

Rational Rules of Thumb in Finite Dynamic Games N-person Backward Induction with Inconsistently Aligned Beliefs and Full Rationality

Yanis Varoufakis

Department of Economics, University of Athens, 8 Pasmazoglou Street, Athens 10559, Greece

Abstract: Recent work has cast considerable doubt on the plausibility of specific assumptions about how rational agents form out-of-equilibrium beliefs in finite extensive games in which beliefs are induced backwards. The point is that the resulting consistently aligned beliefs are incoherent in view of the counterfactuals they rely on. This study asks: How will the possibility of inconsistently aligned beliefs affect the manner in which rational players play such games? It shows that, provided beliefs are aligned monotonically, the more interesting qualitative features of the conventional approach remain unchanged.

Key words: Rational rules, finite dynamic games, full rationality

INTRODUCTION

Imagine a table piled up with G gold sovereigns. Two or more players take turns to collect either one or two coins at a time. If the active player collects one coin, then the next gets a chance to do the same. If on the other hand she collects two coins, the game ends. For this reason taking one coin (or playing ACROSS) will be thought of as a 'cooperative' move thus labelling the taking of two coins (playing DOWN) a 'defection'.

Figure 1 offers the extensive form representation of the game which points to a paradoxical solution: Under the composite assumption that players' beliefs are (a) formed by backward induction and (b) are subject to common knowledge of instrumental Rationality (CKR), the game ends immediately with the first player taking two coins. That this conclusion is paradoxical there is no doubt: Firstly, experimental evidence does not support it^[1]. Secondly, it does not get easier to accept the more intelligently we think about it (especially if G is large). Indeed there has been a number of philosophical and logical objections to the legitimacy of imposing (a) and (b) above simultaneously, an analytical move tantamount to assuming that agents invariably entertain consistently aligned beliefs (CAB)^[2-8].

This study asks: How will rational players who recognise the illegitimacy of the CAB assumption play this game? Will they act in a manner qualitatively different to that prescribed by models which retain CAB after introducing uncertainty about the rationality of one's opponent? The conclusion is that, provided beliefs are aligned monotonically (albeit inconsistently) we can retain the more interesting features of the latter without taking

steps (such as assuming CAB) which are difficult to defend on philosophical grounds. However there is a price to pay as the solution depends on an arbitrary choice regarding the degree of alignment of people's beliefs.

Nevertheless this might be inevitable since the major point of the critical literature mentioned in the previous paragraph is precisely that strategic behaviour yields inherently unpredictable degrees of belief alignment.

The two-person version: Backward induction, together with CKR, leads to the robust conclusion that no instrumentally rational player will ever take just one coin. Yet the paradox here is that in order to work out why, one needs to consider what will happen at the last stage of the game first, then at the penultimate stage ... and so on (that is, it must be pre-supposed that players have chosen only one coin many times already). It is clear that if their beliefs were consistently aligned, the game would have not moved into these later stages.




Fig. 1: Game Strategy

But in order to work out these beliefs we need to consider these stages; a messy sequence of counterfactuals which can only be tamed provided we are prepared to assume that agents go that far into the game as a result of random mistakes (often referred to as ‘trembles’) which occur with tiny probability, are independent of agents’ beliefs and uncorrelated across stages. Evidently the longer the game the greater the amount of ‘trembling’ people must consider as probable before they work out (backwards) what it is rational to believe at the outset and thus, the less convincing the theory. Additionally, the more coins there are on the table the more difficult it is for instrumentally rational players to discern the difference between ‘trembles’ and bluffs (There have been a number of attempts to move away from the implausible assumption of uncorrelated errors-see for example Fudenberg, Kreps and Levine).

Starting with a recognition of these difficulties with CAB, let us begin with a question: “Why would an instrumentally rational player ever choose only one coin when it is her turn to play?” Answer: “Only if she had rational grounds to expect that the next player will also choose one coin with probability at least 1/2”. Consider the stage with k coins left on the table at which player A^k is active and let:

- p_k =Probability that player A^k will choose one coin (i.e., play ACROSS)
- π_k =Probability that A^k is motivated by non-instrumental reasons
- q_k =Probability that A^k is motivated by instrumental reasons but will still choose only one coin

An instrumentally rational player chooses in a manner that maximises her payoffs given the rules of the game and her beliefs about the other player. Hence an instrumental player will always take two coins when there are 3 left on the table. However, she may resist the temptation and pick up only one coin if there are k(>3) coins left and she expects her opponent also to take a single coin during the next stage. This is what is meant by an ‘instrumental reason’ for choosing one rather than two coins at k. By contrast a player who is motivated differently (e.g., is concerned with fairness, or has adopted some universalisable principles of practical reason, or follows a social convention of sharing) is assumed always to choose one coin. This assumption could of course be relaxed by introducing an exogenous probability with which a non-instrumental player chooses two coins. For simplicity, we assume that this probability is zero.

Let us focus on an instrumental player at stage k+1. For the game to have reached k, this means that A^{k+1} set p_k to be greater than 1/2. Eq. 1 captures her expectation:

$$E_A^{k+1}(p_k) = E_A^{k+1}(\pi_k) [\tilde{I}E_A^{k+1}(\pi_k)] \tag{1}$$

$$q_k; q_k = E_A^{k+1} \left[\Pr(E_A^k(p_{k-1}) > 1/2) \right]$$

where, E_A^k(.) denotes the expectation of player A^k who is acting (and thinking) instrumentally. The assumption of CAB ensures that:

$$E_A^{k+1} \left[\Pr(E_A^k(p_{k-1}) > 1/2) \right] = E_A^k \left[\Pr(E_A^k(p_{k-1}) > 1/2) \right] \tag{2i}$$

And:

$$E_A^{k+1}(\pi_k) = E_A^k(\pi_k) \tag{2ii}$$

It is easy to see how, under backward induction, the above two conditions mean one of two things: Either q_i = 0 for i = 2,3,...,G-1, which means that the instrumentally rational player who opens the game (i.e., player A^{G-1}) does not expect the other to take with probability more than 1/2 only a single coin when there are 3 left i.e., E_A^{G-1}(π₃) < 1/2. Or, if E_A^{G-1}(π₃) > 1/2, q_i (> 0) is common knowledge (given CAB) and is computed by means of Bayes’ rule backwards^[9].

Let us now consider the case in which players do not trust that the conditions for CAB [2(i) and 2(ii)] should be taken for granted. As an example, consider first the stage where k = 3. Clearly, q₃ = 0 and therefore p₃ = π₃. At stage k = 5, p₅ will exceed π₅ provided q₅>0. Would it be rational for player A⁵ to entertain such an expectation? The moment we are prepared to accept the possibility that rational players got to stage k = 5 without assuming that they did so as a result of uncorrelated, independently and identically distributed random errors (that is, as long as we allow for the possibility of inconsistently aligned beliefs), then it is inevitable that q₅>0. Thus it turns out that the probability of a ‘cooperative’ move when there are 5 coins on the table is greater than at the later stage when there are 3 coins left (p₅>p₃). If by symmetry q₇>q₅ then, from Eq. 1 it transpires that p₇>p₅. And so on. In effect, we have come to an important conclusion without any controversial assumptions: Since the propensity of rational players to pick up a single coin when it is their turn to play is an increasing function of the expectation on the left hand side of Eq. 1, the more the coins on the table the more likely that the instrumentally rational player will be ‘cooperative’.

To take this observation further, three basic assumptions are required:

Symmetry-S p_k = E_A^{k+1} (p_kI(A^{k+1})) k, where I(A^{k+1}) is the information/belief set of player A^{k+1}.

i.e., instrumentally rational agents will play cooperatively with a probability that others like them would have estimated in an unbiased manner had they had access to their beliefs. This assumption allows for beliefs to be

inconsistently aligned (since q_k is not known with certainty to player A^{k+1}) yet demands that players have the same computational capacities and thus, makes it possible to trace the path of p_k given assumptions R and M below.

Instrumental reflection on non-instrumental agents- R $\pi_k > 0$ and $\Delta^2 \pi_k \geq 0$.

i.e., there is always a possibility the next player will choose to take a single coin non-instrumentally. Moreover, the chances of this happening cannot decrease with the number of coins left on the table. In the simplest case ($\Delta^2 \pi_k = 0$), this probability is constant and corresponds to the proportion of (non-instrumentally) cooperative persons in the population. In the more general case, instrumental agents reflect that the larger the number of coins left, the greater the possibility that normative expectations favouring cooperation will emerge which cannot be explained instrumentally.

Monotonically aligned beliefs-M Eq. (3):

$$q_k \{= E_A^{k+1}[\Pr(E_A^k(p_{k-1}) > 1/2)] = f(p_{k-1}) \text{ where } f'(\cdot) > 0 \quad (3)$$

Condition (3) replaces (2i). Whereas (2i) imposes a strict equality between the beliefs of player A^{k+1} and of A^k viz. the chances that A^k will expect a cooperative move at stage $k-1$, condition (3) issues the far less stringent (and therefore defensible) requirement that their beliefs are linked monotonically. This is equivalent to the thought that, if one is attempting to assess the probability, say γ , of another person predicting that some other probability, say δ , exceeds $1/2$, then it is reasonable to expect that γ will be an increasing function of δ . Clearly, this assumption imposes some alignment between players' beliefs without going to the extremes of the CAB axiom. How much alignment there will be, of course, depends on the precise functional form of $f(\cdot)$. The point of the critical literature on the question of alignment^[2-8] is that due to the inherent unpredictability of human nature, there exists no unique $f(\cdot)$, i.e., one derivable in a uniquely rational manner.

The repercussion of the three assumptions above is simple: Eq. 1 reduces to the difference Eq. (4):

$$p_k = \pi_k + (1-\pi_k)f(p_{k-1}) \quad (4)$$

Given some idea about the form of $f(\cdot)$ and the probability that a player will cooperate for non-instrumental reasons when there are k coins on the table, we can trace the path of the probabilities of cooperative moves by instrumental players. A similar, yet independent, sequence can be found for q_k .

The n-person game: With N players taking turns to collect their one or two coins from the table, it is clear that cooperation requires either a large number of coins or a smaller short term advantage from defection. To extend the analysis so that it applies to a range of payoffs, suppose the rules specify that a player whose

turn it is to act can collect either D or C coins, where $D > C$ (that is, D corresponds to the defection strategy and C to the cooperative move. So far in our game $D = 2$ and $C = 1$). Let $d = (D-C)/D$. Then Eq. 4 generalises to (5) below in which, again, p_k is the probability assessment by player A^k of the likelihood of player A^{k-1} cooperating (if given a chance):

$$p_k = \pi_k + (1-\pi_k)f\left(\prod_{m=2}^M p_{k-m}, d\right) \quad (5)$$

where $M = N$ if $(k - N) \geq \frac{1}{1-d} \vee M = k - 2$ otherwise

where, $f(y, d)$ is the probability with which A^k expects the next player to expect y to be greater than or equal to d , given that she expects it to equal y .

Clearly the condition for an instrumentally rational player A^k to cooperate by choosing C coins (when there are k left on Eq. 6i and 6ii):

$$\prod_{m=0}^M p_{k-m} \geq d \quad (6i)$$

$$p_k = \pi_k \forall k \in \left[\frac{1}{1-d}, N + \frac{1}{1-d} \right] \quad (6ii)$$

The table is given by (6i) while the initial condition of difference Eq. (5) is in (6ii).

Naturally the way in which players' beliefs are aligned-i.e., function $f(\cdot)$ -determines the value of (5). Even though it is a premise of this paper that a unique $f(\cdot)$ ought not be imposed, it is interesting to explore different specifications. Consider those implying that A^k will be certain (or totally undecided) of the next $k-N$ players' decision only if she were totally certain (or undecided) herself if in their position; i.e., $f(0,d) = 0$; $f(1,d) = 1$ and $f(d,d) = 1/2$. It is easy to show that, under these restrictions, cooperative moves are likely by instrumentally rational agents. Table 1 reports on the minimum number of coins that must be left on the table in the two-person game before an instrumentally rational player cooperates (i.e., for condition (6i) to apply). The numbers correspond to the simple case where $f(y,d) = y/2d$.

Table 2 and 3 extend this exploration to the 3-person version of the game. Table 2 utilises a linear function $f(\cdot)$ while Table 3 adopts a non-linear variation of that function. (For an explanation of these tables see the Appendix).

Table 1: $N = 2, f(y,d) = y/2d$

$\pi_j \backslash d$	0.2	0.3	0.4	0.5
0.1	5	6	8	10
0.01	12	22	36	55
0.001	84	184	324	505
0.0001	804	1804	3204	5005

Table 2: $N = 3$, $f(y,d) = y/2d$; Empty cells denote that for cooperation to emerge the number of coins on the table must be infinite

π_k/d	0.025	0.05	0.1	0.2
0.1	6.000	7.00	10.0	58.0
0.01	14.000	24.00		
0.001	88.000			

Table 3: $N=3$; $f(y,d) = \Phi(\alpha+\beta y-d)$ where Φ is a linear probit and α, β are chosen in a way that $f(y,d)$ is never one standard deviation away from the $f(y,d)$ function used in Table 2. Different choices for parameters α and β correspond to different assumptions about the degree of belief alignment between the three players

π_k/d	0.025	0.05	0.1	0.2
0.1	5-6	6-7	9-11	56-60
0.01	13-15	23-25		
0.001	87-93			

CONCLUSION

Unlike models which tackle the same theoretical problem by making particular assumptions which specify detailed stories about the players' out-of-equilibrium beliefs^[10] in which normal form mistakes (or trembles) are introduced, i.e., trembles which are perfectly correlated across information sets. Another example is the popular approach of^[9] which preserves a rigid structure of uncorrelated trembles while introducing more than one type of player, each with a specific probability. The model in this paper is based on very mild assumptions. Indeed its starting point is the recognition that, in this type of game, it is not desirable to start with detailed stories about how deviations from the 'equilibrium' path are to be interpreted by players. Its conclusion is that in addition to being theoretically undesirable such detailed stories/assumptions are not even necessary. Moreover, the analysis offered herein on the basis of the inevitability of at least some inconsistency of rational beliefs seems to be in tune with the most recent results from controlled laboratory experiments^[11]

The reason why stringent assumptions about beliefs are undesirable is that pre-specifying particular patterns of trembles is incompatible with instrumental rationality in view of the counter-factual logic inherent in inducing beliefs via backward induction. On the positive side, the message of the paper is that, even without such detailed stories, the important qualitative results usually derived from restrictive (and thus controversial) assumptions can survive without them. By making only minimalist assumptions (e.g., that people's beliefs are aligned monotonically, rather than consistently), we can still generate the same intuitively appealing predictions as those generated by means of the more controversial assumptions^[8,9]: The probability of a cooperative move by an instrumentally rational player increases with the number of potential future stages (i.e., coins on the table), with decreases in the number of players, with the expectation that agents may be motivated differently, with a narrowing in the gap between the payoffs from defection and cooperation.

Appendix: All three tables were based on the assumption that π_k is constant for all k . Table 1 was

derived as follows: Condition (6ii) tells us that, for $d = 0.5$, non instrumentally rational player would 'cooperate' as long as there are fewer than 5 coins left on the table (3.66 if $d = 0.4$, 3.428 if $d = 0.3$). The best chances for cooperation correspond to $y = 1$, in which case $f(y,d) = 1/2d$. Thus the probability of a cooperative move with 5 coins on the table equals, at most, $\pi(1-\pi)/2d$. For this move to be instrumentally rational, $\pi(1-\pi)/2d$ must exceed d (see condition (6i)). Clearly it does not when there are 5 coins left for any π when $d = 0.5$ (notice that it does when $d = 0.25$). If $m = k-4$, then the probability of a cooperative move can reach a maximum [$m\pi(1-\pi)/2d$]. For this quantity to exceed $1/2$ (i.e., for a cooperative move to be instrumentally rational with $k = m+4$ coins left), $m = 5.55$. Thus the total number of coins be a minimum of 5.55 plus 4, which equals 10 after rounding. Similarly for the rest of Table 1, 2 was compiled in a similar way. Finally, Table 3 generalises by allowing for non-linear alignment between the players beliefs using a probit specification for $f(y,d)$. The range of the minimum number of coins reported corresponds to a choice of the probit's two parameters such that the divergence from the linear case does not exceed one standard deviation.

REFERENCES

- McKelvey, R. and T. Palfrey, 1992. An experimental study of the centipede game. *Econometrica*, 60: 803-836. <http://www.jstor.org/stable/2951567>
- Binmore, K., 1987. Modeling rational players, part I. *Econo. Philosophy*, 3: 179-214. <http://www.citeulike.org/user/jmvidal/article/1585831>
- Bacharach, M., 1987. A theory of rational decision in games. *Erkenntnis*, 27: 17-55. DOI: 10.1007/BF00169710
- Axelrod, R. and D. Dion, 1988. The further evolution of cooperation. *Science*, 242: 1385-1389. DOI: 10.1126/science.242.4884.1385
- Hollis, M., 1991. Penny pinching and backward induction. *J. Philos.*, 88: 473-488. <http://www.jstor.org/stable/2026602>
- Reny, P., 1993. Backward induction, normal form perfection and explicable equilibria. *Econometrica*, 60: 627-649. <http://www.jstor.org/stable/2951586>
- Varoufakis, Y., 1993. Modern and postmodern challenges to game theory. *Erkenntnis*, 38: 371-404. DOI: 10.1007/BF01128237
- Pettit, P. and R. Sugden, 1989. The backward induction paradox. *J. Philos.*, 86: 169-182. <http://www.jstor.org/stable/2026960>
- Kreps, D.M., P. Milgrom, J. Roberts and R. Wilson, 1982. Rational cooperation in the finitely repeated prisoner's dilemma. *J. Econom. Theor.*, 27: 245-252. DOI: 10.1016/0022-0531(82)90029-1
- Fudenberg, D., D. Kreps and A. Levine, 1988. On the robustness of equilibrium refinements. *J. Econom. Theor.*, 44: 354-380. DOI: 10.1016/0022-0531(88)90009-9
- Binmore, K., J. McCarthy, G. Ponti, L. Samuelson and A. Shaked, 2002. A backward induction experiment. *J. Econom. Theor*, 194: 48-88. DOI: 10.1006/jeth.2001.2910