A Generic Bound on Cycles in Two-Player Games

David S. Ahn*

February 2006

Abstract

We provide a bound on the size of simultaneous best response cycles for generic finite two-player games. The bound shows that no cycle will move through the entire strategy space as long as either player has more than two strategies. This bound increases quadratically in the size of the strategy spaces. It is the tightest possible in the sense that we can construct a generic game with a cycle that achieves the bound.

JEL classification: C62, C72

Keywords: Cournot adjustment, best response cycle, bound

*Department of Economics, U.C. Berkeley, 549 Evans Hall #3880, Berkeley, CA 94720, dahn@econ.berkeley.edu

1 Introduction

The idea that players in a repeated two-person game will play the best response to their opponents' actions in the last period dates at least to Cournot's [1] study of oligopoly. He showed this myopic adjustment process may approach a steady state corresponding to Nash equilibrium. In more modern research, variants of this learning process are used to justify the play of Nash strategies. In particular, various studies investigate whether this process is asymptotically stable and approaches a limit strategy, e.g. [2,3].

This asymptotic analysis considers continuous strategy spaces, often through mixed strategies, as the notion of a limit strategy is vacuous otherwise. In contrast, we examine myopic best response dynamics in a finite game. Rather than identifying stable asymptotic behavior, we study the discrete cycles in these finite games directly. Particularly, we quantify a generic bound on their size. This bound implies that no Cournot adjustment process will cycle through the entire strategy space if either player has more than two strategies. Further, we demonstrate that the mentioned bound is the tightest possible, in the sense that we can construct a game producing a best response cycle of that size.

2 Setup

A finite two-player game $G = \langle S_1, S_2, U_1, U_2 \rangle$ consists of a finite action or strategy space S_i and a utility function $U_i : S_1 \times S_2 \to \mathbb{R}$ for each player $i = 1, 2$. Assume $|S_1| \leq |S_2|$. The best response correspondence $BR_i : S_j \to S_i$ is defined by $BR_i(s_j) = \arg\max_{s_i \in S_i} U_i(s_i, s_j)$. The profile of best responses $BR : S \to S$ is $BR(s) = (BR_1(s_2), BR_2(s_1))$ where $S = S_1 \times S_2$. Let $BR^1 = BR$ and $BR^k = BR \circ BR^{k-1}$.

Definition 1. A **best response cycle** is a subset $C \subseteq S$ such that for all $c, c' \in C$, there exists k such that $c' \in BR^k(c)$.

Note that the conclusion of the definition must hold even if $c = c'$. Of course, if $|C| = 1$, then C is a pure strategy Nash equilibrium. There are other intuitive definitions which are equivalent to this one.

Definition 2. A game is **generic** if BR_i is singleton-valued for each player i.

Given fixed strategy spaces, the set of such games is open and dense in the Euclidean space $\mathbb{R}^{2|S_1||S_2|}$ of possible payoffs. In generic games, two different cycles are always disjoint.

We should mention that the cycles we examine result from simultaneous adjustments, as opposed to a process where players alternate updating their strategies. Also, we are not considering the mixed extension of the game, restricting the analysis to pure strategies.

	α	β
A	$1,0$	$0,1$
B	$0,1$	$1,0$

Figure 1: Matching Pennies

	α	β	γ
A	$1,0$	$0,1$	$0,0$
B	$0,0$	$1,0$	$0,1$
C	$0,1$	$0,0$	$1,0$

	α	β	γ
A	C^1	C^2	D^2
B	D^1	C^3	C^4
C	C^6	D^3	C^5

Figure 2: A 3×3 analog of Matching Pennies

The single best response cycle of Matching Pennies is perhaps the best known of all. Because its cycle occupies the entire strategy space, Matching Pennies is sometimes offered as an example of a particularly chaotic game.

A 3×3 analog of Matching Pennies, which is similar to the children's game of Rock-Paper-Scissors, is presented in Figure 2. The game is generic and has two best response cycles, $C = \{(A,\alpha),(A,\beta),(B,\beta),(B,\gamma),(C,\gamma),(C,\alpha)\}$ and $D = \{(B,\alpha),(A,\gamma),(C,\beta)\}$, with six and three elements respectively. The first cycle C involves updated responses by a single, but alternating, player at a time, so each step in the cycle moves either horizontally or vertically in the strategy space. On the other hand, D involves simultaneously updated strategies by both players at each step, so moves diagonally through the strategies.

This 3×3 game suggests two conjectures. First, the largest cycle is twice as large as S_1. Second, the largest cycle involves updating by one player at a time. These intuitions are generally true when $S_1 \leq 4$. They are also particularly confirmed in the 5×5 version of Matching Pennies, which has three cycles, the largest of which is C with ten elements.

However, these conjectures fail whenever $S_1 \geq 5$. A larger cycle in a 5×5 game can be constructed by appending the 3×3 and 2×2 versions of Matching Pennies, as shown in Figure 2. This game has four cycles. In particular, the cycle E has 12 elements and involves

	α	β	γ	δ	ε
A	$1,0$	$0,1$	$0,0$	$0,0$	$0,0$
B	$0,0$	$1,0$	$0,1$	$0,0$	$0,0$
C	$0,0$	$0,0$	$1,0$	$0,1$	$0,0$
D	$0,0$	$0,0$	$0,0$	$1,0$	$0,1$
E	$0,1$	$0,0$	$0,0$	$0,0$	$1,0$

	α	β	γ	δ	ε
A	C^1	C^2	D^2	E^2	D^9
B	D^1	C^3	C^4	D^4	E^4
C	E^1	D^3	C^5	C^6	D^6
D	D^8	E^3	D^5	C^7	C^8
E	C^{10}	D^{10}	E^5	D^7	C^9

Figure 3: A 5×5 analog of Matching Pennies

	α	β	γ	δ	ε
A	$1,0$	$0,1$	$0,0$	$0,0$	$0,0$
B	$0,0$	$1,0$	$0,1$	$0,0$	$0,0$
C	$0,1$	$0,0$	$1,0$	$0,0$	$0,0$
D	$0,0$	$0,0$	$0,0$	$0,1$	$1,0$
E	$0,0$	$0,0$	$0,0$	$1,0$	$0,1$

	α	β	γ	δ	ε
A	C^1	C^2	D^2	E^2	E^8
B	D^1	C^3	C^4	E^{10}	E^4
C	C^6	D^3	C^5	E^6	E^{12}
D	E^1	E^9	E^5	F^1	F^2
E	E^7	E^3	E^{11}	F^4	F^3

Figure 4: A 5×5 game with a larger cycle

simultaneous adjustments by both players at each stage.

3 Results

The main result provides a bound on the size of best response cycles in a generic finite two-player game. The proof decomposes the cycle into two sequences: each player's initial action, her opponent's response to that action, his response to her response, and so on. The bound is computed by examining the rate of recurrence for both sequences. For example, the E cycle in Figure 2 is:

$$\begin{array}{c|cccccccccccc} S_1 & D & A & E & B & D & C & E & A & D & B & E & C \\ S_2 & \alpha & \delta & \beta & \varepsilon & \gamma & \delta & \alpha & \varepsilon & \beta & \delta & \gamma & \varepsilon \end{array}$$

One sequence can be constructed as the first player's initial action D, then the second player's best response δ to D, the first players best response E to δ, and so on. This essentially proceeds diagonally along the array. Doing similarly for the initial action α for the second player, we can construct two sequences, a and b:

$$\begin{array}{c|cccccccccccccc} & 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 & 8 & 9 & 10 & 11 & 12 & \cdots \\ \hline a & D & \delta & E & \varepsilon & D & \delta & E & \varepsilon & D & \delta & E & \varepsilon & D & \cdots \\ b & \alpha & A & \beta & B & \gamma & C & \alpha & A & \beta & B & \gamma & C & \alpha & \cdots \end{array}$$

Then D recurs in the first sequence at every fourth index, while α recurs in the second sequence at every sixth index. Then the first recurrence of the profile (D, α) will be at the least common multiple of four and six, namely twelve. That first recurrence corresponds exactly to the length of the cycle. If both sequences had shared identical elements, then the first recurrence of the profile would have been at the first recurrent in either sequence. Our main result formalizes this argument.

For any two positive integers m, n, we denote their least common multiple by $\operatorname{lcm}(m, n) =$

$\min\{x > 0 : x/m$ and x/n are integers$\}$. The integers are relatively prime if $\text{lcm}(m, n) = mn$.

Proposition 1. *If G is generic, then the size of any best response cycle is bounded by*

$$2(\max\{\text{lcm}(m, n) : m + n \leq |S_1|\} \vee |S_1|).$$

Proof. Genericity of G guarantees that the best response correspondence is singleton-valued, so we treat $BR : S \to S$ as a function in the proof. Consider a best response cycle C and select any strategy profile $s \in C$. Initialize $C^0 = s$ and let $C^k = BR^k(s)$. We recursively define a sequence $\{a^k\} \subseteq S_1 \cup S_2$. Let $a^0 = s_1$, the first player's strategy at the beginning of the cycle. For any odd index k, let $a^k = BR_2(a^{k-1})$, which is the second player's best response to the previous element of the sequence. For any even index k, let $a^k = BR_1(a^{k-1})$, which is the first player's best response to the previous element of the sequence. Similarly, define the sequence b^k by starting $b^0 = s_2$, $b^k = BR_1(a^{k-1})$ for k odd, and $b^k = BR_2(a^{k-1})$ for k even.

Since C is a cycle, the first player's initial action a^0 must recur in the $\{a^k\}$ sequence. Moreover, since the odd elements of a are from the second player's strategy set, the first recurrence of a^0 must be at an even index $m = \{k > 0 : a^k = a^0\}$. So there exists an even integer m such that $C^m = (a^0, b)$ for some $b \in S_2$. Finally, observe that a^0 recurs at any positive integer multiple zm of m: $a^{zm} = a^0$ for any positive integer z.

Similarly, let n be the minimal $n > 1$ with $b^n = s_2$. Then, mutatis mutandis, $n/2 \leq |S_1|$ and $b^{zn} = b^0$ for all positive integer multiples zn of n. Since the least common multiple $\text{lcm}(m, n)$ is a positive integer multiple of both m and n, we have $a^{\text{lcm}(m,n)} = a^0$ and $b^{\text{lcm}(m,n)} = b^0$. Hence $C^{\text{lcm}(m,n)} = C^0$, so $|C| \leq \text{lcm}(m, n)$.

Let $A_1 = \{a^k : k \text{ even}\} \subseteq S_1$ and $B_1 = \{b^k : k \text{ odd}\} \subseteq S_1$. Similarly, define $A_2 = \{a^k : k \text{ odd}\} \subseteq S_2$ and $B_2 = \{b^k : k \text{ even}\} \subseteq S_2$. Recalling the construction of a and b, $A_2 = BR_2(A_1)$ and $A_1 = BR_1(A_2)$. The former implies $|A_2| \leq |A_1|$ and the latter implies $|A_1| \leq |A_2|$. So $|A_1| = |A_2|$. Similarly, $|B_1| = |B_2|$.

Notice $m \leq 2|A_1|$. If not and $m > 2|A_1|$, then there must exist even integers $m', m'' < m$ such that $a^{m'} = a^{m''}$, implying a^0 never recurs in the $\{a^k\}$ sequence. This in turn contradicts the assumption that C is a cycle. Also, $m \geq 2|A_1|$. If not and $m < 2|A_1|$, then there must exist some strategy $a^k \in A_1$ such that a^0 appears twice before a^k appears once, which is also a contradiction. So $m = 2|A_1|$. Similarly, $n = 2|B_2|$. Since we just showed $|B_1| = |B_2|$, this means $n = 2|B_1|$.

Case 1: $A_1 \cap B_1 \neq \emptyset$. Then there exist indices j, k such that $b^j = a^k$. Select any strategy $s_1 \in A_1$. Since s_1 is infinitely recurrent, there exists some even integer i such that $s_1 = a^{k+i}$.

But, since $b^{j+i} = a^{k+i}$, s_1 is also an element of B_1. So $A_1 \subseteq B_1$. Symmetrically, $B_1 \subseteq A_1$. Hence $A_1 = B_1$. Then $m = 2|A_1| = 2|B_1| = n$. So $\mathrm{lcm}(m,n) = m = 2|A_1| \leq 2|S_1|$. Thus $|C| \leq 2|S_1| \leq 2(\max\{\mathrm{lcm}(m,n) : m+n \leq |S_1|\} \vee |S_1|)$.

Case 2: $A_1 \cap B_1 = \emptyset$. Then $\mathrm{lcm}(m,n) = \mathrm{lcm}(2|A_1|, 2|B_1|) = 2 \cdot \mathrm{lcm}(|A_1|, |B_1|)$, and $|A_1| + |B_1| \leq |S_1|$ since A_1 and B_1 are disjoint subsets of S_1. Thus $|C| \leq \mathrm{lcm}(|A_1|, |B_1|) \leq 2(\max\{\mathrm{lcm}(m,n) : m+n \leq |S_1|\} \vee |S_1|)$. □

The genericity assumption cannot be dropped. A game where both players are indifferent to all outcomes obviously has a cycle that covers the entire strategy space and exceeds the bound.

The cycle examined at the beginning of this section, the E cycle of Figure 2 falls into the second case of the proof. The first player's actions in the top sequence a are disjoint from his actions in the bottom sequence b.

Corollary 2. *If G is generic and $|S_2| > 2$, then no best response cycle will cycle through the entire strategy space.*

Proof. Case 1: $|C| \leq 2|S_1| < |S_1||S_2|$. Case 2: $|C| \leq 2 \cdot \mathrm{lcm}(m,n) \leq 2mn < m^2 + 2mn + n^2 = (m+n)^2 \leq |S_1|^2 \leq |S_1||S_2|$. □

Matching Pennies, shown in Figure 1, is often given as an example of a chaotic game. The best response dynamic cycles through the entire strategy space. The Corollary shows that Matching Pennies is exceptionally pathological; all games with cycles that cover their strategy spaces are ordinally equivalent to Matching Pennies. It is still possible for a union of cycles to cover the strategy space, as the game in Figure 2 demonstrates.

The expression of Proposition 1 is a bit untidy. But, if the game is large enough, the expression simplifies to a quadratic function of $|S_1|$. The simplification implies that the bound increases quadratically with the size of the smaller strategy space. It also implies that any best response cycle will be strictly less than half the size of the entire game: $|C| < \frac{|S_1||S_2|}{2}$.

Corollary 3. *If G is generic, then the size of any best response cycle is bounded by*

$$\begin{array}{cl} 2|S_1| & \textit{if } |S_1| \leq 4 \\ \frac{|S_1|^2 - 1}{2} & \textit{if } |S_1| > 4 \textit{ and } |S_1| \textit{ odd} \\ \frac{|S_1|^2 - 2|S_1|}{2} & \textit{if } |S_1| > 4 \textit{ and } |S_1| \textit{ even} \end{array}$$

Proof. Any two consecutive integers are relatively prime, so $\mathrm{lcm}(n, n+1) = n(n+1)$. If $|S_1|$ is odd, then $\max\{\mathrm{lcm}(m,n) : m+n \leq |S_1|\} = \frac{|S_1|+1}{2}\frac{|S_1|-1}{2} = (|S_1|^2 - 1)/4$. If $|S_1|$ is even, $\max\{\mathrm{lcm}(m,n) : m+n \leq |S_1|\} = \frac{|S_1|}{2}\frac{|S_1|-2}{2} = (|S_1|^2 - 2|S_1|)/4$. Verifying maximality in either case is straightforward. Since $|S_1| > 4$, either of these quotients is larger than $|S_1|$. □

The final result shows that Proposition 1 cannot be improved. It gives the tightest possible bound, because there always exists a generic game which achieves the bound. Intuitively,if $S_1 \leq 4$, then the $|S_1| \times |S_1|$ analog to Matching Pennies will satisfy the bound. If $|S_1| \geq 5$, select integers m and n which satisfy the bound. Then construct a game with the $m \times m$ analog of Matching Pennies in the northwest portion of the matrix, and the $n \times n$ analog in the southeast portion. This is exactly how the bound is achieved in Figure 2.

Proposition 4. *For fixed strategy spaces S_1 and S_2, there exists a generic game $\langle S_1, S_2, U_1, U_2\rangle$ with a best response cycle achieving the bound in Proposition 1 and Corollary 3.*

Proof. Case 1: $\max\{\mathrm{lcm}(m,n) : m+n \leq |S_1|\} \geq |S_1|$. Let m, n refer to the maximizers. Then pick two subsets of the first player's actions $R_1, T_1 \subset S_1$ such that $|R_1| = m$ and $|T_1| = n$. Similarly construct $R_2, T_2 \subset S_2$. Index R_i as $r_i^1, r_i^2, \ldots, r_i^m$ and T_i as $t_i^1, t_i^2 \ldots t_i^n$. Set $U_2(r_1^k, r_2^k) = 1$ and $U_1(r_1^k, s_2) = 0$ for all $s_2 \neq r_2^k$. Set $U_1(r_1^{k+1}, r_2^k) = 1$ and $U_1(s_1, r_2^k) = 0$ for all $s_1 \neq r_1^{k+1}$, where we adopt the convention that $r_i^{m+1} = r_i^1$. Similarly construct $U_2(t_1^k, \cdot)$ and $U_1(\cdot, t_2^k)$, substituting t for s and n for m in the expressions. These utility functions are generic and the arguments in the proof of Proposition 1 show that these payoffs will produce a best response cycle of size $\mathrm{lcm}(m,n)$.

Case 2: $\max\{\mathrm{lcm}(m,n) : m+n \leq |S_1|\} < |S_1|$. Then $|S_1| \leq 5$ and simple verification confirms that embedding the $S_1 \times S_1$ analog of Matching Pennies satisfies the bound. □

References

[1] Cournot, A., 1971. Researches into the Mathematical Principles of the Theory of Wealth, 1838. A.M. Kelley, New York.

[2] Fisher, F.M., 1961. The stability of the Cournot solution of the oligopoly problem. Rev. Econ. Stud. 28, 125-135.

[3] Hahn, F.H., 1962. The stability of the Cournot oligopoly solution. Rev. Econ. Stud. 29, 329-331.