Games and Economic Behavior 99 (2016) 82–88

Contents lists available at ScienceDirect

Games and Economic Behavior

www.elsevier.com/locate/geb

Note

Identification of payoffs in repeated games ☆

Byung Soo Lee, Colin Stewart *

University of Toronto, Canada

ARTICLE INFO

Article history:
Received 5 February 2016
Available online 26 July 2016

JEL classification:
C72
C73

Keywords:
Payoff identification
Best-response equivalence
Repeated games

ABSTRACT

In one-shot games, an analyst who knows the best response correspondence can only make limited inferences about the players' payoffs. In repeated games with full monitoring, this is not true: we show that, under a weak condition, if the game is repeated sufficiently many times and players are sufficiently patient, the best response correspondence completely determines the payoffs (up to positive affine transformations).

1. Introduction

How much can one infer about players' payoffs in a game based only on their best response correspondences? In static games, such inferences are quite limited; while best responses convey some information about a player's preferences over her own actions for any given profile of the other players' actions, they say nothing about that player's preferences as the others' actions vary. Among other things, this makes welfare comparisons essentially impossible: one can show that for any profile of best response correspondences and any action profile $\mathbf{a}$ in a finite game, there exist payoffs according to which $\mathbf{a}$ is Pareto efficient, and payoffs according to which $\mathbf{a}$ is Pareto dominated, both of which lead to the given best response correspondences.

In repeated games with full monitoring, one can potentially infer much more. To the extent that other players' future actions depend on one's own current action, best responses convey information about preferences over others' actions. We show that this can be enough to fully identify payoffs (up to positive affine transformations). More precisely, as long as no player has an action ensuring that—regardless of others' actions—she obtains her highest possible payoff, the best response correspondences uniquely determine the payoffs when the game is repeated sufficiently many times and players are sufficiently patient.[1]

To illustrate, consider the 2×2 stage game depicted in Fig. 1. First suppose this game is played once, and the best response correspondence for the row player is such that T is a best response if and only if the probability p that the column player assigns to L is at least $p^* \in (0, 1)$, while B is best response if and only $p \leq p^*$. What can we infer about the payoffs? First, $a > c$ and $d > b$. Without loss of generality, through an appropriate positive affine transformation, we may

☆ We are grateful to the editor, Françoise Forges, and two anonymous referees for helpful comments, and to Young Wu for excellent research assistance. Stewart thanks SSHRC for financial support.

* Corresponding author.
E-mail addresses: byungsoolee@gmail.com (B.S. Lee), colinbstewart@gmail.com (C. Stewart).

[1] In the terminology of Myerson (2013), best-response equivalence implies full equivalence under these conditions.

http://dx.doi.org/10.1016/j.geb.2016.07.004
0899-8256/

	L	R
T	a	b
B	c	d

	L	R
T	1	θ
B	0	$\theta + \frac{p^*}{1-p^*}$

Fig. 1. A simple 2×2 game. Payoffs are for the row player. The matrix on the right illustrates the extent to which payoffs can be identified if the game is not repeated.

	L	C	R
T	2	2	2
M	0	3	0
B	3	0	0

	L	C	R
T	2	2	2
M	0	3	0
B	4	-2	-2

Fig. 2. Two games with the same best response correspondence for the row player (when the game is not repeated).

normalize a and c to be 1 and 0, respectively. Second, given this normalization, the row player's indifference between T and B when the column player chooses $p = p^*$ implies that $d - b = p^*/(1 - p^*)$. These conditions determine payoffs up to a constant parameter capturing the row player's preferences across the two columns, and capture all that can be inferred from best responses. In particular, there is a continuum of distinct games having this best response correspondence.

Now suppose the game is played twice without discounting. Consider the row player's best responses to strategies that play L in the first period, followed by L in the second period if the row player played T in the first period, and a mixture assigning probability p to L and $1 - p$ to R otherwise. Suppose we observe that the row player is indifferent between T and B in the first period when the column player uses this strategy with $p = p^{**}$. If $p^{**} > p^*$, then T is the best response for the row player in the second period regardless of her first-period action. Hence the indifference condition is

$$1 + 1 = 0 + p^{**} + (1 - p^{**})\theta,$$

from which we obtain $\theta = (2 - p^{**})/(1 - p^{**})$. Thus we can pin down the exact payoffs. This approach succeeds whenever $\theta > (2 - p^*)/(1 - p^*)$, ensuring that p^{**} is indeed greater than p^*.

By varying the column player's strategy and checking for indifferences for the row player between her first-period actions, one can identify θ in this way regardless of its value. More generally, however, if there is no strategy for the column player that makes the row player indifferent in the one-shot game, then more periods may be needed. For example, if in the game depicted in Fig. 1 we have $a = 1$, $b = 3/2$, $c = 0$, and $d = 1/2$, then three periods are needed; with only two periods, varying the column player's action in the second period does not provide a strong enough incentive for the row player ever to prefer B in the first period. In general, although the number of repetitions needed depends on the payoffs, one can see from the best responses whether the payoffs can be identified.

This example is relatively simple, in part because the row player's payoffs can be identified in the static game up to the addition of a constant to each outcome in one column. In general, this may not be possible, as Fig. 2 indicates: the row player's set of best responses to any mixed strategy of the column player is identical in the two games depicted in the figure, but neither game can be obtained from the other by adding constants to columns.[2]

A number of papers have examined the testable restrictions of equilibrium notions in certain classes of games with various assumptions about what is observable to the analyst (Bossert and Sprumont, 2013; Chambers et al., 2010; Ledyard, 1986; Ray and Zhou, 2001; Ray and Snyder, 2013; Sprumont, 2000). We depart from this line of work in several respects. First, we take the game form as fixed and focus on identification rather than testable restrictions. Second, we do not assume that only equilibrium play is observable.[3] Experimental evidence suggests that subjects are often rational in the sense that they maximize expected utility with respect to some belief, but do not form correct beliefs about others' strategies (see, e.g., Costa-Gomes and Crawford, 2006, Kneeland, 2015). In this case, although players may not play Nash equilibrium, best responses can be observed if beliefs are elicited (as in Nyarko and Schotter, 2002) or determined by experimental design (as in Agranov et al., 2012). Although the assumption that the analyst can observe the full best response correspondence is quite strong, as we discuss below, our results require only knowledge of best responses to a small class of strategies. We do, however, require that the analyst know the extensive form structure of the game (in particular, that payoffs are constant across repetitions of the stage game).

Our work can be viewed as a strategic analogue of the classical problem of identifying preferences based on choices from menus (see, e.g., Arrow, 1959). In the classical model, if the set of menus is rich enough, one-shot choices are sufficient to fully identify preferences. The strategic structure of our setting effectively limits the kinds of menus from which the agent can choose, in which case making future menus contingent on the agent's choice can help to recover more information about preferences.

[2] Morris and Ui (2004) discuss a similar example.

[3] Abito (2015) studies partial identification of payoffs in repeated games based on equilibrium play. Nishimura (2014) considers the testable implications of individual rationality in extensive games when other players may not be rational.