prediction threshold, which is p_t = ½ in all five games. *P* will choose t_1 with probability *p*, which will lead to the cooperative outcome (a_{11}, b_{11}) with this probability, and to (a_{12}, b_{12}) with probability 1–*p*. As *p* approaches 1, the cooperative outcome will be achieved with ever greater certainty, even though it is not a pure-strategy Nash equilibrium.

More generally, if $p < 1$, the imperfect predictions of P imply that he will not always choose t_1, but sometimes t_2, when Q chooses s_1. Consequently, the outcome if Q chooses s_1 will be the following mix of payoffs:

$$[p \times (a_{11}, b_{11})] + [(1\text{–}p) \times (a_{12}, b_{12})] = [a_{12} + p(a_{11} - a_{12}), b_{12} + p(b_{11} - b_{12})].$$

If $p = 1$, as noted above, the cooperative outcome (a_{11}, b_{11}) occurs with certainty.

We turn next to the three Class II games, which are the 2×2 strict ordinal games that satisfy the following conditions:

- There is no mutually best (4,4) outcome.
- P has no dominant strategy.
- Q has a dominant strategy that includes a Pareto-inferior pure-strategy Nash equilibrium.
- The (s_1, t_1) outcome is cooperative, Pareto-optimal, and preferred by Q to the (s_2, t_2) outcome (which is the Pareto-inferior Nash equilibrium).
- $\Delta_1 > 0$ and $\Delta_2 > 0$.

In Figure 4, the Nash equilibrium is starred, and the non-Nash outcome that sufficiently accurate predictions can stabilize is underscored.[7]

Figure 4. Three Class II Games (One Pareto-Inferior Nash Equilibrium)

[7] Because games 27, 28, and 48 are not symmetric, they have counterparts in which the roles of P and Q are interchanged. However, in these counterpart games, the Predictor has a dominant strategy, so they are not prediction games. In these games, P will presumably choose his dominant strategy in lieu of trying to predict the Q's choice, even though it can be expected to lead to a Pareto-inferior Nash-equilibrium outcome. The same problem bedevils Prisoners' Dilemma (game 32, which is not shown), a symmetric game in which both players have dominant strategies that yield a Pareto-inferior Nash equilibrium outcome. While good predictions in our model do not offer an escape from games in which the Predictor has a dominant strategy, other game-theoretic models with prediction probabilities do; see, for example, the model of arms races in Brams, Davis, and Straffin (1979).

Game 27 (p_t = ¾)

(3,4)	(1,2)
(4,1)	(2,3)*

Game 28 (p_t = ¾)

(3,4)	(1,3)
(4,1)	(2,2)*

Game 48 (p_t = ¾)

(3,4)	(1,1)
(4,2)	(2,3)*

What is noteworthy in the Class II prediction games is that *Q* is best served by abandoning her dominant strategy, s_2, and instead choosing her dominated strategy, s_1, if *p* is sufficiently high. If the ranks are assumed to be cardinal utilities, the threshold probability p_t is ¾ in each of these games. The outcome that results with probability *p*, (3,4), favors *P*, giving him his best outcome and *Q* her next-best outcome.

Altogether, there are a total of eight prediction games in Figures 3 and 4. They constitute 14% of the 57 conflict games and include *all* games in which there is (i) either no Nash equilibrium in pure strategies or a Pareto-inferior Nash equilibrium and (ii) a non-Nash Pareto-optimal cooperative outcome. All these games yield the cooperative outcome with probability *p*, and the adjacent outcome in the same row (s_1, t_2) with probability (1–*p*), if $p > p_t$. At the non-Nash cooperative outcomes in these games, *P* does at least as well and sometimes better—in terms of comparative rankings of payoffs—as *Q*.

5. The Cuban Missile Crisis: A Class I Game

There is a prodigious literature on the Cuban missile crisis, which several analysts have modeled as a game of Chicken in which the two superpowers are seen to be on a collision course.[8] Brams (1994, 2011) has disputed the Chicken interpretation and proposed an alternative game, which is a comparison we will discuss shortly.

[8] Supporting this interpretation is Secretary of State Dean Rusk's statement, "We're eyeball to eyeball, and I think the other fellow just blinked"—spoken at the climatic moment of the crisis and reported several

First, however, some background. Before the breakup of the Soviet Union in 1991 and its demise as a superpower, the Cuban missile crisis was surely the most dangerous confrontation between the superpowers ever to occur (known as the Caribbean crisis in the Soviet Union). It was precipitated by a Soviet attempt in October 1962 to install medium-range and intermediate-range nuclear-armed ballistic missiles in Cuba capable of hitting a large portion of the United States. How the superpowers—or, more accurately, their leaders—managed this crisis has been described in great detail, but this scrutiny has involved little in the way of formal strategic analysis.[9]

After the presence of such missiles was confirmed on October 14, the Central Intelligence Agency estimated that they would be operational in about ten days. A so-called Executive Committee (ExCom) of high-level officials was convened to decide on a course of action for the United States. ExCom met in secret for six days. The comments of Theodore Sorensen, special counsel and advisor to President John F. Kennedy, on its deliberations reflect well the game-theoretic thinking of its members:

> We discussed what the Soviet reaction would be to any possible move by the United States, what our reaction with them would have to be to that Soviet reaction, and so on, trying to follow each of these roads to their ultimate conclusion (quoted in Holsti, 1964, p. 188).

The goal of the United States was immediate removal of the Soviet missiles, and U.S. policy makers considered several strategies, eventually narrowed to two:

---

weeks later. *Collision Course* is actually the title of the first book on this crisis (Pachter, 1963), and *Eyeball to Eyeball* it the title of another book on the crisis (Brugoni, 1992),

[9] Exceptions include the game-theoretic models of Snyder and Diesing (1977), Fraser and Hipel (1982-83), Brams (1985, 1990), and Dixit, Skeath, and Reilly (2015).

1. **A naval blockade (*B*)**, or "quarantine" as it was euphemistically called, to prevent shipment of further missiles, possibly followed by stronger action to induce the Soviet Union to withdraw those missiles already installed.
2. **A "surgical" air strike (*A*)** to wipe out the missiles already installed, insofar as possible, perhaps followed by an invasion of the island.

The alternatives open to Soviet policy makers were:

1. **Withdrawal (*W*)** of the missiles already in Cuba.
2. **Maintenance (*M*)** of the missiles already installed.

The game of Chicken (game 57), would, at first blush, seem an appropriate model of this conflict.[10] As applied to the Cuban missile crisis, the strategies and ranking of outcomes by the United States and the Soviet Union are shown in Figure 5.[11]

Figure 5. Cuban Missile Crisis (Game 57: Chicken)

		United States	
		Blockade: *B*	Air strike: *A*
Soviet Union	Withdraw: *W*	Compromise (3,3)	U.S. victory, Soviet defeat (2,4)
	Maintain: *M*	Soviet victory, U.S. defeat (4,2)	Nuclear war (1,1)

[10] The description of Chicken, and the alternative game we propose to model the crisis, are adapted from Brams (1994, 2011), but our explanation for the outcome of the crisis is based on our prediction model, not on the "theory of moves" that is developed and applied in these books.
[11] The superpowers are modeled as unitary actors. This simplification has been rectified in part by constructing alternative models that emphasize different features of crisis decision-making (e.g., in Allison, 1971).

Of course, there is no way to verify that the strategy choices and probable outcomes given in Figure 5 were the most likely ones, or that the Chicken rankings are the most appropriate. For example, if the Soviet Union had viewed an air strike on its missiles as jeopardizing vital national interests, the (4,2) outcome may very well have ended in nuclear war between the two sides, giving it the same value as (1,1). Still another simplification relates to the assumption that the players chose their actions simultaneously, when in fact a continual exchange of messages, occasionally backed up by actions, occurred over those fateful days of October 1962.[12]

Most observers of this crisis agree that neither side was eager to take any irreversible step, such as a driver in Chicken might do by defiantly ripping off his or her steering wheel in full view of the other driver, thereby foreclosing the option of swerving and communicating this to the other driver. While the United States "won," in a sense, by getting the Soviets to withdraw their missiles, Premier Nikita Khrushchev, the Soviet leader, extracted from President Kennedy a promise not to invade Cuba in the future, which seems to indicate that the eventual outcome was a compromise resolution of sorts.[13] Moreover, even though the Soviets responded specifically to the blockade that was imposed and, therefore, did not make the choice of their strategy independently of the American strategy choice, the fact that the United States held out the possibility of escalating the conflict to at least an air strike indicates that the initial blockade decision was not considered final—that is, the United States considered other strategic options still open after the blockade was imposed.

---

[12] For several books on the crisis not cited here, which use different sources and emphasize different aspects of the crisis, see Brams (2011, p. 228, note 13).

[13] A release of letters between Kennedy and Khrushchev, however, indicates that Kennedy's promise was conditional on the good behavior of Cuba (Pear, 1992).

Truly, this was a game of sequential bargaining, in which each side did not make an all-or-nothing choice but instead considered alternatives should the other side fail to respond in a manner considered appropriate. Facing the most serious breakdown in the deterrence relationship between the superpowers that had persisted since World War II, each side was carefully feeling its way, step by ominous step.

Before the crisis, the Soviets, fearing an invasion of Cuba by the United States and also the need to bolster their international strategic position, concluded that it was worth the risk of installing the missiles; they believed that, confronted by a *fait accompli*, the United States, in all likelihood, would be deterred from invading Cuba and would not attempt any other severe reprisals (Garthoff, 1989). The Soviet calculation suggests that it thought there was a relatively low probability of nuclear war,[14] thereby making it rational to risk provoking the United States.

Although this thinking may be more or less correct, there are good reasons to believe that U.S. policy makers viewed the game not to be Chicken at all, at least as far as they ranked the possible outcomes. In Figure 6, we offer an alternative representation of the Cuban missile crisis, retaining the same strategies for both players as given in the Chicken representation (Figure 5), but we assume a different ranking of outcomes by the United States.

[14] This probability seems to have been higher than initially thought, according to some disclosures. They indicate that the Soviets had four times as many troops in Cuba as U.S. intelligence estimated as well as tactical nuclear weapons that they would have used in the event of a U.S. invasion (Tolchin, 1992). At the height of the crisis. President Kennedy estimated the chances of war to be between one-third and one-half (Sorensen, 1965, p. 705). For a theoretical analysis of the probability of nuclear war, see Avenhaus, Brams, Fichtner, and Kilgour (1989).

The resulting game is game 30*t* which, unlike Chicken, is a Class I prediction game.[15] Its outcomes may be interpreted as follows:

***BW*: (3,3).** The choice of blockade by the United States and withdrawal by the Soviet Union remains the cooperative outcome for both players.

***BM*: (1,4).** In the face of a U.S. blockade, Soviet maintenance of their missiles leads to a Soviet victory (its best outcome) and U.S. capitulation (its worst outcome).

***AM*: (4,1).** An air strike that destroys the missiles that the Soviets were maintaining is an "honorable" U.S. action (its best outcome) and thwarts the Soviets (their worst outcome).

***AW*: (2,2).** An air strike that destroys the missiles that the Soviets were withdrawing is a "dishonorable" U.S. action (its next-worst outcome) and thwarts the Soviets (their next-worst outcome).

Figure 6. Cuban Missile Crisis (Game 30*t*)

		United States	
		Blockade: *B*	Air strike: *A*
Soviet Union	Withdraw: *W*	Compromise (3,3)	"Dishonorable" U.S. action, Soviets thwarted (2,2)
	Maintain: *M*	Soviet victory, U.S. capitulation (4,1)	"Honorable" U.S. action, Soviets thwarted (1,4)

[15] Chicken is neither a Class I nor a Class II prediction game. Although the United States, as *P*, does not have a dominant strategy, its preferred outcomes for each choice of the Soviet Union do not lie along the main diagonal: When the Soviets choose *W*, the United States prefers *B*, whereas the opposite is true in game 30*t*.

Even though an air strike thwarts the Soviets in the case of the outcomes with payoffs of (2,2) and (4,1), we interpret (2,2) to be a less damaging outcome for the Soviet Union. This is because world opinion, we surmise, would severely condemn the air strike as a flagrant overreaction—and hence a "dishonorable" U.S. choice—if there were clear evidence that the Soviets were in the process of withdrawing their missiles anyway. On the other hand, given no such evidence, a U.S. air strike, perhaps followed by an invasion, would probably be viewed by U.S. policy makers as a necessary, if not an "honorable," action to dislodge the Soviet missiles.

Before analyzing these possibilities, however, we offer a brief justification—mainly in the words of the participants—for the alternative representation given by game 30*t*. The principal protagonists, of course, were President Kennedy and Premier Khrushchev. Their public and private communications over the thirteen days of the crisis indicate that both understood the dire consequences of precipitous action and shared, in general terms, a common interest in preventing nuclear war. For the purpose of the present analysis, however, what is relevant are their specific preferences for each outcome.

Did the United States prefer an air strike (and possible invasion) to the blockade, given that the Soviets would withdraw their missiles? In responding to a letter from Khrushchev, Kennedy said:

> If you would agree to remove these weapons systems from Cuba . . . we, on our part, would agree . . . (a) to remove promptly the quarantine measures now in effect and (b) to give assurances against an invasion of Cuba (Allison, 1971, p. 228).

This statement is consistent with the game 30*t* representation of the crisis [because (3,3) is preferred to (2,2) by the United States] but not consistent with the Chicken (game 57) representation [because (4,2) is preferred to (3,3) by the United States].

Did the United States prefer an air strike to the blockade, given that the Soviets would maintain their missiles? According to Robert Kennedy, a close adviser to his brother during the crisis, “If they did not remove those bases, we would remove them” (Kennedy, 1969, p. 170). This statement is consistent with the game 30*t* representation [because (4,1) is preferred to (1,4) by the United States] but not with the Chicken representation [because (2,4) is preferred to (1,1) by the United States].

Finally, it is well known that several of President Kennedy’s advisers felt very reluctant about initiating an attack against Cuba without exhausting less belligerent courses of action that might bring about the removal of the missiles with less risk and greater sensitivity to American ideals and values. As Robert Kennedy put it, an immediate attack would be looked upon as “a Pearl Harbor in reverse, and it would blacken the name of the United States in the pages of history” (Sorensen, 1965, p. 684). This statement is consistent with the United States’s ranking *AW* next worst (2)—a “dishonorable” U.S. action in the game 30*t* representation—rather than best (4)—a U.S. victory in the Chicken representation.

If game 30*t* (Figure 6) provides a more realistic representation of the participants’ perceptions than does Chicken (Figure 5), game theory offers little in the way of explaining how the cooperative (3,3) outcome was achieved and rendered stable. After all, this outcome, as in Chicken, is not Nash equilibrium outcome and, as noted earlier, no other pure-strategy outcome is.

How then can we explain the choice of (3,3) by both sides? As a prediction game, assume the United States is *P* and attempts to predict whether the Soviet Union will withdraw or maintain its missiles after the blockade.

We believe President Kennedy chose blockade after predicting, or at least hoping, that it would be sufficient to induce the Soviet Union to withdraw its missiles: It would force Khrushchev to choose between the serious embarrassment of withdrawal and the dire consequence of war if he did not do so, including a possible nuclear exchange.

Obviously, the Soviets would prefer (3,3) in game 30*t* to being thwarted at (4,1). As for Kennedy's choice, it was probably driven by his desire to show his resolve after the April 1961 Bay of Pigs debacle in Cuba and his June 1961 meeting with Khrushchev in Vienna, in which he believed he had appeared weak and indecisive.[16]

At the Vienna meeting, a crisis over Berlin was precipitated when Khrushchev demanded withdrawal of Western forces from West Berlin. Two months later, when construction of the Berlin Wall commenced in August, there was no forceful U.S. reaction, which reinforced Khrushchev's belief that Kennedy was not a formidable opponent.

In fact, Khrushchev believed that if the United States protested but took no countermeasures after the deployment of missiles in Cuba, he could make their withdrawal contingent on ceding all of Berlin to East Germany, which he considered strategically far more important than Cuba in the Cold War. Khrushchev said he wished to confront the Americans "with more than words … the logical answer was missiles" (Weldes, 1999, p. 29).

When Kennedy demanded withdrawal of Soviet missiles from Cuba, a conventional war in the region was not a viable option for the Soviets. For one thing, the United States had overwhelming military superiority in the Caribbean.[17] For another, if there were escalation of the conflict to the nuclear level, the Soviet Union had only a few intercontinental ballistic missiles it could launch against the United States, and it was at a severe disadvantage in nuclear-armed submarines.

---

[16] Kennedy said about the meeting, "He beat the hell out of me," and he told *New York Times* reporter James Reston, "It was the worst thing in my life. He savaged me." Khrushchev told his son, Sergei, that if missiles were installed in Cuba, Kennedy "would make a fuss, make more of a fuss, and then agree" to their deployment (Kempe, 2011, p. 257).

[17] Of course, the Soviet Union could have fomented trouble elsewhere, as it did by blockading Western ground access to its sector of Berlin, beginning in June 1948 and lasting almost a year. But a massive airlift of supplies to West Berlin, involving over 200,000 flights, induced Josef Stalin at that time to lift the blockade in May 1949.

Consequently, it seemed predictable that the Soviets would choose to withdraw their missiles while trying to wring concessions from the United States. This included, as noted earlier, a promise by the United States not to invade Cuba as well as a secret agreement to withdraw its missiles from Turkey.

As terrifying as the Cuban missile crisis was, it led to the establishment of a hot line between the superpowers—and other stabilizing measures—which played a significant role in alleviating subsequent international crises. By correctly predicting that Khrushchev would back down from his belligerent stance at the outset of the crisis, Kennedy was able not only to defuse it but set in motion later steps that clearly had favorable long-term consequences.

6. The Iran Agreement on Nuclear Weapons: A Class II Game

In 2012, fifty years after the Cuban missile crisis, several countries and the International Atomic Energy Agency feared that Iran might be attempting to develop a nuclear capability that could be used for military purposes. Israel, in particular, believed that Iran was enriching uranium in order to develop nuclear weapons that could be used against it. It had suspected such surreptitious activities earlier, but it claimed in 2012 that they posed an imminent threat to its existence.

Iran denied that developing nuclear weapons was its intention, despite the discovery of previously hidden nuclear-production facilities. It said that it desired to enrich uranium only as an alternative energy source to be used for civilian purposes.

Israeli Prime Minister Benjamin Netanyahu threatened to attack Iran and destroy its nuclear capability unless there was proof, based on the rigorous inspection of its suspected nuclear facilities, that Iran was not developing nuclear weapons. (A number of Israeli leaders opposed such an attack, arguing that at best it might delay but would not stop Iran's acquisition of nuclear weapons.[18]) Israel and Iran were at an impasse when Iran denied international inspectors access to the facilities in question.

Because of its refusal, Iran suffered ever more severe economic sanctions imposed by the United States, the European Union, and other countries. But a carrot was held out, with the sanctioners offering to relax or lift the sanctions if Iran agreed to allow inspections and to credibly commit to halting any efforts that could lead to the production of nuclear weapons. However, a number of countries, including China and Russia, opposed the use of sanctions.

The most immediate danger of armed conflict arose from Israel's threat to attack Iran's nuclear-production facilities. More specifically, Israel's position was that, failing an agreement, it would attack Iran's facilities before a point of no return—called a "zone of immunity" by Israeli Defense Minister Ehud Barak—was reached. The point that would trigger an attack, Israel said, would be the time just before these facilities became sufficiently hardened (they were inside a mountain) to be effectively impregnable.

Whether the United States would actively participate in such an attack, or covertly facilitate it, was unclear. On March 8, 2012, President Barack Obama said the United States "will always have Israel's back," which signaled that he was supportive of Israel's concern but did not spell out exactly what the United States would do to aid Israel.

To go back a few decades to the1970s, Israel began producing, but never publicly acknowledged possessing, nuclear weapons, though it is now presumed to have about eighty nuclear warheads (Arms Control Association, 2016). It did say, however, that it would not be the first party to introduce them into a conflict.

---

[18] Prominent among these critics was Meir Dagan, the director of Mossad until 2011 (Sanger, 2016).

In claiming that Iran's acquisition of nuclear weapons threatened its existence, Israel implied that it would use every means short of nuclear weapons to arrest Iran's development of them if economic sanctions or covert actions failed. The latter had included assassinations and cyberwarfare. Indeed, the latter had significantly disrupted Iran's enrichment of uranium.

Unlike the superpowers during the Cold War, Israel was unwilling to rely on its own nuclear deterrent and MAD, perhaps in part because it feared that terrorists could gain control of Iranian nuclear weapons and act "crazily," without concern for what Israel's response might be. Also, Israel's small physical size made its survival an issue, even after it retaliated from an attack, whereas Iran's ability to absorb a retaliatory strike was greater, possibly giving it an incentive to preempt with nuclear weapons.

We present in Figure 7 a Class II game (game 27 in Figure 4) to model the conflict between Iran and Israel. In this game, Iran chooses between developing (D) or not developing ($\overline{D}$) nuclear weapons, and Israel chooses between attacking (A) or not attacking ($\overline{A}$) Iran's nuclear facilities.

Figure 7. Iran-Israel Conflict (Game 27)

		Israel	
		Don't Attack: $\overline{A}$	Attack: A
Iran	Don't Develop: $\overline{D}$	Israel's threat effective (3,4)	Unprovoked attack (1,2)
	Develop: D	Israel's threat empty (4,1)	Justified attack (2,3)

We assume Israel's ranking to be $\overline{D}\ \overline{A} > DA > \overline{D}A > D\overline{A}$. As justification, there is little doubt that Israel would most prefer a cooperative solution ($\overline{D}\ \overline{A}$), in which Iran does not develop nuclear weapons and Israel does not attack them. Israel could claim that its threat was effective, so no attack was required. Israel would least prefer that Iran develop nuclear weapons without Israel's making an effort to stop their production ($D\overline{A}$), proving Israel's threat was empty.

Between attacking weapons that are being developed (*DA*) and mistakenly attacking weapons that are not being developed ($\overline{D}A$), we assume that Israel would prefer the former: An unprovoked attack on weapons being developed would certainly create a major crisis, but it would be seen by Israel as a measure that was essential to its survival.[19]

As for Iran, we assume that its most preferred outcome is to develop nuclear weapons without being attacked ($D\overline{A}$), and its least preferred outcome is not to develop nuclear weapons and be attacked anyway ($\overline{D}A$). In between, we assume that Iran prefers the cooperative outcome ($\overline{D}\ \overline{A}$) to the noncooperative outcome (*DA*), which could lead to a major conflict and possibly war after the attack.

D is a dominant strategy for Iran, and the unique Nash equilibrium in game 27 is the noncooperative outcome (*DA*). Unfortunately for both countries, this outcome, (2,3), is Pareto-inferior to (3,4), but the strategy pair $\overline{D}\ \overline{A}$ that yields (3,4) is not a Nash equilibrium.

[19] But there is the question of whether Israel's detection of uranium enrichment would constitute a *casus belli* for Israel to attack Iranian nuclear facilities. After all, enrichment does not imply that the enriched uranium will be used for weaponization.

In July 2015, P5+1—the five permanent members of the U.N. Security Council plus Germany and the European Union—reached an agreement with Iran for the robust inspection of its nuclear facilities for fifteen years. These inspections would ensure that any significant enrichment of uranium by Iran, which could lead to the production of nuclear weapons, would be detected, triggering strong countermeasures.

The agreement also included several other provisions to inhibit Iran's development of such weapons. Although the Israeli government, though not several leading Israeli critics, vehemently opposed this agreement, the agreement has so far defused a volatile situation that could have led to an Israeli attack.

Why didn't Israel attack earlier, as it continually threatened to do from 2012 to 2015? We suggest that it probably had good intelligence that Iran was not approaching the zone of immunity and, in addition, that Iran did not have the capability, or even the intention, of producing nuclear weapons. If this was the case, Israel would prefer $\overline{A}$ to A, obtaining (3,4) instead of (1,2) in game 27.

Iran, we presume, knew that Israel, as well as the United States, could closely track its progress in its nuclear program. While not knowing exactly what these countries knew about its activities, Iran was certainly aware that they could predict, with a high probability p, its choice of a strategy.

Knowing p, even if not exactly, Iran's expected payoff from choosing $\overline{D}$—including the benefit of getting at least some economic sanctions lifted[20]—appears to have been greater than its expected payoff from choosing D. Thereby both sides were dissuaded from choosing their strategies associated with the Nash equilibrium outcome of

[20] The 2015 agreement, known as the Joint Comprehensive Plan of Action (JCPOA), provided for the gradual lifting of sanctions if Iran verifiably abided by its commitment not to develop nuclear weapons.

(2,3) that, in the absence of Israel's prediction capability and Iran's knowledge of it, would be the outcome that game theory would prescribe, especially because this outcome is associated with Iran's dominant strategy.

In summary, we have argued that the conflict between Iran and Israel over Iran's possible development of nuclear weapons can plausibly be represented by game 27, in which Israel is P and Iran is Q. Given Israel's probability of correct prediction is sufficiently high, which we assume is known by Iran, and it chooses in accordance with its preferences in game 27, it is in Iran's interest to choose not to develop nuclear weapons for two reasons: (i) its fear of an attack on its production facilities; and (ii) the continued tightening of economic sanctions (this is not in our game as such but seems to have been a major factor in inducing Iran to reach a settlement[21]).

7. Conclusions

We postulated a 2-person game in which one player, Predictor (P), does not have a dominant strategy but can predict with probability $p > 1/2$ the strategy choice of an opponent, Predictee (Q). Q chooses a strategy that maximizes her expected payoff, given that she knows p—but not P's prediction—and that P will act according to his prediction.

In all 2×2 strict ordinal games in which there is a unique Pareto-inferior Nash equilibrium (Class I) or no pure-strategy equilibrium (Class II), and a non-Nash Pareto-optimal cooperative outcome, P can induce the cooperative outcome if his predictions are

[21] Tightening of sanctions could be modeled as a deterioration of Iran's payoff at (4,1) in game 27 to (3,1), and a complementary increase of its payoff at (3,4) to (4,4), yielding a no-conflict game that has two Nash equilibrium outcomes, (4,4) and (2,3). Obviously, the Pareto-superior (4,4) outcome would be more appealing to both sides, but neither equilibrium in this game, as in the symmetric game, Stag Hunt, is supported by a dominant strategy, as D is in game 27. It seems probable that this deterioration for Iran had not yet occurred in 2015, so the prediction model seems the more plausible explanation of why a settlement was reached.

sufficiently accurate. This outcome gives each player at least a next-best payoff. It seems to have occurred in a Class I game that models the 1962 Cuban missile crisis between the United States and the Soviet Union, and in a Class II game that models the 2015 conflict between Iran and Israel over Iran's possible development of nuclear weapons.

We first considered the representation of the Cuban missile crisis as a game of Chicken (game 57) but then suggested that the non-Nash cooperative (3,3) outcome was realized in a different game (game 30*t*), which better models the preferences of the two players. In the latter game, which is a Class I prediction game in which there is no Nash equilibrium in pure strategies, the United States correctly predicted that a blockade would qbe sufficient to induce the Soviet Union to withdraw its missiles from Cuba, though it held out the possibility of escalation to an air strike if the Soviets did not withdraw their missiles. Aware of U.S. intelligence capabilities and making a probable expected-payoff calculation, the Soviet Union acceded to the U.S. demand for withdrawal of its missiles after the United States agreed to certain concessions.

In the Iran-Israel conflict over the Iran's possible development of nuclear weapons, the agreement reached in 2015 seems to have been fostered by Israel's strong intelligence and Iran's awareness of it. It seems that Iran abandoned its dominant strategy in a Class II prediction game (game 27), moving the conflict from a possible Pareto-inferior Nash equilibrium to a Pareto-superior non-Nash cooperative outcome.

The lesson from both our theoretical model and the empirical cases is that good predictions by P, and knowledge of P's capability by Q, can stabilize a non-Nash Pareto-optimal cooperative outcome in eight prediction games. Surprisingly, Q does not need to know *P*'s prediction but only the probability that it is correct for the players to choose, with this probability, the cooperative outcome in these games.

That this salutary result occurs not only in theory but also in practice underscores the critical role that information can play in helping to resolve serious international conflicts. To induce cooperative outcomes in prediction games, it is in the interest of *P*, who does not have a dominant strategy, to be able to predict with reasonable accuracy the choice of *Q*, and for *Q* to understand that *P* has this capability.

References

Allison, Graham T. (1971). *Essence of Decision: Explaining the Cuban Missile Crisis*. Boston: Little, Brown.

Arms Control Association (2016). "Nuclear Weapons: Who Has What at a Glance." https://www.armscontrol.org/factsheets/Nuclearweaponswhohaswhat.

Avenhaus, Rudolf, Steven J. Brams, John Fichtner, and D. Marc Kilgour (1989). "The Probability of Nuclear War." *Journal of Peace Research* 26 (February): 91-99.

Brams, Steven J. (1975). "Newcomb's Problem and Prisoners' Dilemma." *Journal of Conflict Resolution* 19, no. 4 (December): 596-611.

Brams, Steven J. (1994). *Theory of Moves*. Cambridge, UK: Cambridge University Press.

Brams, Steven J. (2011). *Game Theory and the Humanities*. Cambridge, MA: MIT Press.

Brams, Steven J., Morton D. Davis, and Phillip D. Straffin, Jr. (1979). "The Geometry of the Arms Race." *International Studies Quarterly* 23, no. 4 (December): 567-588.

Brugoni, Dina A. (1992). *Eyeball to Eyeball: The Inside Story of the Cuban Missile Crisis*. New York: Random House.

Bruns, Bryan Randolph (2015). "Names for Games: Locating 2 × 2 Games." *Games* 6, no 4: 495-520.

Dixit, Avinash, Susan Skeath, and David Reiley (2015). *Games of Strategy*, 4th ed. New York: W. W. Norton, 2015.

Fraser, Neill M., and Keith W. Hipel (1962-63). "Dynamic Modeling of the Cuban

Missile Crisis." *Conflict Management and Peace Science* 6, no. 2 (Spring): 1-18.

Garthoff, Raymond L. (1989). *Reflections on the Cuban Missile Crisis*, rev. ed. Washington, DC: Brookings.

Kempe, Frederick (20110. *Berlin 1961*. New York: Penguin.

Kennedy, Robert F. (1969). *Thirteen Days: A Memoir of the Cuban Missile Crisis*. New York: W. W. Norton.

Pachter, Henry M. (1963). *Collision Course: The Cuban Crisis and Coexistence*. New York: Praeger.

Pear, Robert (1992). *The Cuba Missile Crisis*. "Kennedy Left a Loophole." *New York Times*, January 7, p. A5.

Rapoport, Anatol, David G. Gordon, and Melvin J. Guyer (1976). *The 2 × 2 Game*. Ann Arbor, MI: University of Michigan Press.

Sanger, David E. (2016). A Spymaster Who Saw Cyberattacks as Israel's Best Weapon Against Iran." *New York Times*, March 22.

Snyder, Glenn H., and Paul Diesing (1977). *Conflict among Nations: Bargaining, Decision Making, and Systems Structure in International Crises*. Princeton, NJ: Princeton University Press.

Sorenson, Theodore C. (1965). *Kennedy*. New York: Harper and Row.

Tolchin, Martin (1992). "U.S. Underestimated Soviet Force in Cuba During '62 Missile Crisis." *New York Times*, February 25, p. A11.

Weldes, Juda (1999). *Constructing National Interests: The United States and the Cuban Missile Crisis*. Minneapolis: University of Minnesota Press.