Economics 142: Behavioral Game Theory **Winter 2008**
Vincent Crawford

Many strategic situations in business, international relations, politics, or war are well approximated by games. A *game* is a multi-person decision situation, defined by its *structure*: the players, its "rules" (the order of players' decisions, their feasible decisions at each point, and the information they have when making them); how players' decisions determine the outcome; and players' preferences over outcomes.

Behavioral game theory is a blend of traditional game theory and empirical knowledge whose goal is the understanding of strategic behavior needed for applications.

Such understanding includes topics from behavioral decision theory plus two topics that are specific to multi-person settings: (1) preference interdependence (such as altruism, envy, spite, or reciprocity); and (2) players' models of other players.

Here I narrow the focus to (2), assuming that behavior is (mostly) rational in the decision-theoretic sense and self-interested. I further subdivide (2) into: (2a) how players model others' decisions in initial responses to games with no clear precedents; and (2b) how players learn to predict others' decisions in repeated play of analogous games.

We'll start with a Game Survey designed to highlight some of the issues studied in behavioral game theory. The results will be anonymously tabulated and used as empirical support for some of the ideas developed below.

As just noted, analyses of behavior in games must confront all the issues that arise with individual decisions, plus one that is unique to games:

> Because the outcome is influenced by others' decisions as well as your own, to do well in a game you need to predict others' decisions, taking their incentives into account. This may require a mental model of other players (including a model of their models of you!).

Game theory has a standard model of how people decide what to do in games, which rests on the assumption that people can perfectly predict others' decisions:

> *Nash equilibrium* (often shortened to *equilibrium*) in which each player chooses a decision that is best for himself, given correct expectations about others' decisions

Equilibrium makes clear predictions of game outcomes, which are often accurate when players have learned to predict others' decisions from experience with analogous games (for example, Walker and Wooders, "Minimax Play at Wimbledon," 2001 *AER*).

But in novel situations there may be no analogous games, and equilibrium must then come from sophisticated strategic thinking rather than learning from direct experience.

This makes equilibrium a less plausible assumption, and equilibrium predictions are often much less reliable for initial responses to games than when learning is possible.

In these notes I will briefly discuss the strategic issues addressed by game theory and how the standard theory addresses them, using equilibrium and related notions.

I then outline a behavioral game theory synthesis of models of thinking and learning.

I then compare equilibrium predictions in some simple situations with history, experimental data, or intuitions regarding initial responses to games, highlighting situations where there are systematic deviations from equilibrium predictions.

I then describe a structural but non-equilibrium model of initial responses to games that has emerged from recent experimental work, based on something called "level-*k*" thinking, and compare its predictions with intuition and experimental data.

In simple games a level-*k* model's predictions tend to coincide with equilibrium, so equilibrium predictions rest on a broader and more plausible set of behavioral assumptions, and are correspondingly more reliable.

In more complex games a level-*k* model's predictions can deviate systematically from equilibrium, but in predictable ways. These deviations often bring the model's predictions closer to evidence and intuition, resolving puzzles left open by equilibrium analysis.

I conclude with a brief introduction to learning in games.

Ideas and Issues

Something is *mutual knowledge* if all players know it; and *common knowledge* if all know it, all know that all know it, and so on. I focus on players' problem of predicting others' decisions by assuming that they have common knowledge of the structure of the game. This allows game outcomes to be uncertain if their distributions are common knowledge. The theory does not require common knowledge of the structure, but is easier to explain with it.

Rationality and dominance

	L	**R**
T	2, 2	2, 1
B	1, 2	1, 1

Crusoe vs Crusoe

	Confess	**Don't**
Confess	-5, -5	-1, -10
Don't	-10, -1	-2, -2

Prisoner's Dilemma

Crusoe vs. Crusoe is just two decision problems "traveling together," not really a game; each player has a best decision independent of the other's (a *dominant* decision or strategy; the dominant decision *dominates* (*strictly*, in this case) the other decision).

In Prisoner's Dilemma, players' decisions affect each other's payoffs but each player still has a dominant decision. The game is interesting because individually optimal decisions yield a Pareto-inefficient outcome, highlighting an important distinction between individual and group rationality when there are payoff interactions, even when there are no interactions between players' choices.

Iterated knowledge of rationality and iterated dominance

	Push	Wait
Push	1 5	5 3
Wait	-1 9	0 0

Pigs in a Box

In Pigs in a Box, think of Row (R) as a big pig and Column (C) as a little pig. (The box is a Skinner box, named for the famous behavioral psychologist B.F. Skinner.)

Pushing a lever at one end yields 10 units of grain at the other. Pushing "costs" either pig the equivalent of 2 units of grain. (That is, a pig's utility is his grain consumption, minus 2 if he pushed the lever and minus 0 otherwise.)

If R pushes while C waits, C can eat 5 units before R runs down and shoves C aside.

If C pushes while R waits, C cannot shove R aside, so R gets all but one unit of grain.

If both push and then arrive at the grain together, C gets 3 units and R gets 7. If both wait, both get 0 units of grain.

	Push	**Wait**
Push	**5**, **1**	**3**, **5**
Wait	**9**, **-1**	**0**, **0**

Pigs in a Box

In experiments with real pigs playing the game over and over, if a stable behavior pattern emerges it tends to be at (R Push, C Wait), the equilibrium outcome.

Like some things in game theory, this is surprising, because C gets a better outcome even though R can do anything C can do, and more.

It happens here because C's weakness means that it gets no reward from Push, so Wait dominates Push for C. But when C plays Wait, R does have an incentive to Push.

If the pigs were rational and had studied game theory, they wouldn't have to wait for a stable pattern to emerge: C could figure out that it should Wait no matter what R does; and R could figure out that a rational C will Wait and that R himself should therefore Push.

As this example suggests, strategic thinking tends to yield the same outcome as learning in the long run. (How does learning work in this example?)

Iterated or common knowledge of rationality and rationalizability

	L	C	R
T	0 7	5 0	3 0
M	0 5	2 2	0 5
B	7 0	5 0	3 7

Dominance-solvable game

Now consider Game Survey question 1ab. *Iterated (strict) dominance* yields a unique prediction (which is why the game is called dominance-solvable):

> R is strictly dominated by C for Column.
> When R is eliminated then B is strictly dominated by M for Row.
> When R and B are eliminated then L is strictly dominated by C for Column.
> When R, B, and L are eliminated then T is strictly dominated by M for Row.

So (M, C) is the only outcome that survives iterated elimination of strictly dominated strategies. A game with a unique outcome that does this is called *dominance-solvable*. The set of strategies that survive iterated strict dominance is independent of the order in which dominance is performed. (However, the set of strategies that survive iterated *weak* dominance is *not* independent of the order; and some other results and methods don’t work for iterated weak dominance. In these notes I will focus on strict dominance.)

To see more clearly how outcomes relate to behavioral assumptions, assume that players are *rational* in the sense that they maximize their expected payoffs given *beliefs* about other players' strategies that are not logically inconsistent with anything they know.

Define a *rationalizable* strategy as one that survives iterated elimination of strategies that are *never weak best responses* in that there are no beliefs that make them one of the player's best responses.

	Push	**Wait**
Push	5, 1	3, 5
Wait	9, -1	0, 0

Pigs in a Box

In Pigs in a Box, for example, Push is a best response for R (strict and so also weak) to any beliefs that assign high enough probability to C playing Wait.

But Push is never even a weak best response for C because no beliefs about R's strategy can make it yield as high an expected payoff as Wait. However once Push is eliminated for C, Wait is never a weak best response for R.

Thus the only rationalizable strategies in Pigs in a Box are Push for R and Wait for C. (R Push, C Wait) is also, not coincidentally (why?), the game's unique equilibrium.

More generally, in a two-person game a player's strategies that are never weak best responses are the same as those that are strictly dominated for that player. Thus the strategies that survive iterated elimination of never weak best responses are the same as those that survive iterated elimination of strictly dominated strategies. (However, in three- or more-person games the two ideas are not quite the same. In these notes I will focus on two-person games and use the simpler notion of iterated dominance.)

In the dominance-solvable game from Survey question 1ab:

R is strictly dominated by C for Column.
When R is eliminated then B is strictly dominated by M for Row.
When R and B are eliminated then L is strictly dominated by C for Column.
When R, B, and L are eliminated then T is strictly dominated by M for Row.

Thus the only rationalizable strategies are those that survive iterated elimination of strictly dominated strategies: M for Row and C for Column. (M, C) is also the game's unique equilibrium.

	L	**C**	**R**
T	7, 0	0, 5	0, 3
M	5, 0	2, 2	5, 0
B	0, 7	0, 5	7, 3

Dominance-solvable game

Theorem: If players' rationality is common knowledge, each player must choose a rationalizable strategy. Conversely, any profile of rationalizable strategies is consistent with common knowledge of rationality.

Proof: Consider the first part in the dominance-solvable game. R is strictly dominated by C for Column, so a rational Column will never choose R. When R is eliminated then B is strictly dominated by M for Row, so a rational Row who knows that Column is rational will never choose B. When R and B are eliminated then L is strictly dominated by C for Column, so a rational Column who knows that Row knows that Column is rational will never choose L. When R, B, and L are eliminated then T is strictly dominated by M for Row, so a rational Row who knows that Column knows that Row knows that Column is rational will never choose T. (Thus the number of levels of iterated knowledge of rationality needed is just the number of rounds of iterated dominance, here four; you need *common* knowledge only for indefinitely "large" games.)

	L	**C**	**R**
T	0 7	5 0	3 0
M	0 5	2 2	0 5
B	7 0	5 0	3 7

Dominance-solvable game

The second part of the proof follows by building a “tower” of beliefs in the game from Survey question 2ab, in which there is no strict dominance so any strategy for either player is rationalizable, to support any rationalizable outcome. Assume common knowledge of rationality and consider (T, L) for example:

> Row will play T if he thinks it sufficiently likely that Column will play L.
> Row will think it sufficiently likely that Column will play L if Row thinks Column thinks it likely that Row will play B.
> Row will think it sufficiently likely that Column thinks it likely that Row will play B if Row thinks Column thinks Row thinks it likely that Column will play R.
> And so on.

	L	**C**	**R**
T	**0** **7**	**5** **0**	**7** **0**
M	**0** **5**	**2** **2**	**0** **5**
B	**7** **0**	**5** **0**	**0** **7**

Unique Equilibrium without Dominance

It is easy to check that such a tower can be built for any outcome in this game, which corresponds to the fact that all outcomes survive iterated elimination of strategies that are never weak best responses or, equivalently, strictly dominated strategies.

Rationality, the same beliefs, and Nash equilibrium

The game from Survey question 2ab is typical of economically interesting games in that not only is the game not dominance-solvable, rationalizability implies few (here, no) restrictions on behavior. In such games useful predictions depend on more than rationality or even common knowledge of rationality.

	L	C	R
T	7, 0	0, 5	0, 7
M	5, 0	2, 2	5, 0
B	0, 7	0, 5	7, 0

Unique Equilibrium without Dominance

Note that most of the towers used in the second part of the proof have players believing inconsistent things. For example, in the tower for (T, L) Row plays T because he thinks Column will probably play L, but a rational Column would not play L if he expected Row to play T. The only outcome whose tower has both players believing *consistent* things is (M, C): A rational Row will play M if he thinks Column will probably play C, a rational Column will play C if he thinks Row will probably play M, and so on.

Traditional game theory gets specific predictions in such games by adding an ingredient: the assumption that players have the same beliefs about how the game will be played.

This leads to the idea of *Nash equilibrium* or *equilibrium*, defined as a combination of strategies such that each player's strategy is best for him, given the other's strategy.

If rational players expect a given equilibrium, then their best responses are to play their equilibrium strategies. In the game from question 2ab the unique equilibrium is (M, C), supported by beliefs in which Row expects Column to play C and therefore plays M, Column expects Row to play M and therefore plays C, and so on, at all levels.

Theorem: If players are rational and have the same beliefs about how the game will be played, then their beliefs and (if they are unrandomized, or *pure*) their strategies must be in Nash equilibrium.

Theorem: If players are rational and have the same beliefs about how the game will be played, then their beliefs and (if they are *pure*) their strategies must be in Nash equilibrium.

This theorem leads to the question: Why should players—even rational players—have the same beliefs about how a game will be played?

Traditional game theory offers two complementary answers:

> Thinking: If players have perfect models of each other's decisions, strategic thinking can lead them to have the same beliefs, and so play an equilibrium, even in their initial responses to a game.
>
> Learning: If players don't have perfect models but repeatedly play analogous games, direct experience can eventually allow them to predict others' decisions, and so play an equilibrium in the limit.

Behavioral game theory qualifies these answers by questioning whether people have perfect models of others' decisions and studying the models of others that appear to determine their decisions empirically—mostly experimentally.

Behavioral game theory also studies the interaction between peoples' models of others, how people learn from experience with analogous games, and how their initial responses and learning interact to determine the dynamics and limiting outcome of their strategy choices.

Mixed-strategy equilibrium and equilibrium in beliefs

In Matching Pennies, two players, Row and Column, choose simultaneously between two actions, Heads and Tails; Column wins if they match and Row wins if they mismatch.

Assume players choose their actions to win, which yields a *payoff* of 1, while losing yields -1. Further assume that these payoffs are von Neumann-Morgenstern utilities, so that we can describe players' choices under uncertainty as maximizing their *expected* payoffs.

	Heads	**Tails**
Heads	**-1** **1**	**1** **-1**
Tails	**1** **-1**	**-1** **1**

Matching Pennies

Any strategy in Matching Pennies is rationalizable, so we need more than rationality.

But if players choose only between the *pure* (unrandomized) strategies Heads and Tails, Matching Pennies has no equilibrium; any strategy combination is not best for one player.

But in Matching Pennies it is important to be unpredictable, so it is natural to interpret "choice" to include *mixed* (randomized) strategies as well as pure strategies. Think of choosing a mixed strategy as choosing the *probabilities* p and q of Heads and Tails.

	Heads (q)	**Tails**
Heads (p)	**1**, **-1**	**-1**, **1**
Tails	**-1**, **1**	**1**, **-1**

Matching Pennies

With mixed strategies Matching Pennies has an equilibrium, in which each player plays each pure strategy with probability ½.

If players choose their strategies with probabilities $p = q = ½$ and a player correctly anticipates the other's probability, then Heads and Tails yield him the same expected payoff: $–1p +1(1–p) = 1p – 1(1–p)$ and $1q –1(1–q) = –1q +1(1–q)$. (p makes Column indifferent between Heads and Tails, and q makes Row indifferent.)

No strategy (pure or mixed) has higher expected payoff, and p or $q = ½$ are among players' best choices; so these mixed strategies are in equilibrium.

$p = q = 1/2$ is the only equilibrium in Matching Pennies: If a player could predict a choice probability different than 1/2 for the other player, then one of his pure strategies would yield a higher expected payoff; but Matching Pennies has no equilibrium in pure actions.

The equilibrium in Matching Pennies is best thought of as an *equilibrium in beliefs*, in which each player's mixed strategy represents the other player's beliefs about the first player's realized pure strategy (that is, how the coin lands).

Equilibrium in beliefs is a kind of "rational expectations" equilibrium, in which players form correct expectations about each other's realized pure strategy (not about a market aggregate, as in the notion of rational expectations used in macroeconomics).

The player himself need not be uncertain about his realized pure strategy: It could be nonrandom, for example based on private discussions with subordinates. His realized pure strategy only needs to be unpredictable to the other player.

Further, players' realized pure strategies need not be in equilibrium—even if players' mixed strategies are in equilibrium, they could both end up playing Heads. Only players' beliefs must be in equilibrium, as implied by the last theorem:

Theorem: If players are rational and have the same beliefs about how the game will be played, then their beliefs and (if they are *pure*) their strategies must be in Nash equilibrium.

Now consider the Perturbed Matching Pennies game from Survey Question 7ab.

	Heads (q)	**Tails**
Heads (p)	**2**, **-2**	**-1**, **1**
Tails	**-1**, **1**	**1**, **-1**

Perturbed Matching Pennies

Compare Perturbed Matching Pennies with Matching Pennies and record your intuitions about how to play (pure or mixed strategy, as you prefer) as Row or Column.

Once again there's a unique equilibrium, in mixed strategies. The equilibrium p and q solve:

$-2p + 1(1-p) = 1p - 1(1-p)$, which yields $p = 2/5$ and
$2q - 1(1-q) = -1q + 1(1-q))$, which yields $q = 2/5$.

The relationship between these mixed-strategy probabilities and those for the original, symmetric version of Matching Pennies probably matches your intuition (at least qualitatively) for Column because its better-on-average action, Tails, has probability $3/5 > ½$. But the relationship probably goes against your intuition for Row because its better-on-average action, Heads, has probability $2/5 < ½$.

("Intuition" here comes from decision theory, where increasing the possible reward to a decision would never make a rational person choose it less often.)

The equilibrium must be counterintuitive because if Row tried to exploit the high payoff of Heads too much, and this was predictable, Column could neutralize it by setting $q = 0$). With the predictability that equilibrium assumes, Row can exploit the high payoff of Heads only by setting $p < ½$. The equilibrium p of 2/5 yields Row payoff 1/5, greater than its equilibrium payoff of 0 in Matching Pennies.

This principle is general (see Crawford and Smallwood, *Theory and Decision* 1984) but it seems too subtle to be identified in bridge or poker textbooks or in other informal writing on strategy (it is mentioned in von Neumann and Morgenstern's *Theory of Games and Economic Behavior*).

More issues and ideas: Coordination and multiple equilibria

So far I have focused on games with unique rationalizable outcomes or at least unique equilibria. Yet many economically interesting games have multiple equilibria. Here I discuss only coordination games, which are particularly important in economics. If economics is “about” coordination, we should study coordination in games—not just the coordination that happens (magically) in competitive markets.

	Go	**Wait**
Go	0 0	1 1
Wait	1 1	0 0

Alphonse and Gaston

	Fights	**Ballet**
Fights	1 2	0 0
Ballet	0 0	2 1

Battle of the Sexes

In Alphonse and Gaston and Battle of the Sexes, there are two pure-strategy equilibria (and one mixed-strategy equilibrium), reflecting the *two* ways to solve the coordination problem. Each of the two ways requires them to behave differently when there may be no cues to break the symmetry. Battle of the Sexes complicates the problem with different preferences about *how* to coordinate.

How would you play as Alphonse? As Gaston? As Row or Column in Battle of the Sexes?

(Why “Alphonse and Gaston”? In the early 1900s Frederick B. Opper created the *Alphonse and Gaston* comic strip, with two excessively polite fellows saying “after you, my dear Gaston” or “…Alphonse” and never getting through the doorway. The characters are mostly forgotten, but we still have Alphonse-Gaston games in the dual-control “three-way” lighting circuits in our homes.)

Alphonse and Gaston

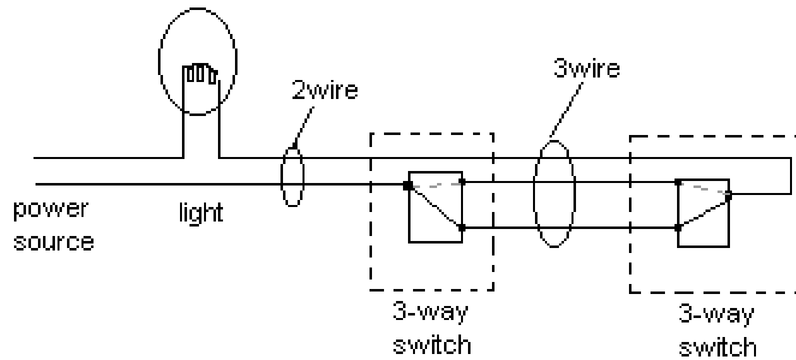

Alphonse and Gaston in your home

Two-Person Stag Hunt

	Other Player: Stag	Other Player: Rabbit
Stag	2, 2	0, 1
Rabbit	1, 0	1, 1

n-Person Stag Hunt

	All Other Players: All-Stag	All Other Players: Not All-Stag
Stag	2	0
Rabbit	1	1

Consider the Stag Hunt game from Survey Question 6ab. The game models a story from Rousseau's *Discourse on Inequality*. It's like an assembly line that won't move faster than the slowest person on the line, a meeting that can't start until everyone is there, or a choice between joining a productive but fragile society and autarky, which is less rewarding but safer because less dependent on coordination. How would you play with one other person? With many other people?

With two or *n* players, there are two symmetric pure-strategy equilibria, "all-Stag" and "all-Rabbit." (There's also an uninteresting mixed-strategy equilibrium.) All-Stag is better for all than all-Rabbit, which is therefore "payoff-dominant". But playing Stag is riskier in that unless all others play Stag, a player would do better playing Rabbit, the more so the more other players there are. All-Rabbit is therefore "risk-dominant" when $n > 2$ and borderline risk-dominant when $n = 2$. In an *n*-person game like Stag Hunt but in which the median rather than minimum choice determined the outcome, all-Stag is much less fragile, even if *n* is large, and may then be risk-dominant as well as payoff-dominant.

Behavioral game theory as a synthesis of strategic thinking and learning

As noted above, traditional game theory offers two complementary answers to the question Why should players have the same beliefs about how a game will be played?:

> Thinking: If players have perfect models of each other's decisions, strategic thinking can lead them to have the same beliefs, and so play an equilibrium, even in their initial responses to a game.
>
> Learning: If players don't have perfect models but repeatedly play analogous games, direct experience can eventually allow them to predict others' decisions, and so to play an equilibrium in the limit.

Behavioral game theory qualifies these answers by questioning whether people have perfect models of others' decisions and studying their models empirically—mostly experimentally.

Behavioral game theory also studies the interaction between peoples' models of others, how they learn from experience, and how initial responses and learning interact to determine the dynamics and limiting outcome.

I now give a concrete example of this interaction.

Consider a “Continental Divide” coordination game from Van Huyck, Cook, and Battalio’s (1997 *JEBO*) experiment.

Seven subjects choose simultaneously and anonymously among “effort” levels from 1 to 14, with each subject’s payoff determined by his own effort and a summary statistic, the median, of all players’ efforts, in a publicly announced way.

The group median is then publicly announced, subjects choose new effort levels, and the process continues.

The relation between a subject’s payoff, his own effort, and the median of all players’ efforts was announced to the subjects by giving them a table like the one on the next page.

In the table as shown on the next page, the payoffs of a player’s best response(s) to each possible median is(are) highlighted in bold; and the payoffs of (symmetric, pure-strategy) equilibria “all–3” and “all–12” are highlighted in large bold.

(For the subjects, of course, there was no highlighting in the table.)

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**	**66**	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 group. Half the groups had an initial median of eight or above, and half had an initial median of seven or below. (I suspect the experimenters cleverly chose the design to make this happen, but it’s not uncommon.)

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.

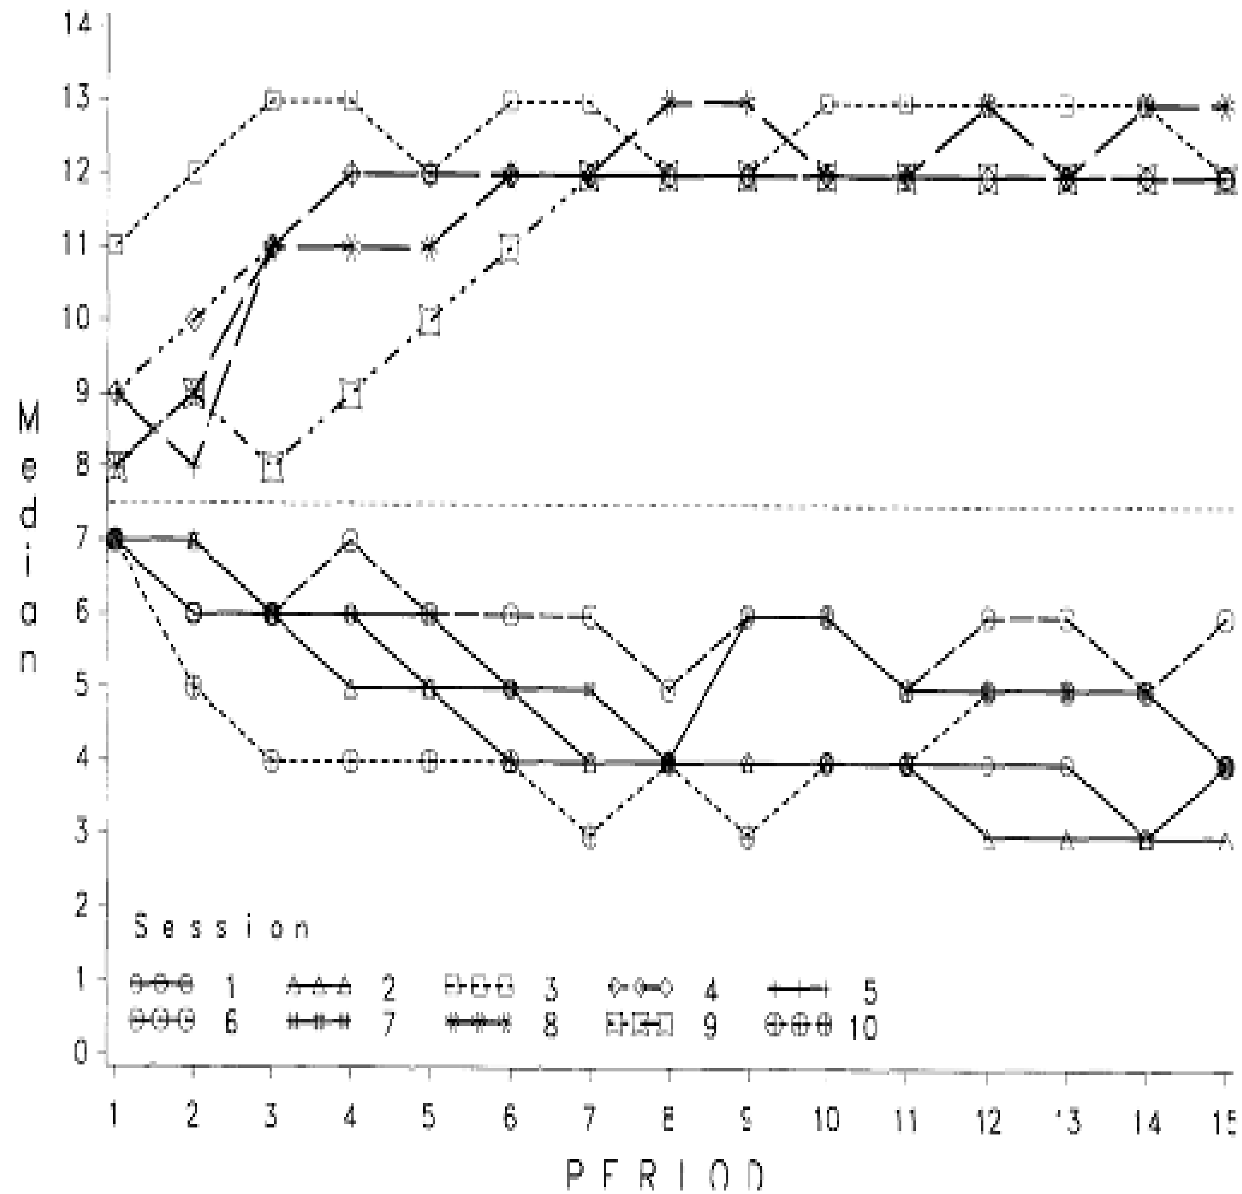

Fig. 3. Median choice in sessions 1 to 10 by period.

As this example makes clear, it’s not enough to know that learning will eventually yield convergence to some equilibrium, even if we are only interested in the final outcome.

To predict the final outcome, we need to know something about the distribution of subjects’ initial responses and the structure of their learning rules.

(Here it seems that we “only” need to know the initial group median, but sometimes, as illustrated below, we need to know much more than that.)

Because subjects had no prior experience, their initial responses are entirely the product of strategic thinking, the main focus of this part of the class.

Learning is also important, and I will spend some time on it at the end. But you can often make an educated guess about the outcome of learning dynamics by using simple graphical techniques, which I will illustrate below.

Experimental studies of strategic thinking in simultaneous-move games with unique pure-strategy equilibria

Consider subjects' initial responses in Nagel's (1995 *AER*) "guessing games" (Survey question 3, the same for groups a and b).

> 15-18 subjects simultaneously guess between [0,100]
>
> The subject whose guess is closest to a p (= 1/2 or 2/3), times the group average guess wins a prize, say \$50
>
> The structure is publicly announced

Record your intuition about what to guess if p = ½. If p = 1/3.

Nagel's games have a unique equilibrium, in which all players guess 0. The games are dominance-solvable, so the equilibrium can be found by repeatedly eliminating stupid (*dominated*, to game theorists) guesses.

For example, if $p = 1/2$:

> It's stupid to guess more than 50 ($1/2 \times 100 \leq 50$)
>
> Unless you think other people are stupid, it's also stupid to guess more than 25 ($1/2 \times 50 \leq 25$)
>
> Unless you think other people think other people are stupid, it's also stupid to guess more than 12.5 ($1/2 \times 25 \leq 12.5$)
>
> And so on, down to 6.25, 3.125, and eventually all the way 0

The rationality-based argument for this "all–0" equilibrium is stronger than the arguments for equilibrium in the other examples, because it depends "only" on iterated knowledge of rationality, not on players having the same beliefs.

But even people who are rational themselves are seldom certain that others are rational, or that others believe that they themselves are rational, and so on; so they probably won't (and shouldn't) guess 0. But what do they do?

Nagel's subjects played the games repeatedly, but we can view their initial guesses as responses to games played in isolation if they treated their influences on the future as negligible. They never guessed 0 initially; their responses were heterogeneous, respecting 0 to 3 rounds of repeated dominance (first picture p = 1/2; second picture p = 2/3):

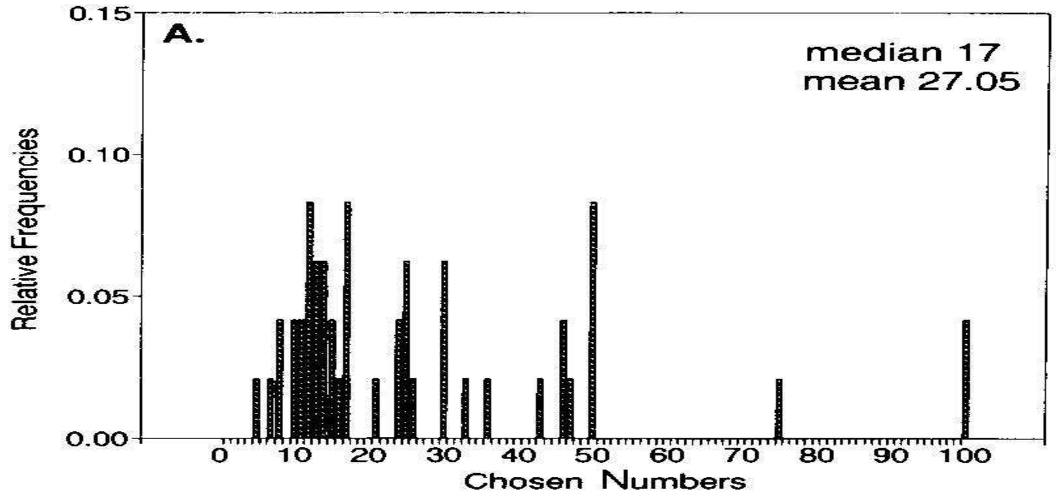

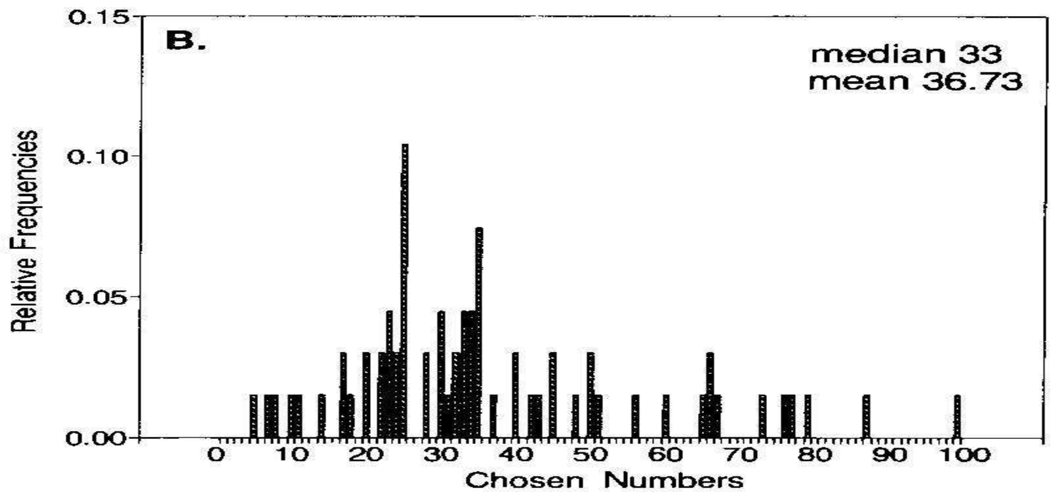

Subjects’ initial responses are coherent and often “strategic”: they respond to the difference between p = ½ and p = 2/3 in the way anyone but a traditional game theorist would expect (equilibrium, and only equilibrium, predicts no response); and they make undominated guesses 85-95% of the time.

But their guesses don’t come close to equilibrium, or even to random deviations from equilibrium (they are systematically biased above equilibrium; and more recent research shows that this kind of bias persists even if equilibrium is not the lowest possible guess).

The data do suggest that the deviations from equilibrium have a coherent but individually heterogeneous structure: spikes are clearly visible (amid the noise) at $50p^k$ for target p and k = 1,2,3,…. (The spikes are like the spectrograph peaks that suggested the existence of discrete chemical elements.)

Similar patterns of heterogeneous but structured non-equilibrium strategic behavior have been found in initial responses to several other kinds of games.

The experimental evidence suggests that although subjects respect dominance for themselves most of the time, they are much less likely to rely on dominance for others.

Further, their reliance on iterated dominance seems to stop at only 1-3 rounds.

Thus equilibrium is too strong to describe behavior, and even rationalizability is too strong.

People often assume that the spikes in Nagel’s data are evidence of subjects doing a finite number of rounds of iterated deletion of dominated strategies.

But in the most recent and most comprehensive experiments on strategic thinking, Costa-Gomes and Crawford (2006 *AER*) showed that many subjects are following decision rules based on “level-*k* thinking,” that one can explain a large fraction of subjects’ deviations from equilibrium using a level-*k* model, and no other model explains a significant fraction.

(Iterated deletion of dominated strategies is not separated from level-*k* thinking in Nagel’s design, but the two notions are strongly separated in Costa-Gomes and Crawford’s design.)

The reason for its empirical success may be that level-*k* thinking yields a workable model of others’ decisions while avoiding most of the cognitive complexity of equilibrium analysis.

As Selten (1998 *EER*) says:

> Basic concepts in game theory are often circular in the sense that they are based on definitions by implicit properties…. Boundedly…rational strategic reasoning seems to avoid circular concepts. It directly results in a procedure by which a problem solution is found. Each step of the procedure is simple, even if many case distinctions by simple criteria may have to be made.

Level-*k* models allow behavior to be heterogeneous, but they assume that each player follows a rule drawn from a common distribution over a particular hierarchy of decision rules or *types*.

Type *Lk* anchors its beliefs in a nonstrategic *L0* type and adjusts them via thought-experiments with iterated best responses: *L1* best responds to *L0*, thus it has a perfect model of the game but a naïve model of others; *L2* best responds to *L1*, thus it has a perfect model of the game and a less naïve models of others; and so on.

In applications the type frequencies are treated as behavioral parameters, estimated or translated from previous analyses. The estimated distribution is fairly stable across games, with most weight on *L1*, *L2*, and *L3*.

Even though few subjects follow the "anchoring type" *L0*, its specification is crucial, representing *L1*'s beliefs, *L2*'s beliefs about *L1*'s beliefs, and so on.

In applications, *L0* is often taken to be uniform random over the feasible decisions.

In Nagel's games, a uniform random (over [0, 100]) *L0* guesses 50 on average.

L1 therefore guesses $50p$, *L2* guesses $50p^2$, and so on, just as in the spikes in Nagel's data.

Costa-Gomes and Crawford (2006 *AER*) elicited subjects' initial responses to a series of 16 dominance-solvable two-person guessing games like those in Survey questions 4ab and 5ab, which are close relatives of Nagel's guessing games.

In each game, two players make simultaneous guesses. Each player has a lower and upper limit, both strictly positive. Each player also has a target, and his payoff increases with the closeness of his guess to his target times the other's guess.

The targets and limits vary independently across players and games, with targets either both less than one, both greater than one, or mixed. The 16 games are dominance-solvable in 3-52 rounds, with unique equilibria determined by the targets and limits.

For example, in game γ4δ3, player *i*'s limits and target are [300, 500] and 1.5 and player *j*'s are [300, 900] and 1.3. The product of targets $1.5 \times 1.3 > 1$, so players' equilibrium guesses are determined by their upper limits; *i*'s equilibrium adjusted guess equals his upper limit of 500, but *j*'s is below his upper limit at 650 (in the figure, guesses in $R(k)$ are eliminated in round *k* of iterated dominance).

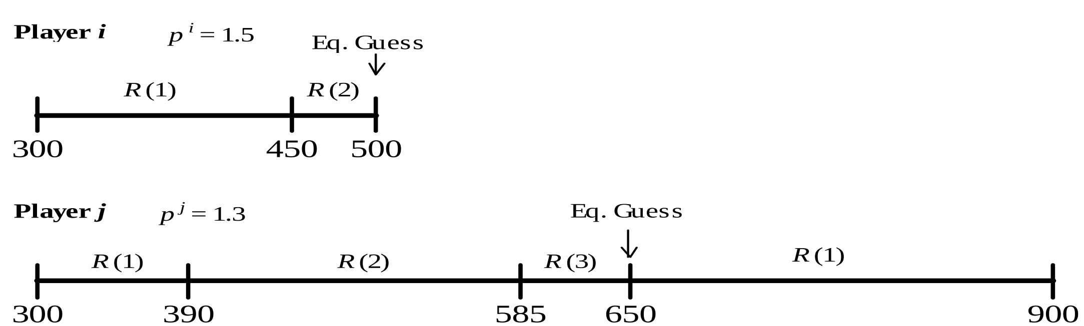

The large strategy spaces and independent variation of targets and limits enhance separation of types' implications for decisions, so many subjects' types can be clearly identified from their guesses. (a_i, b_i, a_j, and b_j are players' lower and upper limits, p_i and p_j are their targets, *D1* and *D2* are "iterated dominance" types, *E* is a type that makes its equilibrium guess, and *S* is a hypothetical "sophisticated" type that can accurately predict others' responses to the games.)

Game	a_i	b_i	p_i	a_j	b_j	p_j	*L1*	*L2*	*L3*	*D1*	*D2*	*E*	*S*
1	100	900	1.5	300	500	0.7	600	525	630	600	611.25	750	630
2	300	900	1.3	300	500	1.5	520	650	650	617.5	650	650	650
3	300	900	1.3	300	900	1.3	780	900	900	838.5	900	900	900
4	300	900	0.7	100	900	1.3	350	546	318.5	451.5	423.15	300	420
5	100	500	1.5	100	500	0.7	450	315	472.5	337.5	341.25	500	375
6	100	500	0.7	100	900	0.5	350	105	122.5	122.5	122.5	100	122
7	100	500	0.7	100	500	1.5	210	315	220.5	227.5	227.5	350	262
8	300	500	0.7	100	900	1.5	350	420	367.5	420	420	500	420
9	300	500	1.5	300	900	1.3	500	500	500	500	500	500	500
10	300	500	0.7	100	900	0.5	350	300	300	300	300	300	300
11	100	500	1.5	100	900	0.5	500	225	375	262.5	262.5	150	300
12	300	900	1.3	300	900	1.3	780	900	900	838.5	900	900	900
13	100	900	1.3	300	900	0.7	780	455	709.8	604.5	604.5	390	695
14	100	900	0.5	300	500	0.7	200	175	150	200	150	150	162
15	100	900	0.5	100	500	0.7	150	175	100	150	100	100	132
16	100	900	0.5	100	500	1.5	150	250	112.5	162.5	131.25	100	187

On average 90% of subjects' guesses respected simple dominance, much more than random (which would be approximately 60%) and typical of initial responses in other experiments.

All but 12 subjects respected dominance in 13 or more games, suggesting that they understood the games and maximized self-interested payoffs, given coherent beliefs.

Of the 88 subjects, 43 made guesses that complied *exactly* (within 0.5) with one type's guesses in 7-16 of the games (20 *L1*, 12 *L2*, 3 *L3*, and 8 *Equilibrium*): far more than could occur by chance, given the strong separation of types' guesses and the fact that guesses could take from 200 to 800 different rounded values.

But 35 of those 43 subjects conformed closely to types other than *Equilibrium*: 20 to *L1*, 12 to *L2*, and 3 to *L3*.

Given the type definitions, those subjects' deviations from equilibrium can be confidently ascribed to non-equilibrium beliefs rather than altruism, spite, confusion, or irrationality.

The other 45 subjects' types are less apparent from their guesses; but econometric estimates still turn up only *L1*, *L2*, *L3*, and *Equilibrium* in significant numbers.

Thus there are no iterated dominance or sophisticated subjects in this population. Subjects seem to find level-*k* thinking quite natural, and iterated dominance thinking rather awkward.

Nagel's and Costa-Gomes and Crawford's analyses were inspired by the famous passage in chapter 12 of Keynes' *General Theory*, in which he likened professional investment

> . . . to those newspaper competitions in which the competitors have to pick out the six prettiest faces from a hundred photographs, the prize being awarded to the competitor whose choice most nearly corresponds to the average preferences of the competitors as a whole; so that each competitor has to pick, not those faces which he himself finds prettiest, but those which he thinks likeliest to catch the fancy of the other competitors, all of whom are looking at the problem from the same point of view. . . . It is not a case of choosing those which, to the best of one's judgment, are really the prettiest, nor even those which average opinion genuinely thinks the prettiest. We have reached the third degree where we devote our intelligences to anticipating what average opinion expects the average opinion to be. And there are some, I believe, who practice the fourth, fifth and higher degrees.

Keynes' wording here suggests finite iteration of best responses, initially anchored by players' true aesthetic preferences: a different, social context-dependent specification of *L0*.

Another intriguing quotation comes from Benjamin Graham (of Graham and Dodd's *Security Analysis*), in *The Intelligent Investor* (thanks to Steven Scroggin for the reference):

> …imagine you are partners in a private business with a man named Mr. Market. Each day, he comes to your office or home and offers to buy your interest in the company or sell you his [the choice is yours]. The catch is, Mr. Market is an emotional wreck. At times, he suffers from excessive highs and at others, suicidal lows. When he is on one of his manic highs, his offering price for the business is high as well…. His outlook for the company is wonderful, so he is only willing to sell you his stake in the company at a premium. At other times, his mood goes south and all he sees is a dismal future for the company. In fact… he is willing to sell you his part of the company for far less than it is worth. All the while, the underlying value of the company may not have changed - just Mr. Market's mood.

Here, Graham is suggesting a best response to Mr. Market, which is a simplified model of other investors. (Although in context, his main goal in the passage from which this quotation comes is to keep you from becoming too emotionally involved with your own portfolio.)

Thus Mr. Market is Graham's *L0* (random, though probably not uniform). So he is advocating being *L1*…. But he published this, so he may actually be *L2*….

And if you ever find yourself in a situation where you need to outguess him, maybe you should be *L3*…but not higher: it can be just as bad to be too sophisticated as to be too unsophisticated.

Level-*k* analyses of strategic thinking

Level-*k* models are a simple, tractable alternative to equilibrium models of initial responses.

Lk for $k > 0$ is rational and *k*-level rationalizable: Its decisions coincide with equilibrium decisions in games that are *k*-dominance solvable. For $k = 2$, 3, or 4 at most, which is empirically plausible, this means level-*k* rules yield equilibrium decisions in games as simple as Pigs in a Box and (for $k = 3$ or 4) the 3×3 dominance-solvable game in the example above.

Although level-*k* rules' simplified models of others sometimes yield the same decisions as equilibrium, so equilibrium predictions can be based on weaker, more plausible assumptions, in other games level-*k* rules deviate systematically from equilibrium.

As a result, a model in which people follow a distribution of *Lk* rules can often predict people's initial responses better than equilibrium.

A level-*k* model usually predicts a distribution of outcomes, but this uncertainty is due to the analyst's inability to observe players' types, not to players' uncertainty about each other; thus the resemblance to mixed equilibrium is superficial.

I now consider some examples that illustrate the potential for using level-*k* models to understand behavior in games.

L0 must often be adapted to the setting; but defining *Lk*, $k > 0$, by iterating best responses "works" in most settings.

Fiction as data? Outguessing in *The Far Pavilions*

In M. M. Kaye's novel *The Far Pavilions*, the main male character, Ash, is trying to escape from his Pursuers along a North-South road; both have a single, *strategically simultaneous* choice between North and South—that is, their choices are time-sequenced, but the Pursuers must make their choice irrevocably before they learn Ash's choice.

> If the pursuers catch Ash, they gain 2 and he loses 2.
>
> But South is warm, and North is the Himalayas with winter coming, so both Ash and the Pursuers gain an extra 1 for choosing South, whether or not Ash is caught.

		Pursuers: **South (q)**	Pursuers: **North**
Ash	**South (p)**	-1, 3	1, 0
	North	0, 1	-2, 2

Far Pavilions Escape

Record your intuitions about what to do, as Ash or Pursuers.

Far Pavilions Escape has a unique equilibrium in mixed strategies, in which $3p + 1(1 - p) = 0p + 2(1 - p)$ or $p = 1/4$, and $-1q + 1(1 - q) = 0q - 2(1 - q)$ or $q = ¾$.

As in Perturbed Matching Pennies, this equilibrium is intuitive for the Pursuers, but not for Ash.

But Ash overcomes his intuition and goes North. The Pursuers unimaginatively go South, so Ash escapes…and the novel can continue…romantically…for 900 more pages.

In equilibrium Ash North, Pursuers South has probability $(1 - p)q = 9/16$, not a bad fit; but try a level-*k* model with uniform random *L0*:

Types	Ash	Pursuers
L0	uniform random	uniform random
L1	South	South
L2	North	South
L3	North	North
L4	South	North
L5	South	South

***Lk* types' decisions in Far Pavilions Escape**

Thus the level-*k* model correctly predicts the outcome provided that Ash is *L2* or *L3* and the Pursuers are *L1* or *L2*.

How do we know which type Ash is? Here fiction provides data on cognition as well: Kaye recounts Ash's mentor's (Koda Dad, played by Omar Sharif in the miniseries) advice (p. 97: "ride hard for the north, since they will be sure you will go southward where the climate is kinder…").

If we take the mentor's "where" to mean "because", then Ash is *L3*:

> Ash thinks the Pursuers are *L2*, and so that the Pursuers think Ash is *L1*, so that the Pursuers think Ash thinks the Pursuers are *L0*.
>
> Thus Ash thinks the Pursuers expect him to go South (because it's "kinder" and the Pursuers are no more likely to pursue him there).
>
> So Ash goes North.

L3 is my record-high *k* for a clearly explained *Lk* type in fiction. (I offer you a $100 reward for the first clearly explained *L4* or higher in fiction.)

Poe's *The Purloined Letter* (http://xroads.virginia.edu/%7EHYPER/POE/purloine.html) has another *L3*, but Conan Doyle doesn't even have an *L1*!

I suspect that even postmodern fiction may have no higher *Lk*s, because they wouldn't be credible.

Outguessing in games like Perturbed Matching Pennies

Camerer reports some (informally gathered) data for a game closely related to the Perturbed Matching Pennies game from Survey Question 7ab (see also Rosenthal, Shachat and Walker, *IJGT* 2003).

	L(33%)	**R(67%)**
T(72%)	2, 0	0, 1
B(28%)	0, 1	1, 0

Perturbed Matching Pennies

The equilibrium mixed-strategy probabilities are Pr{T} = Pr{B} = 0.5 for Row and Pr{L} = 0.33 and Pr{R} = 0.67 for Column.

An *L1* Row plays T and an *L1* Column plays L and R with equal probabilities. An *L2* Row plays T and an *L2* Column plays R. An *L3* Row plays B and an *L3* Column plays R.

With a mixture of 50% *L1*s, 30% *L2*s, and 20% *L3*s in both player roles, the level-*k* model's predicted choice frequencies are 80% T for Row and 25% L for Column: not a perfect fit, but reasonable.

Note that the distribution of heterogeneous types "purifies" the mixed equilibrium.

Coordination in Market Entry/Battle of the Sexes games

In market entry experiments, a number of subjects choose simultaneously between entering ("In") and staying out ("Out") of a market with given capacity. In yields a given positive profit if no more subjects enter than capacity allows; but a given negative profit if too many subjects enter. Out yields 0 profit, no matter what other subjects do.

The natural equilibrium prediction is the symmetric mixed-strategy equilibrium, in which each player enters with a given probability that makes all indifferent between In and Out.

This mixed-strategy equilibrium makes the expected number of entrants approximately equal market capacity, but there is a probability that too many or too few will enter.

Even so, subjects in market-entry experiments have better ex post coordination (number of entrants closer to market capacity) than in the symmetric equilibrium.

This led Kahneman to remark, "…to a psychologist, it looks like magic." (But actually, no one would be at all surprised by this unless he believed in equilibrium, so it would only really look like magic to a game theorist.)

Camerer, Ho, and Chong's (2004 QJE, Section III.C) analysis shows that Kahneman's magic can be explained by a level-*k* model. I now do a similar level-*k* analysis in a simple two-person market-entry game with capacity one, which is like Battle of the Sexes.

	In	Out
In	**0**, **0**	***a***, **1**
Out	**1**, ***a***	**0**, **0**

Market Entry (a > 1)

The unique symmetric equilibrium is in mixed strategies, with $p \equiv \Pr\{\text{In}\} = a/(1+a)$ for both players.

The expected coordination rate is $2p(1 - p) = 2a/(1+a)2$; and players' payoffs are $a/(1+a) < 1$, worse for each than his worst pure-strategy equilibrium.

In the level-k model, each player follows one of four types, *L1*, *L2*, *L3*, or *L4*, with each player role filled by a draw from the same distribution. I assume for simplicity that the frequency of *L0* is 0, and that *L0* chooses its action randomly, with Pr{In}= Pr{Out}= ½.

Higher types' best responses are easily calculated: *L1*s mentally simulate *L0*s' random decisions and best respond, choosing In; similarly, *L2*s choose Out, *L3*s choose In, and *L4*s choose Out.

Types	*L1*	*L2*	*L3*	*L4*
L1	In, In	In, Out	In, In	In, Out
L2	Out, In	Out, Out	Out, In	Out, Out
L3	In, In	In, Out	In, In	In, Out
L4	Out, In	Out, Out	Out, In	Out, Out

The predicted outcome distribution is determined by the outcomes of the possible type pairings and the type frequencies. If both roles are filled from the same distribution of types, players have equal ex ante payoffs, proportional to the expected coordination rate.

L3 behaves like *L1*, and *L4* like *L2*. Lumping *L1* and *L3* together and letting v denote their total probability, and lumping *L2* and *L4* together and letting $(1 – v)$ denote their total probability, the expected coordination rate is $2v(1 – v)$. This is maximized at $v = ½$ where it takes the value ½. Thus for v near ½, which is plausible, the coordination rate is close to ½. (For more extreme values the rate is worse, actually falling to 0 as $v \to 0$ or 1.)

By contrast, the mixed-strategy equilibrium coordination rate, $2a/(1 + a)2$, is maximized when $a = 1$, where it takes the value ½. As $a \to \infty$, the mixed-strategy equilibrium coordination rate converges to 0 like $1/a$. Even for moderate values of a, the level-k coordination rate is higher than the equilibrium rate.

The level-*k* model yields a very different view of coordination than the traditional equilibrium model.

Equilibrium (and equilibrium selection principles like risk- and payoff-dominance) play no role at all in players' strategic thinking.

Coordination, when it occurs, is an accidental (though statistically predictable) by-product of non-equilibrium decision rules.

Finally, even though decisions are simultaneous and there is no possibility of observation of the other player's decision or communication with him, the predictable heterogeneity of strategic thinking allows more sophisticated players such as *L2*s to mentally simulate the decisions of less sophisticated players such as *L1*s and accommodate them, just as Stackelberg followers would, with coordination benefits for all.

This mental simulation doesn't work perfectly, so an *L2* doesn't do as well as if he were really a Stackelberg follower: An *L2* models his partner as an *L1*, but his partner is an *L1* only some of the time.

Neither would it work if strategic thinking were not predictably heterogeneous: As the table shows, if everyone were the same type they would always miscoordinate.

But that it works at all, without communication or observation, is very surprising.

Coordination in Battles of the Sexes with Non-Neutral Framing of Decisions

Crawford, Gneezy, and Rottenstreich (2008 *AER*) randomly paired large subject groups to play games whose payoff structures (except for a symmetric game) were like Battle of the Sexes, but in which there was a commonly observable labeling of decisions, *X* and *Y*, with *X* more salient than *Y*. (Compare Schelling's (1960) classic "meeting in NYC" experiments.)

Although the salience of the *X* label makes it easy and obvious in principle for subjects to coordinate on the "both-*X*" equilibrium, the game still poses a nontrivial strategic problem because both-*X* is one player's favorite way to coordinate but not the other's, and its asymmetric relation to the game's payoffs tempts players to respond asymmetrically.

Just as in a society of men and women playing Battle of the Sexes in which (for cultural reasons) Ballet is more salient than Fights or women's preferences are more salient than men's, there is a tension between the "label salience" of *X* and the "payoff-salience" of a player's favorite way to coordinate: payoff salience reinforces label salience for one player role (Column players or P2s on the next page) but opposes it for players for the other (P1s).

This tension has large and surprising consequences for coordination: Since Schelling's experiments with symmetric games, people have assumed that slight payoff asymmetries would not interfere with coordination. But in these results they have a very strong effect.

The table gives the observed choice frequencies of *X* for both player roles, with subjects in both roles pooled in the symmetric game but not in the other games.

		P2(76%) X	P2(76%) Y
P1(76%)	X	5,5	0,0
	Y	0,0	5,5

Symmetric

		P2(28%) X	P2(28%) Y
P1(78%)	X	5,5.1	0,0
	Y	0,0	5.1,5

Slight Asymmetry

		P2(61%) X	P2(61%) Y
P1(33%)	X	5,6	0,0
	Y	0,0	6,5

Moderate Asymmetry

		P2(60%) X	P2(60%) Y
P1(36%)	X	5,10	0,0
	Y	0,0	10,5

Large Asymmetry

***X*-*Y* games**

The first thing to note is that even tiny payoff asymmetries cause a large drop in the expected coordination rate, from 64% in the symmetric game to 38%, 46%, and 47% in the others.

But even more surprisingly, the pattern of miscoordination completely reversed as the asymmetric *X*-*Y* games progressed from small to large payoff differences:

> With slightly asymmetric payoffs, most subjects in both roles favored their partners' payoff-salient decisions.
>
> But with moderate or large asymmetries, most subjects in both roles switched to favoring their own payoff-salient decisions.

Unless we can understand the reasons for the reversed pattern of miscoordination, we won't really understand why payoff asymmetries cause a large drop in the coordination rate.

I now sketch a level-*k* model, with an *L0* that responds to payoff- and label-salience in a particular, realistic way, that gracefully explains the patterns in the data.

Assume that *L0* responds to both label and payoff salience, but with a "payoffs bias" that favors payoff over label salience, other things equal. In symmetric games *L0* chooses *X* with some probability greater than ½. In any asymmetric game, whether or not label-salience opposes payoff-salience, *L0* chooses its payoff-salient decision with probability p > ½.

Although *L0*'s choice probabilities are the same for P1s and P2s, they imply *L1* and *L2* choice probabilities that differ across player roles due to the asymmetric relations between label and payoff salience for P1s and P2s.

L1's and *L2*'s choices for P1 and P2 are completely determined by p, the extent of *L0*'s payoff bias. A level-*k* model can track the reversal of the pattern of miscoordination between the slightly asymmetric game and the games with moderate or large payoff asymmetries if (and only if) 0.505 (= 5.10/[5.10+5]) < p < 0.545 (= 6/[6+5]), so *L0* has only a modest payoff bias.

Assuming that p falls into this range and that the population frequency of *L1* is 0.7, close to most previous estimates, the model's predicted choice frequencies differ from the observed frequencies by more than 10% only in the symmetric game (where the model somewhat overstates the homogeneity of the subject pool).

In the symmetric game, with no payoff salience, *L0* favors the salience of *X*. *L1* P1s and P2s therefore both choose *X*, and *L2* P1s and P2s follow suit. (Thus in this case the model makes the same prediction as equilibrium selection based on salience as in a Schelling focal point.)

In the slightly asymmetric game, the payoff differences are small enough that *L1* P1s choose *X*, P2s' payoff-salient decision, because *L1* P1s think it is sufficiently likely that *L0* P2s will choose *X* that choosing *X* yields them higher expected payoffs. *L2* P2s, who best respond to *L1* P1s, thus choose *X* as well. By contrast, *L1* P2s choose *Y*, P1s' payoff-salient decision, because *L1* P2s think it is sufficiently likely that *L0* P1s will choose *Y*. *L2* P1s thus choose *Y* as well. In sum, *L1* P1s choose *X* and *L2* P1s choose *Y*, while *L1* P2s choose *Y* and *L2* P2s choose *X*. When $q = 0.7$, the model predicts that 70% of P1s will choose *X* but only 30% of P2s will choose *X*, coming reasonably close to the observed frequencies of 78% and 28%.

Finally, in the games with moderate or large payoff asymmetries, *L0*'s payoffs bias is just as strong. But because the payoffs bias is not *too* strong ($p < 0.545$), the payoff differences are large enough that *L1* P1s and P2s now both choose their own instead of their partners' payoff-salient decisions, *Y* for P1s and *X* for P2s. Because *L2*s best respond to *L1*s in the opposite role, *L2* P1s choose *X* and *L2* P2s choose *Y*. In sum, *L1* P1s choose *Y* and *L2* P1s choose *X*, while *L1* P2s choose *X* and *L2* P2s choose *Y*. When $q = 0.7$, the model predicts that only 30% of P1s will choose *X* but 70% of P2s will choose *X*, again close to the observed frequencies of 33-36% and 61-60%.

Outguessing in Hide and Seek games with non-neutral framing of locations

In Rubinstein and Tversky's experimental Hide and Seek games (see Crawford and Iriberri (2007 *AER*), seekers (Survey question 8a) were told the following story:

> You and another student are playing the following game: Your opponent has hidden a prize in one of four boxes arranged in a row. The boxes are marked as follows: A, B, A, A. Your goal is, of course, to find the prize. His goal is that you will not find it. You are allowed to open only one box. Which box are you going to open?

Hiders (Survey question 8b) were told an analogous story.

Record your intuitions about how to play as Hider. As Seeker.

The ABAA framing of locations is non-neutral in two ways:

> The *B* location is distinguished by its label.
>
> The two *end A* locations are inherently salient.
>
> Together these two saliencies distinguish *central A* as “the least salient location.”

The framing (order and labeling) of the four locations is a tractable abstract model of a cultural or geographic landscape like those that play important roles in real Hide and Seek games.

With a payoff of 1 for winning, RTH's Hide and Seek game translates into:

		Seeker			
		A	**B**	**A**	**A**
Hider	**A**	1 0	0 1	1 0	0 1
	B	0 1	1 0	1 0	0 1
	A	0 1	0 1	1 0	0 1
	A	0 1	0 1	0 1	1 0

Hide and Seek

Like Matching Pennies, Hide and Seek has a unique, mixed-strategy equilibrium, with equal probabilities on all four locations for both players.

Equilibrium leaves no room for the non-neutral framing to influence people's choices.

But in Rubinstein and Tversky's experiments, *central A* was most prevalent for Hiders (37%) and even more prevalent for Seekers (46%); as a result Seekers can expect find a Treasure 32% of the time, more than the 25% with which they would find it in equilibrium.

This raises three puzzles, none of which are resolved by equilibrium (or noisy generalizations of equilibrium like "quantal response equilibrium"):

> If seekers are as smart as hiders on average, why don't hiders who are tempted to hide in *central A* realize that seekers will be just as tempted to look there?
>
> Why do hiders choose actions that allow seekers to find them more than 25% of the time, when they could hold it down to 25% via the equilibrium mixed action?
>
> Why do seekers choose *central A* even more than hiders?

These puzzles can all be gracefully resolved by a level-*k* model in which *L0* is sensitive to the framing of locations.

Assume that with given probabilities, each player role is filled by one of five level-*k* types: *L0*, *L1*, *L2*, *L3*, or *L4*.

Lk, $k > 0$, anchors its beliefs in a nonstrategic *L0* type and adjusts them via thought-experiments involving iterated best responses.

L0 (for hiders and seekers) reflects the simplest hypothesis a player can make about his opponent's instinctive response: that he will choose a salient location, simply because it is salient. Assume that *L0* plays A, B, A, A with probabilities $p/2$, q, $1–p–q$, $p/2$, where $p > 1/2$ and $q > ¼$. (A uniform random *L0* would make *Lk* coincide with equilibrium.) Thus *L0* favors focally labeled and/or end locations, to an equal extent for hiders and seekers. But the model allows the data to decide which is more salient, B or the end locations (the ends seem to be).

Hiders' and seekers' strategic responses to the framing are confined to *Lk*, $k > 0$, which ignores the framing except as it influences *L0*'s choice probabilities.

Given this specification of *L0*, it's not hard to show that:

> *L1* Hiders choose *central A* to avoid *L0* Seekers and *L1* Seekers avoid *central A* in searches for *L0* Hiders
>
> *L2* Hiders choose *central A* with probability between 0 and 1 and *L2* Seekers choose it with probability 1
>
> *L3* Hiders avoid *central A* and *L3* Seekers choose it with probability between 0 and 1
>
> *L4* Hiders and Seekers both avoid *central A*

With a plausible distribution of types estimated from Rubinstein and Tversky's data (0% *L0*, 19% *L1*, 32% *L2*, 24% *L3*, 25% *L4*), the level-*k* model explains their results, including the prevalence of *central A* for hiders and its greater prevalence for seekers.

The asymmetry in hiders' and seekers' behavior follows naturally from their role-asymmetric responses to *L0*, with no asymmetry in behavioral assumptions across roles.

Learning Models

Learning models describe how players adjust their decisions over time in response to experience with analogous games. The learning process is usually modeled as repetition of the same "stage game" (usually with different player groups), so that the analogies are perfect.

The game is played either by a small group randomly selected from one or more populations—for example, random pairing to play a two-person game, with player roles filled either from the same or from identifiable separate populations—or sometimes by the entire population at once as in Van Huyck, Cook, and Battalio's "Continental Divide" game.

Players' decisions and roles in the game are distinguished by commonly understood labels, which are the "language" in which they code their experience, and in which any convention that emerges will be expressed.

Players view their decisions in the stage game as the objects of choice, and the dynamics of their decisions are modeled directly (or indirectly in terms of their beliefs, with decisions modeled as best replies) rather than determined by an equilibrium in the stage game or the repeated game that describes the entire learning process.

Learning is "adaptive" in that strategies adjust in a direction that would increase payoffs, other things (including others' adjustments) equal, given the current state of the system.

Pugilists, Dancers, or Birds of Different Feathers

To make these ideas more concrete, imagine a large population of men and women repeatedly and anonymously paired (with gender publicly observable in each pair, so they can base their strategies on gender if they so choose) to play Battle of the Sexes.

	Fights	**Ballet**
Fights	**1** **2**	**0** **0**
Ballet	**0** **0**	**2** **1**

Battle of the Sexes

Now draw a differential equation "phase diagram" with the population frequency of men playing Fights, *m*, on the horizontal axis and the frequency of women playing Fights, *w*, on the vertical axis. We will use this diagram to analyze the dynamics of simple learning rules.

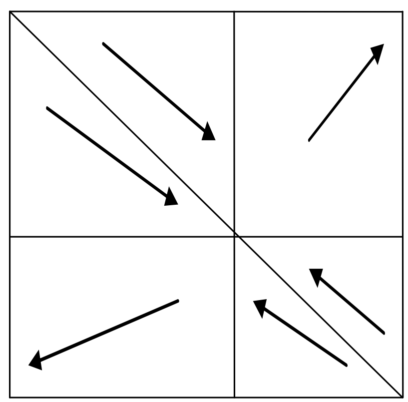

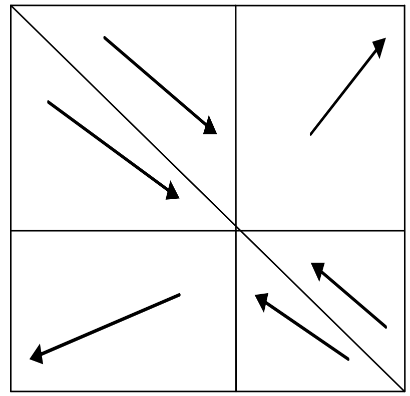

For men the expected payoff of Fights is higher than Ballet whenever $w > 1/3$ ($2w > 1 - w$). For women the payoff of Fights is higher whenever $m > 2/3$ ($m > 2(1 - m)$). There are four regions: ($m > 2/3$, $w > 1/3$), ($m > 2/3$, $w < 1/3$), ($m < 2/3$, $w > 1/3$), ($m < 2/3$, $w < 1/3$). For plausible learning rules, when ($m > 2/3$, $w > 1/3$), m and w rise. When ($m > 2/3$, $w < 1/3$), m falls and w rises. When ($m < 2/3$, $w > 1/3$), m rises and w falls. And when ($m < 2/3$, $w < 1/3$), m and w fall. When ($m > 2/3$, $w > 1/3$), $m \to 1$ and $w \to 1$; and when ($m < 2/3$, $w < 1/3$), $m \to 0$ and $w \to 0$. When ($m > 2/3$, $w < 1/3$) or ($m < 2/3$, $w > 1/3$), if (with symmetry) the initial condition is above the diagonal—$m + w > 1$—the system enters ($m > 2/3$, $w > 1/3$) and $m \to 1$ and $w \to 1$; if it's below the diagonal, the system enters ($m < 2/3$, $w < 1/3$) and $m \to 0$ and $w \to 0$.

In this setting the limiting outcome must be one of the two pure-strategy equilibria, in each of which all people follow a convention based on the commonly understood Fights versus Ballet labeling of their decisions. Which one they will follow is determined by whether the frequencies of initially arrogant men and wimpy women sum to more than half the population.

Now consider a large population repeatedly and anonymously paired to play the same kind of game, with two pure-strategy equilibria, one favored by one player and the other favored by the other; but now with no observable labeling of players or decisions.

Players in this game can still use the payoffs to distinguish their strategies according to which one would yield them the more favorable outcome if their partner coordinated with it.

I follow the evolutionary game theory literature in calling these strategies Hawk (choose the strategy that would yield you the more favorable outcome if your partner coordinates with it, as Fights previously did for men and Ballet did for women) and Dove (choose the decision that would yield your partner the more favorable outcome if he coordinates with it).

With this redescription, in terms of labels that reflect the symmetry of men's and women's strategic positions, we can represent the game symmetrically like this:

	Hawk	**Dove**
Hawk	**0** **0**	**1** **2**
Dove	**2** **1**	**0** **0**

Hawk-Dove Game

There are two equivalent ways to analyze the learning dynamics in this game.

The first is to recycle the phase diagram used to analyze Battle of the Sexes, but to impose the added restriction that the frequency of players playing Hawk must be equal in both player roles. This is just as if in Battle of the Sexes the frequency of men playing Fights, m in my notation, must be equal to the frequency of women playing Ballet, $1 - w$.

Because $m = 1 - w$ is equivalent to $m + w = 1$, this restriction limits the dynamics to the diagonal running from northwest to southeast in the previous two-dimensional phase diagram.

As the diagram suggests, the dynamics will now converge to the intersection of lines in the center, which represents the mixed-strategy equilibrium of the game at Pr{Hawk} = 2/3.

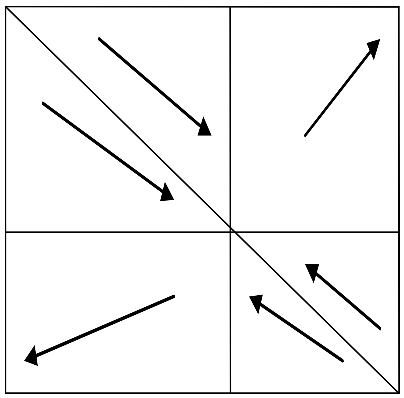

	Hawk	Dove
Hawk	0 0	1 2
Dove	2 1	0 0

Hawk-Dove Game

The second, less magical way to analyze the learning dynamics is to graph the expected payoffs of Hawk and Dove (in either player role) against the population frequency of Hawk. This "builds in" the restriction that the frequency of players playing for their favorite equilibrium must be the same in both roles, and allows us to represent the dynamics in a one-dimensional phase diagram, with expected payoffs of Hawk and Dove on the vertical axis and population frequency of Hawk on the horizontal axis.

When the frequency of Hawk is low, Hawk has higher payoff than Dove, and vice versa. Thus the dynamics follow the arrows on the horizontal axis, converging to the frequency of Hawk where the payoff lines cross, which is 2/3, representing the mixed-strategy equilibrium.

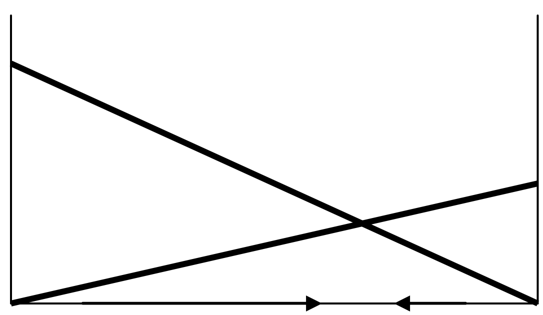

Minimum- and Median-Effort Coordination Games

John Van Huyck, Ray Battalio, and Richard Beil (1990 *AER*, 1991 Q*JE*, 1993 *GEB*; "VHBB") studied games that are like larger versions of Stag Hunt (seven efforts), analyzed in Crawford (1991 *GEB*, 1995 *Econometrica*) and Crawford and Broseta (1998 *AER*).

VHBB's 1990 and 1991 experimental designs

Repeated play of symmetric coordination games in populations of subjects, interacting all at once ("large groups") or in pairs drawn randomly ("random pairing").

Subjects chose simultaneously among 7 efforts, with payoffs and ex post optimal choices determined by own efforts and an order statistic, the population median or minimum effort in large groups or the current pair's minimum with random pairing.

There were five leading treatments, varying the order statistic (minimum in 1990, median in 1991), the size of the subject population, and the patterns in which they interact (minimum games were played either by the entire population of 14-16 or by random pairs, median games were played by the entire population of 9); each population was large enough to make subjects treat their own influences on the order statistic as negligible.

Explicit communication was prohibited throughout, the order statistic was publicly announced after each play (with random pairs told only pair minima), and the structure was publicly announced at the start, so subjects were uncertain only about others' efforts.

Payoff Table Γ

		Median value of X chosen						
		7	6	5	4	3	2	1
Your choice of X	7	1.30	1.15	0.90	0.55	0.10	−0.45	−1.10
	6	1.25	1.20	1.05	0.80	0.45	0.00	−0.55
	5	1.10	1.15	1.10	0.95	0.70	0.35	−0.10
	4	0.85	1.00	1.05	1.00	0.85	0.60	0.25
	3	0.50	0.75	0.90	0.95	0.90	0.75	0.50
	2	0.05	0.40	0.65	0.80	0.85	0.80	0.65
	1	−0.50	−0.05	0.30	0.55	0.70	0.75	0.70

Payoff Table A

		Smallest Value of X Chosen						
		7	6	5	4	3	2	1
Your Choice of X	7	1.30	1.10	0.90	0.70	0.50	0.30	0.10
	6	–	1.20	1.00	0.80	0.60	0.40	0.20
	5	–	–	1.10	0.90	0.70	0.50	0.30
	4	–	–	–	1.00	0.80	0.60	0.40
	3	–	–	–	–	0.90	0.70	0.50
	2	–	–	–	–	–	0.80	0.60
	1	–	–	–	–	–	–	0.70

Other Player

	Stag	Rabbit
Stag	2, 2	0, 1
Rabbit	1, 0	1, 1

Two-Person Stag Hunt

All Other Players

	All-Stag	Not All-Stag
Stag	2	0
Rabbit	1	1

***n*-Person Stag Hunt**

The random-pairing and large-group minimum games are like larger versions of the two-effort Stag Hunt games seen earlier.

The stage games all have seven strict, symmetric, Pareto-ranked equilibria, with players' best responses an order statistic of population efforts.

The games are like a meeting that can't start until a given quorum is achieved—100% in the large-group minimum game, 50% in the large-group median games.

Intuitively, efficient coordination is more difficult, the larger the quorum or the larger the group, but traditional equilibrium analysis and its refinements don't fully reflect this.

VHBB's 1990 and 1991 results

The five leading treatments all evoked similar initial responses (table in Crawford 1991 *GEB*).

TABLE I

		Minimum treatment				
		A (%)	*B* (%)	*A′* (%)	C_d (%)	C_f (%)
Subject's	7	33 (31)	76 (84)	23 (25)	11 (37)	13 (42)
initial	6	10 (9)	1 (1)	1 (1)	1 (3)	0 (0)
effort	5	34 (32)	2 (2)	2 (2)	2 (7)	6 (19)
	4	18 (17)	5 (5)	7 (8)	5 (17)	2 (6)
	3	5 (5)	1 (1)	7 (8)	3 (10)	1 (3)
	2	5 (5)	1 (1)	17 (19)	1 (3)	1 (3)
	1	2 (2)	5 (5)	34 (37)	7 (23)	8 (26)
Totals		107 (101)	91 (99)	91 (100)	30 (100)	31 (99)

		Median treatment		
		Γ, Γdm (%)	Ω (%)	Φ (%)
Subject's	7	8 (15)	14 (52)	2 (7)
initial	6	4 (7)	1 (4)	3 (11)
effort	5	15 (28)	9 (33)	9 (33)
	4	19 (35)	3 (11)	11 (41)
	3	8 (15)	0 (0)	2 (7)
	2	0 (0)	0 (0)	0 (0)
	1	0 (0)	0 (0)	0 (0)
Totals		54 (100)	27 (100)	27 (99)

Subjects almost always converged to some equilibrium, but the dynamics varied with the treatment variables (order statistic, number of players, interaction pattern), with large differences in drift, history-dependence, rate of convergence, and equilibrium selection:

In 12 out of 12 large-group median trials, there was near-perfect "lock-in" on the initial median (even though it varied across runs and was usually inefficient)

In 9 out of 9 large-group minimum trials, there was very strong downward drift, with subjects always approaching the least efficient equilibrium

In 2 out of 2 random-pairing minimum trials, there was very slow convergence, no discernible drift, and moderate inefficiency

Comparing the first two reveals an "order statistic" or "robustness" effect, with coordination less efficient the smaller the groups that can disrupt desirable outcomes.

Comparing the last two reveals a "group size" effect, in which coordination is less efficient in larger groups (holding the order statistic constant, measured from the "bottom").

TABLE III

MEDIAN CHOICE FOR THE FIRST TEN PERIODS OF ALL EXPERIMENTS

	Period									
Treatment	1	2	3	4	5	6	7	8	9	10
Gamma										
Exp. 1	4	4	4	4	4	4	4*	4	4*	4*
Exp. 2	5	5	5	5	5	5	5	5	5	5
Exp. 3	5	5	5	5	5	5	5	5	5	5*
Gammadm										
Exp. 4	4	4	4	4	4	4*	4*	4*	4*	4*
Exp. 5	4	4	4	4*	4*	4*	4*	4*	4*	4*
Exp. 6	5	5	5	5	5	5	5	5*	5*	5*
Omega										
Exp. 7	7	7	7	7*	7*	7*	7*	7*	7*	7*
Exp. 8	5	5	5	5	5*	5*	5*	5*	5*	5*
Exp. 9	7	7	7*	7*	7*	7*	7*	7*	7*	7*
Phi										
Exp. 10	4	4	4	4	4*	4*	4*	4*	4*	4*
Exp. 11	5	5	5	5*	5*	5*	5*	5*	5*	5*
Exp. 12	5	5	5	5*	5*	5*	5*	5*	5*	5*

Notes. Exp. = experiment. * = indicates a mutual best response outcome.

Table 2—Experimental Results for Treatment *A*

	Period									
	1	2	3	4	5	6	7	8	9	10
Experiment 1										
No. of 7's	8	1	1	0	0	0	0	0	0	1
No. of 6's	3	2	1	0	0	0	0	0	0	0
No. of 5's	2	3	2	1	0	0	1	0	0	0
No. of 4's	1	6	5	4	1	1	1	0	0	0
No. of 3's	1	2	5	5	4	1	1	1	0	1
No. of 2's	1	2	2	4	8	7	8	6	4	1
No. of 1's	0	0	0	2	3	7	5	9	12	13
Minimum	2	2	2	1	1	1	1	1	1	1
Experiment 2										
No. of 7's	4	0	1	0	0	0	0	0	0	1
No. of 6's	1	0	1	0	0	1	0	0	0	0
No. of 5's	3	3	2	1	0	0	1	1	0	1
No. of 4's	4	6	2	3	3	0	0	0	0	0
No. of 3's	1	4	2	5	0	1	1	0	1	0
No. of 2's	3	2	6	5	5	9	3	4	3	1
No. of 1's	0	1	2	2	8	5	11	11	12	13
Minimum	2	1	1	1	1	1	1	1	1	1
Experiment 3										
No. of 7's	4	4	1	0	1	1	1	0	0	2
No. of 6's	2	0	2	0	0	0	0	0	0	0
No. of 5's	5	6	1	1	1	0	0	0	0	0
No. of 4's	3	3	2	1	2	1	0	0	0	1
No. of 3's	0	0	7	6	0	2	3	0	0	0
No. of 2's	0	1	1	4	5	3	6	3	2	2
No. of 1's	0	0	0	2	5	7	4	11	12	9
Minimum	4	2	2	1	1	1	1	1	1	1
Experiment 4										
No. of 7's	6	0	1	1	0	0	1	0	0	0
No. of 6's	0	6	2	0	0	1	0	0	0	0
No. of 5's	8	5	5	5	0	1	0	0	0	0
No. of 4's	1	1	4	6	7	1	2	1	1	0
No. of 3's	0	2	3	2	4	3	2	2	1	0
No. of 2's	0	1	0	0	2	3	7	4	2	2
No. of 1's	0	0	0	1	2	6	3	8	11	13
Minimum	4	2	3	1	1	1	1	1	1	1

Table 2—Experimental Results for Treatment *A*, Continued

	Period									
	1	2	3	4	5	6	7	8	9	10
Experiment 5										
No. of 7's	2	2	3	1	1	1	1	0	0	0
No. of 6's	1	3	1	0	0	0	0	0	0	0
No. of 5's	9	3	0	4	1	0	2	0	0	0
No. of 4's	3	4	6	2	1	2	0	2	1	1
No. of 3's	1	2	2	4	6	0	0	0	0	1
No. of 2's	0	2	2	3	4	6	5	2	5	3
No. of 1's	0	0	2	2	3	7	8	12	10	11
Minimum	3	2	1	1	1	1	1	1	1	1
Experiment 6										
No. of 7's	5	3	1	1	1	1	2	2	2	3
No. of 6's	2	0	0	0	1	0	0	0	0	0
No. of 5's	5	1	0	0	0	1	0	0	0	0
No. of 4's	2	3	4	0	0	0	0	0	0	0
No. of 3's	1	5	4	2	2	2	1	0	2	0
No. of 2's	0	2	4	5	3	3	6	4	5	5
No. of 1's	1	2	3	8	9	9	7	10	7	8
Minimum	1	1	1	1	1	1	1	1	1	1
Experiment 7										
No. of 7's	4	3	1	1	1	1	1	1	1	1
No. of 6's	1	0	0	0	0	0	0	0	0	0
No. of 5's	2	3	0	0	0	0	0	0	0	0
No. of 4's	4	0	1	2	1	0	0	0	0	0
No. of 3's	1	3	2	1	1	0	0	0	0	0
No. of 2's	1	3	2	2	4	4	4	4	5	3
No. of 1's	1	2	8	8	7	9	9	9	8	10
Minimum	1	1	1	1	1	1	1	1	1	1

TABLE 5—DISTRIBUTION OF ACTIONS FOR TREATMENT C: RANDOM PAIRINGS

	Period				
	21	22	23	24	25
Experiment 6					
No. of 7's	5	5	4	10	8
No. of 6's	0	1	3	0	0
No. of 5's	2	5	3	3	4
No. of 4's	3	1	1	1	1
No. of 3's	1	1	1	0	0
No. of 2's	1	1	2	2	2
No. of 1's	4	2	2	0	1
Experiment 7					
No. of 7's	–	–	6	5	5
No. of 6's	–	–	1	0	1
No. of 5's	–	–	0	3	0
No. of 4's	–	–	2	1	4
No. of 3's	–	–	2	0	0
No. of 2's	–	–	0	0	1
No. of 1's	–	–	3	5	3

In case you are wondering what would happen with fixed rather than random pairing, here are the results. There is clear evidence of "strategic teaching," with 12 out of 14 pairs managing to "teach" their way to the most efficient equilibrium. Most subjects seemed to understand that strategic teaching is pointless with random pairing, because it's costly but others reap the benefits. But they used it effectively with fixed pairing.

Table 4—Experimental Results for Treatment C: Fixed Pairings

	Period						
	21	22	23	24	25	26	27
Experiment 5							
Pair 1							
Subject 1	7	7	7	7	7	7	7
Subject 16	7	7	7	7	7	7	7
Minimum	7*	7*	7*	7*	7*	7*	7*
Pair 2							
Subject 2	7	2	7	7	7	7	7
Subject 15	1	7	3	6	7	7	7
Minimum	1	2	7	7	7	7	7
Pair 3							
Subject 3	1	1	1	1	1	1	1
Subject 14	1	1	7	1	1	1	7
Minimum	1*	1*	1	1*	1*	1*	1
Pair 4							
Subject 4	1	7	7	7	7	7	7
Subject 13	7	2	5	7	7	7	7
Minimum	1	2	5	7*	7*	7*	7*
Pair 5							
Subject 5	1	7	4	7	7	7	7
Subject 12	1	4	7	7	7	7	7
Minimum	1	4	4	7*	7*	7*	7*
Pair 6							
Subject 6	5	7	7	7	7	7	7
Subject 11	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*
Pair 7							
Subject 7	1	7	6	7	7	7	7
Subject 10	5	3	6	7	7	7	7
Minimum	1	3	6*	7*	7*	7*	7*

Pair 8							
Subject 8	7	6	6	7	7	7	7
Subject 9	3	5	7	7	7	7	7
Minimum	3	5	6	7*	7*	7*	7*
Experiment 6							
Pair 1							
Subject 2	7	7	4	5	6	6	7
Subject 15	2	3	6	6	7	7	7
Minimum	2	3	4	5	6	6	7*
Pair 2							
Subject 3	5	7	7	7	7	7	7
Subject 14	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*
Pair 3							
Subject 4	1	1	1	1	4	4	1
Subject 13	7	1	1	3	1	1	2
Minimum	1	1*	1*	1	1	1	1
Pair 4							
Subject 5	5	7	7	7	7	7	7
Subject 12	7	7	7	7	7	7	7
Minimum	5	7*	7*	7*	7*	7*	7*

TABLE 4—FIXED PAIRINGS, Continued

	Period						
	21	22	23	24	25	26	27
Pair 5							
Subject 6	4	5	7	7	7	7	7
Subject 11	4	5	7	7	7	7	7
Minimum	4*	5*	7*	7*	7*	7*	7*
Pair 6							
Subject 7	5	7	7	7	7	7	7
Subject 10	5	7	7	7	7	7	7
Minimum	5*	7*	7*	7*	7*	7*	7*

* – Denotes a mutual best-response outcome.

Except for the fixed-pairing results, VHBB’s results can be mostly understood via a simple evolutionary basin of attraction story proposed in Crawford (1991 *GEB*, 1995 *Econometrica*).

Other Player

	Stag	**Rabbit**
Stag	**2** **2**	**1** **0**
Rabbit	**0** **1**	**1** **1**

Two-Person Stag Hunt

All Other Players

	All-Stag	**Not All-Stag**
Stag	**2**	**0**
Rabbit	**1**	**1**

***n*-Person Stag Hunt**

Imagine that there are only two efforts as in Stag Hunt, not seven, and graph the expected payoffs of high (Stag) and low (Rabbit) effort against the population frequency of high effort in the random pairing and large-group minimum games and the large-group median game.

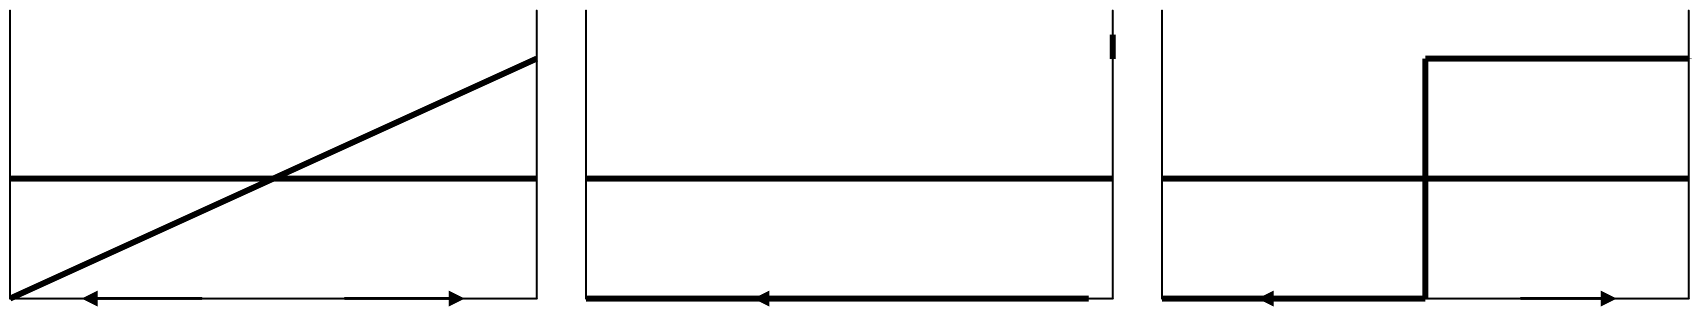

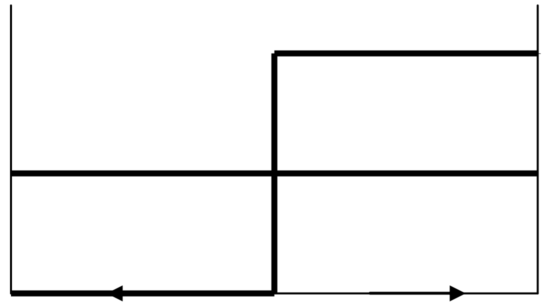

In the large-group median game, the all-Stag and all-Rabbit equilibria are both locally stable.

By symmetry, random shocks are neutral, just as likely to flip the population from all-Stag to all-Rabbit or vice versa.

With random initial conditions, the population would be equally likely to converge to all-Stag or all-Rabbit. If initial conditions favor one equilibrium, its limiting probability is higher.

In the seven-effort version of the game that VHBB studied, if learning always makes subjects adjust their efforts toward the current value of the median, then the population converges to the median without changing it (a general property of order statistics like the median).

Even with random shocks, the median is just as likely to go up as it is to go down.

Either way, the learning dynamics have no up or down trend; and (given the dampening effect of the median on shocks) the population is very likely to "lock in" on the initial median, as in VHBB's experiments.

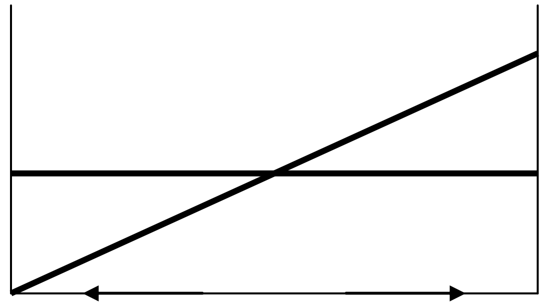

In the random-pairing minimum game, the all-Stag and all-Rabbit equilibria are again both locally stable.

Random shocks are again neutral; and with random initial conditions, the population would be equally likely to converge to all-Stag or all-Rabbit.

Crawford (1995) shows that in the seven-effort version of this game that VHBB studied, it's actually optimal for a (risk-neutral) player to set his effort equal to his forecast of the median effort in the entire population.

Thus, just as in the large-group median game, the learning dynamics have no up or down trend and the population is likely to "lock in" on the initial median.

However, with random pairing a subject samples only a small fraction of the population effort distribution each period (his current partner's effort is an estimate of the population median, but a very noisy one), so convergence will be much slower, as it was in VHBB's experiments.

TABLE 5—DISTRIBUTION OF ACTIONS FOR TREATMENT *C*: RANDOM PAIRINGS

	Period				
	21	22	23	24	25
Experiment 6					
No. of 7's	5	5	4	10	8
No. of 6's	0	1	3	0	0
No. of 5's	2	5	3	3	4
No. of 4's	3	1	1	1	1
No. of 3's	1	1	1	0	0
No. of 2's	1	1	2	2	2
No. of 1's	4	2	2	0	1
Experiment 7					
No. of 7's	–	–	6	5	5
No. of 6's	–	–	1	0	1
No. of 5's	–	–	0	3	0
No. of 4's	–	–	2	1	4
No. of 3's	–	–	2	0	0
No. of 2's	–	–	0	0	1
No. of 1's	–	–	3	5	3

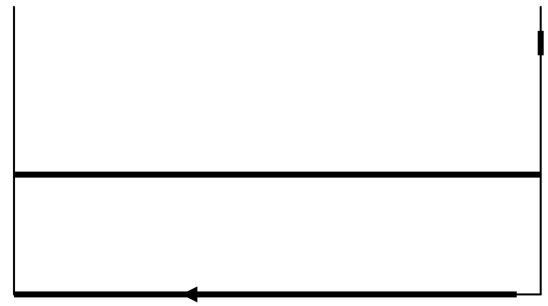

In the large-group minimum game, the all-Rabbit equilibrium is locally stable but the all-Stag equilibrium is locally unstable. Starting from all-Stag, any shock, however small, will make the population converge to all-Rabbit.

This makes the strong convergence to the equilibrium with lowest effort VHBB observed in the large-group minimum game plausible, but in this case the story is more complicated.

In the seven-effort large-group minimum game, if learning always makes subjects adjust their efforts toward the current value of the minimum, then the population converges to initial minimum without changing it. However, in VHBB's experiments the initial minimum was above one in five out of seven sessions, but it always converged quickly down to one.

Crawford (1995) shows that this happens because in the minimum game, random shocks (which represent subjects' inability to perfectly predict others' adjustments) are not neutral as they were in the median game: Instead they tend to make the minimum go down, to an extent that can be approximately quantified. As intuition suggests, the downward trend is stronger, the larger the group or the closer the order statistic (below the median) is to the minimum.

VHBB's 1993 design and results

VHBB's 1993 design was the same as their 1991 design, with repeated play of one of the 1991 median games, but with the right to play auctioned each period to the highest 9 bidders in a population of 18 (English clock auction, same price paid by all winning bidders).

The market-clearing price was publicly announced after each period's auction, the median was publicly announced after each period's play, and the structure was publicly announced at the start.

The stage game has a range of symmetric equilibria, in which all bid the payoff of some equilibrium of the median game and play that equilibrium, unless others bid differently.

In 8 of 8 trials, subjects quickly bid the price to a level that could only be recouped in the most efficient equilibrium and then converged to that equilibrium; the results give strong, precise selection among a range of equilibria.

Auctioning the right to play had a strong efficiency-enhancing effect via focusing subjects' beliefs on more efficient ways to coordinate—a new and potentially important mechanism by which competition promotes efficiency. Crawford and Broseta (1998 *AER*) show that this effect can be understood as following from "order statistic," "optimistic subjects," and "forward induction" intuitions: Auctioning the right to play a 9-person median game in a group of 18 effectively turns the game into a "75^{th} percentile" game, with an upward trend.

TABLE V

Distribution of Actions for Game Γ(9): EC Auction

	Period														
	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
Exp. 10															
Price	1.24	1.24	1.28	1.29	1.30	1.30	1.30	1.30	1.30	1.30	—	—	—	—	—
Undom. actions	≥6	≥6	7	7	7	7	7	7	7	7	—	—	—	—	—
# of 7s	7	8	9	8	9	9	9	9	9	9	—	—	—	—	—
# of 6s	2	1	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 5s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 4s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 3s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 2s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 1s	0	0	0	1	0	0	0	0	0	0	—	—	—	—	—
Median	7	7	7*	7	7*	7*	7*	7*	7*	7*	—	—	—	—	—
Exp. 11															
Price	1.00	1.20	1.29	1.30	1.29	1.30	1.29	1.29	1.30	1.30	—	—	—	—	—
Undom. actions	≥4	≥6	7	7	7	7	7	7	7	7	—	—	—	—	—
# of 7s	4	5	9	9	9	9	9	9	9	9	—	—	—	—	—
# of 6s	1	3	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 5s	2	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 4s	2	1	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 3s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 2s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
# of 1s	0	0	0	0	0	0	0	0	0	0	—	—	—	—	—
Median	6	7	7*	7*	7*	7*	7*	7*	7*	7*	—	—	—	—	—
Exp. 12															
Price	.95	1.04	1.08	1.10	1.15	1.20	1.25	1.25	1.30	1.30	1.30	1.30	1.30	1.30	1.30
Undom. actions	≥3	≥4	≥5	≥5	≥5	≥6	≥6	≥6	7	7	7	7	7	7	7
# of 7s	1	0	1	3	2	5	8	9	9	9	9	9	9	9	9
# of 6s	0	3	5	6	7	4	0	0	0	0	0	0	0	0	0
# of 5s	6	2	2	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	1	4	1	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	1	0	0	0	0	0	1	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	5	5	6	6	6	7	7	7	7*	7*	7*	7*	7*	7*	7*
Exp. 13															
Price	1.05	1.14	1.18	1.25	1.29	1.25	1.25	1.30	1.25	1.30	1.30	1.30	1.30	1.30	1.30
Undom. actions	≥4	≥5	≥6	≥6	7	≥6	≥6	7	≥6	7	7	7	7	7	7
# of 7s	2	2	4	6	9	9	9	9	9	9	9	9	9	9	9
# of 6s	1	6	5	3	0	0	0	0	0	0	0	0	0	0	0
# of 5s	3	1	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	3	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	5	6	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*
Exp. 14															
Price	1.05	1.15	1.27	1.25	1.25	1.30	1.30	1.25	1.30	1.30	1.25	1.30	1.30	1.30	1.30
Undom. actions	≥4	≥5	7	≥6	≥6	7	7	≥6	7	7	7	7	7	7	7
# of 7s	0	5	8	8	9	9	9	9	9	9	9	9	9	9	9
# of 6s	7	4	0	1	0	0	0	0	0	0	0	0	0	0	0
# of 5s	1	0	1	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	6	7	7	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*

Notes. * indicates mutual best response outcome. ___ Partitions actions into $FI(P)$ and the complement of $FI(P)$.

TABLE VI

Distribution of Actions for Game Γ(9): EC Auction and Experienced Subjects

	Period														
	16	17	18	19	20	21	22	23	24	25	26	27	28	29	30
Exp. 7 (M = 5)															
Price	1.09	1.09	1.10	1.19	1.29	1.29	1.30	1.29	1.30	1.30	1.30	1.29	1.30	1.25	1.29
Undom. actions	≥5	≥5	≥5	≥6	7	7	7	7	7	7	7	7	7	≥6	7
# of 7s	0	0	2	5	9	9	9	9	9	9	9	9	9	5	9
# of 6s	2	1	5	4	0	0	0	0	0	0	0	0	0	0	0
# of 5s	6	8	2	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	5	5	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*
Exp. 8 (M = 5)															
Price	1.09	1.25	1.28	1.29	1.30	1.29	1.30	1.30	1.29	1.30	1.29	1.30	1.29	1.30	1.30
Undom. actions	≥5	≥6	7	7	7	7	7	7	7	7	7	7	7	7	7
# of 7s	3	7	9	9	9	9	9	9	9	9	9	9	9	9	9
# of 6s	2	2	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 5s	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	2	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*
Exp. 9 (M = 6)															
Price	1.15	1.21	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29	1.29
Undom. actions	≥5	≥6	7	7	7	7	7	7	7	7	7	7	7	7	7
# of 7s	0	7	9	9	9	9	9	9	9	9	9	9	9	9	9
# of 6s	8	1	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 5s	0	1	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 4s	1	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 3s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 2s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
# of 1s	0	0	0	0	0	0	0	0	0	0	0	0	0	0	0
Median	6	7	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*	7*

Notes. * indicates mutual best response outcome. ____ Partitions actions into $FI(P)$ and the complement of $FI(P)$.

Extensive-form Games (possibly not covered in class, and optional for exam)

To really understand VHBB's 1993 results, we need to know more about how to analyze games with sequences of decisions and counter-decisions, or "extensive-form" games).

A *static* or *simultaneous-move* game has one stage, at which players make simultaneous decisions, like those discussed so far.

A *dynamic* or *extensive-form game* has some sequential decisions.

E.g. Ultimatum Contracting with two feasible contracts, X and Y:

R proposes X or Y to C, who must either accept (a) or reject (r).

If C accepts, the proposed contract is enforced.

If C rejects, the outcome is a third alternative, Z.

R prefers Y to X to Z, and C prefers X to Y to Z.

R's payoffs: u(Y) = 2, u(X) = 1, u(Z) = 0; C's: v(X) = 2, v(Y) = 1, v(Z) = 0.

The game actually depends on whether C can observe R's proposal before deciding whether to accept: With observable proposal it's dynamic; with unobservable proposal it's static.

We can represent either game by its *extensive form* or *game tree*, which shows its sequence of decisions, outcomes, and payoffs.

The order of the *decision nodes* must respect the timing of moves.

Each node belongs to an *information set* (represented by circles), the nodes the player whose decision it is cannot distinguish (and at which he must therefore make the same decision).

All such nodes must belong to same player and have same feasible decisions.

(A game of *perfect information* is one in which a player making a decision can always observe all previous decisions, so every information set contains one decision node, as in Ultimatum Contracting with Observable Proposal.)

For dynamic games it is important to distinguish *strategies* from *decisions* or *actions*. A *strategy* is a complete contingent plan that specifies a decision for each of a player's decision nodes and information sets (like a chess textbook, *not* a move).

In a static game a strategy reduces to a *decision* or *action*. (These definitions apply equally well to *mixed* or *pure* strategies. Specifying a strategy profile—one for each player—must determine an outcome (or probability distribution over outcomes).)

A player’s strategy (or decision) must be feasible independent of others’ strategies.

Players *must* be thought of as choosing strategies *simultaneously* (without observing others’ strategies) at the start of play. Rational, perfect foresight implies that simultaneous choice of strategies yields the same outcome as decision-making in “real time” (this is a testable prediction, which can fail, and does for some real people).

We need complete contingent plans (even for nodes ruled out by prior decisions) to evaluate consequences of alternative strategies, to formalize the idea that the predicted strategy choice is optimal. (0-probability events are endogenously determined by decisions, and so cannot be ignored here as they are in individual decisions.)

With the concept of strategy, we can also represent a game, static or dynamic, by the relationship between its strategy profiles and payoffs: *normal form*, *payoff function*, or (if 2 people) *payoff matrix*.

	a	**r**
X	**2** **1**	**0** **0**
Y	**1** **2**	**0** **0**

Ultimatum Contracting with Unobservable Proposal

	a, a	**a, r**	**r, a**	**r, r**
X	**2** **1**	**2** **1**	**0** **0**	**0** **0**
Y	**1** **2**	**0** **0**	**1** **2**	**0** **0**

Ultimatum Contracting with Observable Proposal

	a	r
X	1, 2	0, 0
Y	2, 1	0, 0

Ultimatum Contracting with Unobservable Proposal

	a, a	a, r	r, a	r, r
X	1, 2	1, 2	0, 0	0, 0
Y	2, 1	0, 0	2, 1	0, 0

Ultimatum Contracting with Observable Proposal

In Ultimatum Contracting, whether or not C can observe R's proposal, R has two pure strategies, "(propose) X" and "Y."

If C cannot observe R's proposal, C has two pure strategies, "a(ccept)" and "r(eject)".

If C can observe R's proposal, C has four pure strategies, "a (if X proposed), a (if Y proposed)", "a, r", "r, a", and "r, r".

C's additional information in Ultimatum Contracting with Observable Proposal "shows up" only in the form of extra strategies for C. But this can affect the outcome.

	a	r
X	2 1	0 0
Y	1 2	0 0

Ultimatum Contracting with Unobservable Proposal

	a, a	a, r	r, a	r, r
X	2 1	2 1	0 0	0 0
Y	1 2	0 0	1 2	0 0

Ultimatum Contracting with Observable Proposal

Suppose the payoffs are as above (R’s: u(Y) = 2, u(X) = 1, u(Z) = 0; C’s: v(X) = 2, v(Y) = 1, v(Z) = 0). Then C prefers either X or Y to Z, so C will accept either X or Y whether or not C can observe R’s proposal. R will then propose Y, his favorite contract, and C will accept.

Now suppose C’s payoffs are changed to: v(X) = 2, v(Y) = 0, v(Z) = 1, so that C now prefers X to Z, but not Y to Z (R’s payoffs are unchanged).

If C can observe R’s proposal, C will accept X but not Y. R will then propose X, which he prefers to Z, and C will accept.

But if C cannot observe R’s proposal, C must accept or reject what R proposes without regard to what it is. If C accepted, R would propose Y, which is worse for C than Z, so C will reject whatever R proposes.

Key theoretical notions in extensive-form games:

Subgame-perfectness or *perfection* requires that players' decisions are in equilibrium not just in the entire game but in every subgame. Like "dynamic consistency" of the solution concept.

E.g. in Ultimatum Contracting with Observable Proposal, the sensible "backward induction" outcome (Y; a, a) is a subgame-perfect equilibrium,

But there are other, not so sensible equilibria, like (X; a, r) or (Y; r, a), that are not subgame-perfect. E.g. (X; a, r) is not in equilibrium in the subgame following a proposal of X because if Row proposed Y, Column would want to accept it.

	a, a	**a, r**	**r, a**	**r, r**
X	**2** **1**	**2** **1**	**0** **0**	**0** **0**
Y	**1** **2**	**0** **0**	**1** **2**	**0** **0**

Ultimatum Contracting with Observable Proposal

Subgame-perfect equilibrium is closely related to iterated weak dominance in the normal form.

Forward induction requires that players draw certain inferences from previous decisions.

Suppose that before playing Battles of the Sexes, the man (but not the woman) has an outside option (poker with the boys?) that would yield him a payoff of 1.5: He can choose the outside option or he can choose to play the Battle of the Sexes game with the woman.

	Fights	**Ballet**
Fights	**1** **2**	**0** **0**
Ballet	**0** **0**	**2** **1**

Battle of the Sexes

If the man chooses to play the Battle of the Sexes game, the woman should infer that he is expecting to get a payoff at least as high as 1.5, and choose Fights. (This argument is related to iterated strict dominance in the normal form. It would be a dominated strategy for the man to give up his outside option and then not choose Fights.)

The subgame-perfect equilibrium in which the man does not exercise his option and then chooses Fights, while the woman chooses Fights, satisfies forward induction. In this equilibrium, the outside option coordinates expectations even though it is not exercised.

There is another subgame-perfect equilibrium in which the man exercises his option and the woman would choose Ballet if the man chose to play the Battle of the Sexes game. This one does not satisfy forward induction.

Experiments with Extensive-form Games

Beard and Beil (1994 *Management Science*) study Rosenthal's (1981 *Journal of Economic Theory*) game. The game tree gives player A the right to opt out (L) with payoffs x for A and y for B; or to give player B the move (R) with two choices, l with payoffs 0 for A, 0 for B; or r with payoffs z for A and w for B; $z > x$ and $w > v$ ($y >$ or $< w$):

	L	**r**
L	**x**, **y**	**x**, **y**
R	**0**, **v**	**z**, **w**

The unique subgame-perfect equilibrium is (R,r) (which uniquely survives iterated weak dominance), but A players who think B is not certain to play r are tempted by L; thus the game is a simple test for reliance on other's dominance.

Intuitively, A players should be more willing to play R when:

(H1) x is lower (R is less risky)

(H2) w – v is higher (B has more incentive to choose r), or

(H3) y is lower (B is less likely to resent A's choice of R and choose l), or w and v are higher (B is more likely to reciprocate A's choice of R by choosing r)

	l	**r**
L	**x**, **y**	**x**, **y**
R	**0**, **v**	**z**, **w**

Subgame-perfect equilibrium doesn't imply any (H1)-(H3), because the unique subgame-perfect equilibrium is always (R, r) provided only that z > x and w > v (y > or < w).

Beard and Beil used a series of treatments to test (H1)-(H3), holding the critical probability that B chooses r that makes A indifferent between L and R constant near one in most treatments (higher than the frequency with which subjects respect dominance, thus making A subjects not rely on dominance)

Treatments			
		Player *A* Chooses *R*	
Treatment	Player *A* plays *L*	Player *B* plays *l* (critical probability)	Player *B* plays *r*
1	(9.75, 3.00)	(3.00, 4.75) (96.4%)	(10.00, 5.00)
2	(9.00, 3.00)	(3.00, 4.75) (85.7%)	(10.00, 5.00)
3	(7.00, 3.00)	(3.00, 4.75) (57.1%)	(10.00, 5.00)
4	(9.75, 3.00)	(3.00, 3.00) (96.4%)	(10.00, 5.00)
5	(9.75, 6.00)	(3.00, 4.75) (96.4%)	(10.00, 5.00)
6	(9.75, 5.00)	(5.00, 9.75) (95.00%)	(10.00, 10.00)
7	(58.50, 18.00)	(18.00, 28.50) (96.4%)	(60.00, 30.00)

Treatments			
		Player *A* Chooses *R*	
Treatment	Player *A* plays *L*	Player *B* plays *l* (critical probability)	Player *B* plays *r*
1	(9.75, 3.00)	(3.00, 4.75) (96.4%)	(10.00, 5.00)
2	(9.00, 3.00)	(3.00, 4.75) (85.7%)	(10.00, 5.00)
3	(7.00, 3.00)	(3.00, 4.75) (57.1%)	(10.00, 5.00)
4	(9.75, 3.00)	(3.00, 3.00) (96.4%)	(10.00, 5.00)
5	(9.75, 6.00)	(3.00, 4.75) (96.4%)	(10.00, 5.00)
6	(9.75, 5.00)	(5.00, 9.75) (95.00%)	(10.00, 10.00)
7	(58.50, 18.00)	(18.00, 28.50) (96.4%)	(60.00, 30.00)

Test (H1) A players should be more willing to play R when x is lower (R is less risky) by comparing Treatments 1, 2, and 3 (x = \$9.75, x = \$9.00, x = \$7.00)

Test (H2) A players should be more willing to play R when w – v is higher (B has more incentive to choose r) by comparing Treatments 1 and 4 (\$0.25, \$2.00)

Test (H3) A players should be more willing to play R when y is lower (B is less likely to resent A's choice of R and choose l), or w and v are higher (B is more likely to reciprocate A's choice of R by choosing r) by comparing Treatments 1 and 5 (B's payoff from A's secure choice L goes from \$3 to \$6)

Results

			A chose *R*		
Treatment	# of pairs	*A* chose *L*	*B* chose *l*	*B* chose *r*	% secure by A
1	35	23	2	10	65.7%
2	31	20	0	11	64.5%
3	25	5	0	20	20.0%
4	32	15	0	17	46.9%
5	21	18	0	3	85.7%
6	26	8	0	18	30.7%
7	30	20	0	10	66.7%

97.8% of B subjects made choices that their maximized own money earnings, suggesting that almost all were self-interested and rational.

Despite the predictability of most subjects' decisions, A subjects opted out in surprisingly large numbers.

(H1)-(H3) were all correct: The rate of opting out varied across treatments in a coherent manner, suggesting that payoffs had a significant, intuitive effect on subjects' willingnesses to rely on the self-interested behavior of others.

Experience as a B player was associated with significantly greater willingness to rely on the other's maximization in the role of an A player.