Behavioural Economics:
Introduction to Behavioral Game Theory and Game Experiments

University of Oxford, Michaelmas Term 2013

and

Asia Summer Institute in Behavioral Economics
National University of Singapore, 22 July-2 August 2013

Vincent P. Crawford, University of Oxford, All Souls College,
and University of California, San Diego

(with thanks to Miguel Costa-Gomes, Nagore Iriberri,
Colin Camerer, Teck-Hua Ho, and Juin-Kuan Chong)

Revised 3 July 2013

Introduction to Behavioral Game Theory and Game Experiments

Behavioral game theory combines theory and empirical (mainly experimental) evidence to develop the understanding of strategic behavior needed to analyze economic, political, and social interactions.

Any aspect of behavioral decision theory—e.g. reference-dependent preferences, present-biased preferences, or heuristics and biases in probabilistic judgment—is equally relevant to behavioral game theory.

But behavioral game theory as it has developed so far is quite different from “behavioral decision theory in games”.

Following the “divide-and-conquer” strategy that is often helpful in research, behavioral game theory has considered issues that are unique to games, taking individual decisions as self-interested and rational.

The hope is that behavioral game theory done this way will combine usefully with behavioral decision theory—that behavioral decision theory will be “plug and play” in behavioral game theory.

What is unique about games is that good decisions in them normally depend on predicting other people's decisions.

(Games also raise the issues of social preferences and reciprocity, but their strategic aspects have mostly been modeled using traditional game theory, and they are usually grouped with decisions rather than games.)

Behavioral analyses of the problem of predicting others' decisions fall naturally into two groups, "thinking" and "learning".

When people lack sufficient prior experience with analogous games to have learned how their partners are likely to play in the current game, people's predictions must rely on some form of strategic thinking, by which they model how other players are likely to respond to a game.

When people do have enough prior experience, learning has a strong tendency to override strategic thinking, replacing "thinking" models of others' decision processes with learning models in which other players are predicted to respond as others did to analogous games in the past.

Game experiments

Progress in behavioral game theory depends on empirical evidence, but strategic behaviour is notoriously sensitive to the details of the environment, making theories of behaviour in games “information hogs”.

Field settings therefore usually lack the control of preferences and information and observability of the rules needed to discriminate sharply among theories of strategic behaviour. (There are important exceptions.)

Most empirical studies of strategic behavior therefore rely on laboratory experiments, which with good design often have a decisive advantage.

Field experiments and empirical studies using observational field data are also important, and checking for consistency among their results and the results of laboratory experiments is especially helpful.

But laboratory choices are real choices, and there is no good reason to think theories based on laboratory evidence aren’t useful understanding behaviour in the field (see e.g. Falk and Heckman 2009 *Science*).

Experimental design

I now give an overview of game experiments, and use one to illustrate how thinking and learning often interact to determine outcomes.

The basic design problem is obtaining clear identification of the game that subjects are playing, including their preferences and information.

This is usually done via a design where subjects play a series of games, using the results to test theories of behavior in the games in the series.

In learning studies the games are often simply repeated, with feedback after each play.

In thinking studies the games in the series are often varied, with no feedback to disable learning.

In each case repeated interaction is usually negligible to subjects (e.g. via random pairing from a large group, or large-group interaction with negligible individual influences) to disable repeated-game strategies.

Subjects' preferences are controlled via:

- Salient rewards in money (or binary lotteries) that conform to incentives in the theory being tested

- Neutral framing

- No face-to-face or nonanonymous interaction (unless this is a treatment variable), to minimize "social" effects on preferences.

Subjects' knowledge is controlled via public announcements, practice runs, and tests.

"Public knowledge" is assumed to approximate the theoretical condition of common knowledge, which is often important in game theory but usually cannot be tested for directly.

Case Study: Van Huyck, Cook, and Battalio's (1997 *JEBO*) "continental divide" experiment

I now illustrate design and how thinking and learning interact to determine outcomes via Van Huyck, Cook, and Battalio's experiment.

Seven subjects repeatedly chose simultaneously and anonymously among efforts from 1 to 14, with each subject's payoff determined by his own effort and the median of all subjects' efforts.

After subjects chose their efforts, the group median was publicly announced, subjects chose new efforts, and the process continued Individual decisions were not announced, which probably didn't matter).

The relation between a subject's effort, the median effort, and his payoff was publicly announced via a table as on the next slide.

The payoffs of a player's best responses to each median are highlighted in bold in the table as displayed here (but not to subjects). The payoffs of the (symmetric, pure-strategy) equilibria "all–3" and "all–12" are in large bold.

Continental divide game payoffs

						Median Choice								
Your Choice	1	2	3	4	5	6	7	8	9	10	11	12	13	14
1	45	49	52	55	56	55	46	-59	-88	-105	-117	-127	-135	-142
2	**48**	53	58	62	65	66	61	-27	-52	-67	-77	-86	-92	-98
3	**48**	**54**	**60**	**6**	70	74	72	1	-20	-32	-41	-48	-53	-58
4	43	51	58	65	**71**	**77**	80	26	8	-2	-9	-14	-19	-22
5	35	44	52	60	69	**77**	**83**	46	32	25	19	15	12	10
6	23	33	42	52	62	72	82	62	53	47	43	41	39	38
7	7	18	28	40	51	64	78	75	69	66	64	63	62	62
8	-13	-1	11	23	37	51	69	83	81	80	80	80	81	82
9	-37	-24	-11	3	18	35	57	88	89	91	92	94	96	98
10	-65	-51	-37	-21	-4	15	40	**89**	**94**	98	101	104	107	110
11	-97	-82	-66	-49	-31	-9	20	85	**94**	**100**	105	110	114	119
12	-133	-117	-100	-82	-61	-37	-5	78	91	99	**106**	**112**	**118**	**123**
13	-173	-156	-137	-118	-96	-69	-33	67	83	94	103	110	117	**123**
14	-217	-198	-179	-158	-134	-105	-65	52	72	85	95	104	112	120

There were ten sessions, each with its own, separate subject group.

Half the groups happened to have an initial median of eight or above, and half had an initial median of seven or below.

The results are graphed on the next slide:

- The median-eight-or-above groups converged almost perfectly to the all–12 equilibrium.

- The median-seven-or-below groups converged almost perfectly to the all–3 equilibrium.

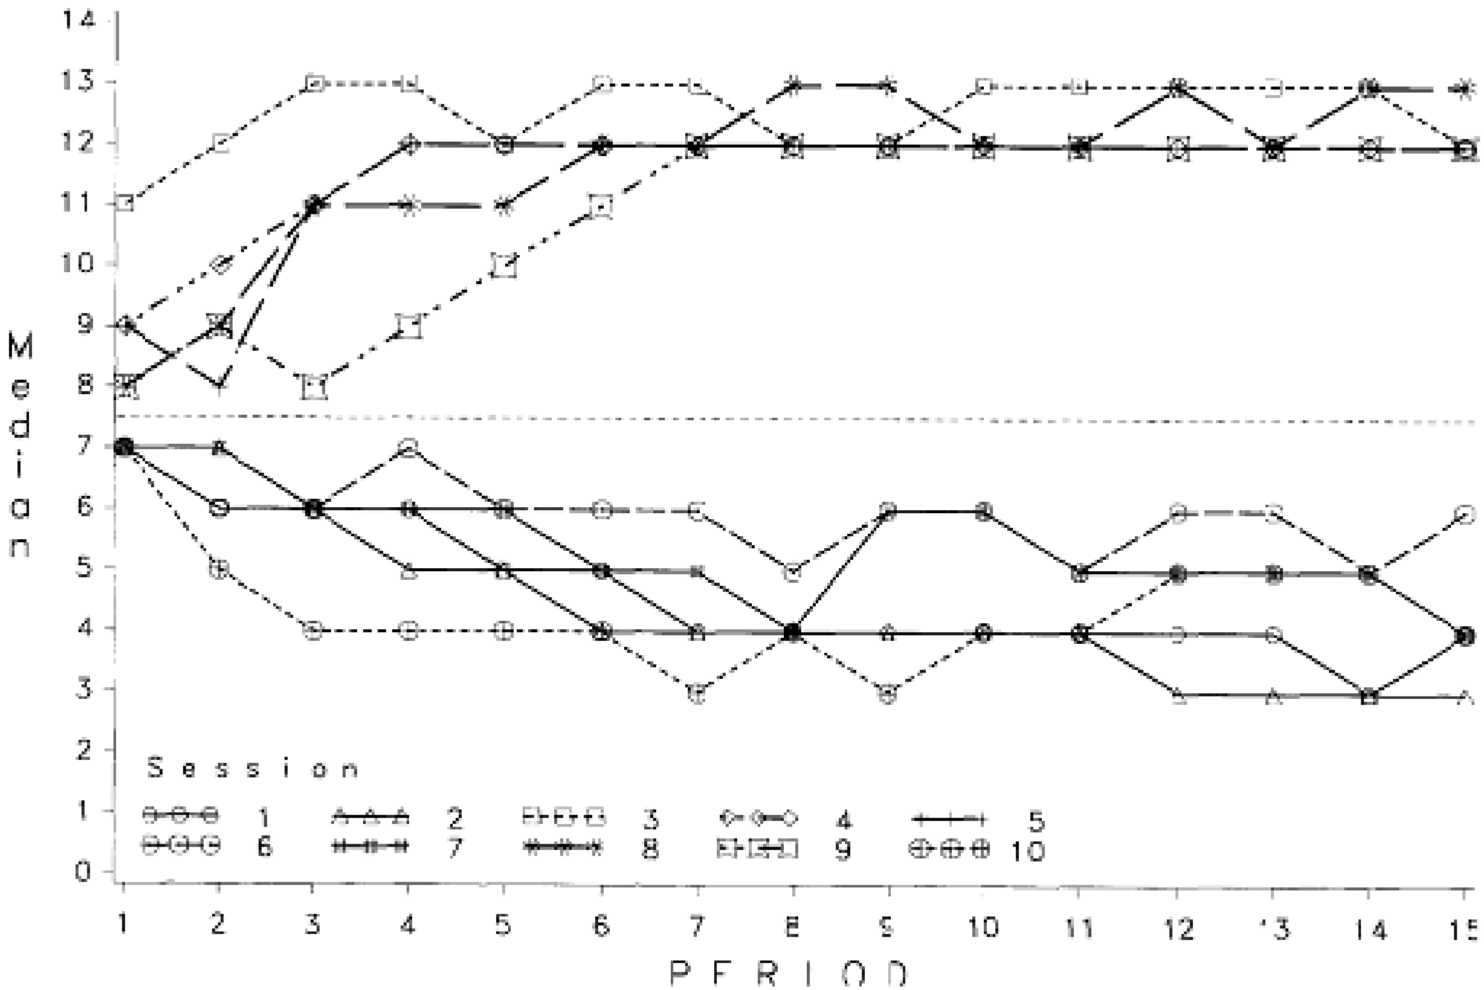

Fig. 3. Median choice in sessions 1 to 10 by period.

A simple adaptive learning model, in which people myopically adjust in the direction of higher payoffs given the most recent value of the group median, seems to describe the dynamics and limiting outcomes well.

The subject groups seemed large enough (and the median robust enough) that subjects treated own influences on the median as negligible.

(Many experiments suggest that the smallest "large" number in behavioral game theory is somewhere between three and five.)

Thus, even though subjects' interaction is formally a repeated game, and phenomena like repeated-game strategies and "strategic teaching" are logically possible in equilibrium, it is appropriate to model subjects' interactions as repeated myopic responses to the stage game.

Equilibrium selection in the limit is completely determined by which equilibrium’s basin of attraction—defined by the myopic best responses described above—subjects’ initial responses fell into.

Thus, it’s not enough to know that learning will eventually converge to some equilibrium, even if we are only interested in (the prior probability distribution of) the final outcome.

We also need to know the prior probability distribution of the median initial response, which determines the final outcome.

In other applications with more complex dynamics, we may also need to know more about the structure of people’s learning rules.