

ECONOMIC RESEARCH REPORTS

Games That End in a Bang or a Whimper

by

**Steven J. Brams
&
D. Marc Kilgour**

RR#: 2001-05

October 2001



C.V. Starr Center for Applied Economics

Department of Economics

Faculty of Arts and Science

New York University

269 Mercer Street, 3rd Floor

New York, New York 10003-6687

To appear in George F. R. Ellis (ed.), *The Far-Future University: Eschatology from a Cosmic Perspective* (Templeton Foundation Press)

Games That End in a Bang or a Whimper

Steven J. Brams
Department of Politics
New York University
New York, NY 10003
USA

E-mail: steven.brams@nyu.edu

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

E-mail: mkilgour@wlu.ca

October 2001

Abstract. Using truels, or three-person duels, as an example, we show that how players perceive a multiple-round game will end can make a big difference in whether it ends non-cooperatively (producing a “bang”) or just peters out (producing a “whimper”):

1. If the players view the number of rounds as *bounded*—reasonable, because the game must end in a finite number of rounds—they will shoot from the start.
2. If the players view the number of rounds as *unbounded*—reasonable, because the horizon of the game is infinite—then a cooperative equilibrium, involving no shooting, can also occur.

Real-life examples are given of players with bounded and unbounded outlooks in truel-like situations. Unbounded outlooks encourage cooperative play, foster hope, and lead to more auspicious outcomes. These outcomes are facilitated by institutions that put no bounds on play—including reprisals—thereby allowing for a day of reckoning for those who violate established norms. Eschatological implications of the analysis, especially for thinking about the future and how it might end, are also discussed.

JEL Classification: N72

Keywords: Truels; backward induction; infinite-horizon games; eschatology

Games That End in a Bang or a Whimper¹

Our subject is the end of games that real human beings play. Because these games can affect how humanity fares in this and possibly other worlds, they have eschatological significance.

Games may be bounded or unbounded. *Bounded games* are those which end after a certain time, or after a specific number of rounds have been played. In *unbounded games*, there is no such limit or bound.

Is life bounded? Although there seems to be no confirmed case of a person's having lived more than 125 years (the confirmed maximum is 122 years, achieved by a French woman who died in 1997), there is no logical reason or scientific barrier to preclude a person, should he or she reach the age of 125, from living to be 126. Hence, to say that life is bounded by a limit, like 125 years, seems unjustified.

But what about extending that limit to 250 years, 1,000 years, or even 1,000,000 years? It seems absurd that any of us will ever approach such an age. On the other hand, the possibility that our genetic material might somehow be preserved or renewed is not so easy to dismiss. Alternatively, living our lives through our descendants—if the definition of life is broadened to include them—renders “ages” like 250, 1,000, or 1,000,000 years conceivable.

¹ Steven J. Brams acknowledges the support of the C. V. Starr Center for Applied Economics at New York University, and D. Marc Kilgour the support of the Social Sciences and Humanities Research Council of Canada. We thank Geir B. Asheim for valuable comments on an earlier version of this paper. Although this paper is relatively nontechnical, there is some discussion, especially in the footnotes, of more technical concepts in game theory that underlie our somewhat informal analysis.

The probability that any of us will, as individuals, live more than 125 years is at present infinitesimal. Practically speaking, it is reasonable to suppose that our lives are bounded by about 125 years.

Basing our actions on this limit, however, may induce very different rational choices from those that reflect the view that life is unbounded. To illustrate this point, we will describe several games, some of which are bounded and some of which are not. The available choices of the players in these games are all the same, but whether the games are perceived to be bounded critically affects how they are rationally played.

The game we conclude with, an *infinite-horizon* game, is perhaps the most realistic: the horizon is infinite—because any pre-specified bound can be exceeded—but the end occurs before the infinite horizon is reached, guaranteeing that any play of the game is finite.² Interestingly enough, if the players' outlooks are unbounded in this game, they have an incentive to be cooperative, causing the game to peter out in a “whimper”; if their outlook is bounded, they will be non-cooperative, causing the game to end in a “bang” (we mean this quite literally, as the game will illustrate).

There is no incontrovertible argument, nor evidence we know of, that either outlook is correct. People can behave rationally under the presumption that an infinite-horizon game will grind to a halt, which it definitely will because it is finite. But because the same game can go on indefinitely, it is impossible to predict this termination exactly.

² Carse's (1986) distinction between finite and infinite games is very different: the former are essentially short-term zero-sum games, whereas the latter are long-term nonzero-sum games that are invariably beneficial to their players, who can even on occasion change the rules to “improve” the game. Carse's analysis of infinite games, which is untouched by game theory, seems to us wistfully naïve; by comparison, we use game theory to try to show how games that may continue indefinitely offer some realistic hope for improving the lots of their players under certain conditions. For more on the application of game theory to the Bible and religion and theology, see Brams (1980, 1983).

Eschatology postulates an end, but it is often quite vague about when it will occur.

This vagueness, we suggest, can be clarified using game theory:

- If people are “forward thinking” (unbounded), they will look ahead to determine how to behave, based on their expectations about the future;
- If people are “backward thinking” (bounded), they will look back from the presumed end, determining optimal choices to make on last round, the next-to-last round, and so on, until they make their initial choices.

Looking ahead, a person calculates expected values, based on events that can happen and probabilities associated with their occurrence (if known). Looking backward, a person calculates that because a game will end at some point, he or she can determine rational choices before then by tracing out the consequences of actions and reactions along the path to this end in a process called *backward induction*. As Theodore Sorensen put it in describing the deliberations of the Executive Committee (Excom) during the October 1962 Cuban missile crisis, “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 those roads to their ultimate conclusion” (quoted in Holsti, Brody, and North, 1964, p. 188).

During the Cuban missile crisis, many people thought that the world might come to an end in a nuclear exchange between the superpowers, rendering their lives decidedly bounded. But when the crisis subsided after thirteen days, life for most people regained its unbounded character.

To be sure, the end of the universe may be little affected by such human events. Rather, the universe seems more an impersonal entity, even if it contains human beings capable of rational thought and action.

Understanding better how human beings view and play games may, nonetheless, illuminate how the universe, which possibly embodies a kind of rationality even if no single entity directs its behavior, moves toward some endstate. At a more personal level, the eschatological view we hold about the boundedness of games may well affect our rational choices in them. As a preview of the distinction between bounded and unbounded games, we start with a simple hypothetical game, whose rules we will progressively change to produce, in the end, an infinite-horizon game.

A Sequential Truel

Imagine three players, A, B, and C, situated at the corners of an equilateral triangle. They engage in a *truel*, or three-person duel, in which each player has a gun with one bullet.³

Assume that each player is a perfect shot and can fire at one other player at any time.⁴ There is no fixed order of play, but any shooting that occurs is sequential: no player fires at the same time as any other. Consequently, if a bullet is fired, the results are always known to all before another bullet is fired.

Finally, assume that each player ranks the outcomes from best to worst as follows:
(1) survive alone, (2) survive with one other player, (3) survive with both other players,

³ Background on truels, including some with rules quite different from those analyzed here, can be found in Kilgour and Brams (1997).

⁴ We rule out the possibility of firing in the air, which would be an optimal choice for a player if it were the first to fire. For once a player has disarmed itself, it would be no threat to its two opponents, which would then have an incentive to shoot each other in a duel. (Why? Because if the second player to choose also

(4) not survive, with no opponents alive, (5) not survive with one opponent alive, and (6) not survive with both opponents alive. Thus, surviving alone is best, dying alone worst.

Who, if anybody, will shoot whom? It is not difficult to see that outcome 3, in which nobody shoots and, therefore, all three players survive, is the rational outcome. Suppose, on the contrary, that A shoots B, hoping for A's outcome 2, whereby it and C survive. A's best outcome, surviving alone, is now impossible—C will not shoot itself. In fact, C, preferring its outcome 1 to outcome 2, will next shoot a disarmed A, leaving itself as the sole survivor.

But this is A's outcome 5, in which A and one opponent (B) are killed while the other opponent (C) lives. To avoid this outcome, A should not fire the first shot; neither, for the same reason, will the other two players. Consequently, nobody will shoot, resulting in outcome 3, in which all three players survive.

Moreover, it will not pay for any two players—say, A and B—to collude and both shoot C, thereby expending their bullets and posing no threat to each other. For if they agree to collude, it would be in each of A's and B's interests to renege and not shoot C—saving its bullet for its partner after that player shoots C—because each player always most prefers its outcome 1.⁵

Thus, thinking ahead about the unpleasant consequences of shooting first or colluding, nobody will shoot or collude. Thereby, all players will survive if the players must act in sequence, giving outcome 3.

This thinking is also rational in the infinite-horizon truel we will describe at the end of the next section. However, there is another point of view, equally rational, that

fired in the air, the third player, acting according to the goals described in the next paragraph, would shoot one of the two disarmed players.)

might be taken in this truel. It yields an ominous outcome, suggesting how conflicts among people, groups, countries, or possibly even larger entities in the universe can lead to death and destruction.

Simultaneous Truels

1. One round. The rules no longer allow the players to choose in sequence, one after another, whereby late choosers learn the choices that other players made earlier. Instead, all three players must now make simultaneous choices of whether or not to shoot, and at which other player, in ignorance of what the other players do (i.e., they cannot communicate with each other to coordinate their choices). This situation is common in life; we must often act before we find out what others are doing.

Now everybody will find it rational to shoot an opponent at the start of play. This is because no player can affect its own fate, but each does at least as well, and sometimes better, by shooting another player—whether the shooter lives or dies—because the number of surviving opponents is reduced.

If each player chooses its target at random, it is easy to see that each has a 25% chance of surviving. Consider player A; it will die if B, C, or both shoot it (3 cases), compared with its surviving if B and C shoot each other (1 case). Altogether, *one* of A, B, or C will survive with 75% probability, and nobody will survive with 25% probability (when each player shoots a different opponent). *Outcome:* There will always be shooting, leaving either one or no survivors.⁶

⁵ Implicitly, we assume that there is no mechanism to enforce agreements, including agreements to collude.

⁶ If there is one survivor, then the three players did not all shoot different opponents. Clearly, the two non-survivors in this situation would be better off if there were no survivors (outcome 4) rather than one surviving opponent (outcome 5). In fact, the strategies associated with outcome 4—each player shoots somebody different—constitute a *Nash equilibrium*, because if any player deviates and shoots the same

2. n rounds ($n = 2$ and known). Assume that nobody has shot an opponent in the first $n - 2$ rounds. We next demonstrate that on the $(n - 1)^{\text{st}}$ round, either at least two players will rationally shoot, or none will.

First, consider the situation in which an opponent shoots A. Clearly, A can never do better than shoot, because A is going to be killed anyway. Moreover, A does better to shoot at whichever opponent (there must be at least one) that is not a target of B or C.⁷

Now suppose that nobody shoots A. If B and C shoot each other, then A has no reason to shoot (though A cannot be harmed by doing so). If one opponent, say B, holds its fire, and C shoots B, A again cannot do better than hold its fire also, because it can eliminate C on the next round. (Note that C, because it has already fired its only bullet, does not threaten A.)

Suppose both B and C hold their fire. If A shoots an opponent, say B, then its other opponent, C, will eliminate A on the n^{th} round. But if A holds its fire, the game passes onto the n^{th} round and, as discussed earlier, A can expect a 25% chance of survival. Thus, if nobody shoots, A again cannot do better than hold its fire.

Whether the players refrain from shooting on the $(n - 1)^{\text{st}}$ round or not—each strategy may be a best response to what the other players do—shooting will be rational on the n^{th} round if there is more than one survivor and at least one player has a bullet remaining. But the anticipation of shooting on the n^{th} round may cause strategies to

opponent as someone else, the deviator does worse (outcome 5 for it). However, to put this Nash equilibrium into effect would require that the players communicate and coordinate their choices, which we have ruled out. In fact, there are three other Nash equilibria in which two players shoot each other and the third holds its fire, but we reject them on the grounds that the strategy of holding one's fire is *dominated*—there is another strategy (in this case, there are two: shooting one or the other of one's opponents) that is never worse and sometimes better.

⁷ As we showed in note 6, when all players fire at different targets, these strategies constitute a Nash equilibrium. This firing occurs immediately, for reasons that will be spelled out in note 8.

“unravel” back to the 1st and 2nd rounds.⁸ *Outcome:* There will always be shooting, leaving one or no survivors.

3. n rounds (n unlimited). The new wrinkle here is that it may be rational for no player to shoot on any round, leading to the survival of all three players. How can this happen?

Our argument earlier that “if you are shot at, you might as well shoot somebody” still applies. But even if you are, say, A, and B shoots C, you cannot do better than shoot B, making yourself the sole survivor (outcome 1). As before, you do best—whether you are shot at or not—if you shoot somebody who is not the target of anybody else, beginning on round 1.

But now suppose that B and C refrain from shooting in round 1, and consider A’s situation. Shooting an opponent is not rational for A on round 1, because the surviving opponent will then shoot A on the next round (there always is a next round if n is unlimited). On the other hand, if all players hold their fire, and if they continue to do so in subsequent rounds, then all three players remain alive.

While there is no “best” strategy in all situations,⁹ the possibilities of survival increase if n is unlimited. *Outcome:* There may be zero, one (any of A, B, or C), or three survivors, but not two survivors.

⁸ Here is the argument for unraveling: On the n^{th} round (n known), players will always shoot if they have any bullets remaining; knowing that this choice is optimal on the last round, players can do no worse than make this choice on the $(n - 1)^{\text{st}}$ round, treating the $(n - 2)^{\text{nd}}$ round as if it were the next -to-last round. Eventually, this reasoning will carry the players back to the 1st round, treating it as if it were the next -to-last round and the 2nd round as if it were the last round. Shooting may therefore be rational on the 1st and 2nd rounds.

⁹ This is because what is best for a player depends on what the other players do. By contrast, not being the first to shoot in the sequential duel we analyzed at the beginning is a *dominant* strategy—it cannot be improved upon, whatever the other players do.

4. Infinite-horizon. This truel is really a variant of situation 3 above that incorporates a more realistic feature. Specifically, at the end of round i and all subsequent rounds, a random event occurs that determines whether the truel continues at least one more round (with probability p_i at the end of round i) or ends immediately (with probability $1 - p_i$). Thus, the probability that a truel ends after exactly k rounds is $p_1 p_2 \dots p_{k-1} (1 - p_k)$. The truel is bounded if and only if $p_i = 0$ for some round i .

If the truel is not bounded (i.e., is infinite horizon), it models games that—like life itself—do not continue forever. While we cannot say at what point such games end, we know they do not continue indefinitely. In such circumstances, if p_i is sufficiently high on each round i , it may be rational never to shoot (Brams and Kilgour, 1998, show that this is also true for a sequential truel with a fixed order of play).

Yet the structure of such games means that the players can anticipate that the truel will end with virtual certainty after several rounds. For example, if $p_i = .51$ for all i , there is a probability of $1 - (.51)^{20} = .9999986$ that, after 20 rounds of play, the game will have terminated. Effectively, then, this can be thought of as an n -round game (n known), à la situation 2, in which there is only slightly more than one chance in a million (i.e., probability .0000014) that the game will not end by round 20.

Applying the reasoning of situation 2 by treating the virtual certainty of termination as a certainty, the players will shoot in rounds 1 or 2, leaving at most one survivor.¹⁰ *Outcome:* How many survivors there are depends on whether the truel is viewed as bounded (at most one player survives) or unbounded (all three players may

¹⁰ This non-cooperative outcome does not depend on how far ahead—2 rounds, 20 rounds, or more—the players project the truel will end. Whatever this point is—even if it is determined probabilistically (as in the infinite-horizon game)—the players' rational choices at the start, applying backward induction, will be to shoot immediately.

survive if p_i is sufficiently high).

A Tale of Two Futures¹¹

Our analysis of the infinite-horizon truel shows that there may be a conflict between two possible futures:

1. Every process must end by some definite point (e.g., each person's lifespan has an upper bound of, say, 125 years);
2. The precise end is unpredictable (it may be highly unlikely that a 125-year-old person will live to be 126, but it is not impossible).

Future 1, in which play is bounded, always leads to shooting in a simultaneous truel, whereas future 2, in which play is unbounded, may induce restraint.

In fact, something akin to future 2 has been argued to be essential in sustaining cooperation in games like repeated Prisoners' Dilemma (PD). If the number of rounds n is known, then play in a repeated PD will, in theory, be non-cooperative, just as it is in the one-round and n -round (n known) simultaneous truels.

But both experimental results and real-life examples of repeated PDs demonstrate that cooperation frequently occurs, which in theory can occur if the "shadow of the future" is sufficiently long.¹² Cooperation may also be rational—even in a one-shot PD and other games, such as Chicken—if the rules of play allow for farsighted thinking

¹¹ The remainder of the paper is adapted from Brams and Kilgour (1998), though the truels analyzed therein are different from those discussed here; see also Bossert, Brams, and Kilgour (2001).

¹² Axelrod (1984) shows that when players follow a strategy of tit-for-tat in repeated PD, the shadow of the future induces cooperation if players do not discount future payoffs too much or, equivalently, the game continues to a new round with a sufficiently high probability. But the accumulation of payoffs, round by round, in a repeated PD is very different from the round-by-round play of a simultaneous truel, wherein there are no payoffs until the game ends. An infinite-horizon truel, we believe, best models the eschatology of lives (and worlds) that will definitely end—and some reckoning, in terms of rewards and punishments, will occur at the end—even though precisely when this end will occur is unknown.

according to the “theory of moves” (Brams, 1994) and some other variants of standard game theory.

More generally, cooperation can be sustained only if there is a sufficient level of *hope*—some reasonable expectation that cooperation will occur in the future. If this hope vanishes, or there is a good prospect of its doing so, then non-cooperative play can be expected of rational players. In games like PD and Chicken, such play will generally end in conflict, although this need not be the case in other games.

As a case in point, outcome 3 in an infinite-horizon truel, in which nobody fires, is consistent with future 2, whereas outcomes 4 and 5 for a player are consistent with either future 1 or future 2. It seems that some real-world players have adhered more to the thinking of future 2, including the United States, Russia, and China: although each has possessed nuclear weapons for more than a generation, all have refrained from using them against each other in anything resembling a truel.

The same self-restraint manifested itself with the non-use of poison gas in World War II, partly in response to revulsion against its use in World War I and partly in fear of reprisal. By contrast, Bosnian Serbs, Bosnian Muslims, and Croats engaged in a very destructive truel in the former Yugoslavia in the early and mid-1990s, mirroring the boundedness of future 1.

Effectively, the Serbs fired the first shot, apparently thinking that quickly conquering territory would give them a big edge. After their early victories, however, they did not fare well because of the reactions of other players, including not only the original parties to the conflict but also new players, like NATO—especially after the conflict expanded to Kosovo in 1998.

Everybody would be better off, we believe, if players did not think they were so clever as to be able to reason backward, from some endpoint, in plotting each other's destruction. Indeed, our results suggest that players would be less aggressive if the future were seen as somewhat murky—as in the infinite-horizon truel—which would render predictions about how many rounds a game will go, or even an upper bound on this number, hazardous. This murkiness, oddly enough, is consistent with hope for the future.

Alternatively, a sequential truel in which the order of choice is endogenous will induce cooperative behavior. As we showed earlier, if any of A, B, or C contemplates shooting first, it ensures its own death when the remaining survivor takes aim. In this case, it is clarity—because there will be retribution—rather than murkiness that induces cooperation.

Conclusions

Two possible eschatological views underlie bounded and unbounded play. To the degree that the future seems to stretch out indefinitely, people probably act more responsibly toward each other, knowing that tomorrow they may pay the price for their untoward behavior today. To sustain themselves, these people may try to develop reputations, often by adhering to certain moral strictures. On the other hand, those who take a more short-term or bounded view may act less responsibly, even immorally.

An important intellectual task is to devise institutions that render destructive behavior unprofitable. But how one makes the future seem to run on smoothly, and instill confidence that the social fabric will not suddenly unravel, is not so clear.

We think the best institutions for this purpose are those that strongly suggest, if not promise, a day of reckoning for those who depart egregiously from norms of fair play.

To return to the Yugoslav example, it is unlikely that the parties who committed the most heinous crimes anticipated the involvement of the International Court of Justice and possible criminal trials.

Likewise, many terrorists seem to look for safe havens from which they will not be extradited. To the extent that international norms of justice not only sanction but also ensure, albeit in the indefinite future, punishment for serious crimes everywhere—including those across national borders—then parties that fire the first shot will be less confident that that shot will be decisive.

Short of ensuring future punishment, institutions that becloud the future, making predictions difficult, may also help to deter reprehensible actions. These institutions range from democracy, with its uncertain electoral futures and other vicissitudes, to extended nuclear deterrence, which offers a good if not certain prospect of protection to allies that might be attacked by an aggressor.

The possibility that these institutions or norms will set in motion forces to reward nonviolent behavior may be analogous to the preventive role of a third player in a truel. Although highly simplified as a social model, the truel does capture an essential feature of social behavior—third parties may play an important role in attenuating conflict.

Their presence, it seems, eases the desperation one often finds in two-player conflicts, which can end up as wars of attrition. The third player, in essence, provides a balancing mechanism that helps to sustain hope, whether the future is murky or clear.

References

- Axelrod, Robert (1984). *The Evolution of Cooperation*. New York: Basic Books.
- Bossert, Walter, Steven J. Brams, and D. Marc Kilgour (2001). "Cooperative vs. NonCooperative Truels: Little Agreement, But Does That Matter?" *Games and Economic Behavior* (forthcoming).
- Brams, Steven J. (1980). *Biblical Games: A Strategic Analysis of Stories in the Old Testament*. Cambridge, MA: MIT Press.
- Brams, Steven J. (1983). *Superior Beings: If They Exist, How Would We Know? Game-Theoretic Implications of Omniscience, Omnipotence, Immortality, and Incomprehensibility*. New York: Springer-Verlag.
- Brams, Steven J. (1994). *Theory of Moves*. Cambridge, UK: Cambridge University Press.
- Brams, Steven J., and D. Marc Kilgour (1998). "Backward Induction Is Not Robust: The Parity Problem and the Uncertainty Problem." *Theory and Decision* 45, no. 3 (December): 263-289.
- Carse, James P. (1986). *Finite and Infinite Games*. New York: Ballantine.
- Holsti, Ole R., Richard A. Brody, and Robert C. North (1964). "Measuring Affect and Action in International Relations Models: Empirical Materials from the 1962 Cuban Missile Crisis." *Journal of Peace Research* 1: 170-189.
- Kilgour, D. Marc, and Steven J. Brams (1997). "The Truel." *Mathematics Magazine* 70, no. 5 (December): 315-326.